

Using an AI Accelerator to Build a Core Competency in AI – Six Case Studies

Three failures and three successes in AI adoption


Stephen Smith

DOI: <https://doi.org/10.66241/iuker>



Stephen Smith

AI Transformations



Abstract

The CEOs of many large companies have recently been given a mandate by their Boards: “Get goin’ on AI!” The CEOs, in turn, typically respond in three distinct ways:

1. Call marketing and generate some AI hype.
2. Gather the senior team and commit to a five-year plan to do it right and start hiring AI PhDs.
3. Find a partner that has an AI competence and tell your Board that you’ve checked off that AI box.

All three of these strategies can work but could also have poor outcomes in terms of lost credibility, lost opportunity, and enormous expense. There is, however, another way to build a corporate competency in AI. We call it an AI Accelerator. The concept has been trialed and tested many times during the last two decades. In this paper I will share the best practices on how to create one at your company.

You've got an AI mandate from your Board. Now what?

I won't go into details here about why you should have an enterprise competence in AI. I'll just say that I've been working in artificial intelligence, machine learning, and data science since 1982 (yes, AI has been around a long time!) and the results that have come about recently are unprecedented. In the forty years that I've been paying attention, I've never seen anything like it.

I'll admit I didn't see the current innovations in generative AI coming and though I've done my research, I'm not entirely sure why it works as well as it does. (But I'm in good company, as Stephen Wolfram, a well-respected scientist and AI researcher, recently admitted at the 2024 TED AI conference that he felt similarly.)

AI is big and, if you're smart, you've either given yourself a mandate to become personally knowledgeable about AI or to build a team and a competence within your organization. Bravo! You've made the right decision. But now you almost certainly find yourself falling behind compared to your competition – and falling further behind because AI innovation is moving at such a rapid pace.

Given this increasing velocity of AI, you need to accelerate your efforts just to begin to keep up (let alone catch up). You'd like to have a corporate competence in AI today but building an AI team is difficult and expensive. And despite the fact that you may have been given a command to “make AI happen now!” there is no instantaneous way to do that. There are currently four typical paths that a company will take if they don't already have an AI competency and they need one yesterday:

1. **Hype** – This strategy involves getting your marketing team together and brainstorming about how to redefine your existing company and product as already AI driven. Then you pick a tagline and start the hype engine. The downside of this strategy is loss of credibility and lost time that could have been used to build a real AI competency.
2. **Build it well and slowly** – The downside of this approach is that it is slow and expensive. It also may not be what you need, given that innovation in the field of AI is moving so quickly. What you aim at today in AI, may be

the wrong thing tomorrow and you need a nimble organization that can react to change quickly.

3. **Partner** – The upside of this approach is that this can be a lower-cost and more rapid way to get your “AI checkbox” checked off. The downside is that you are investing in someone else to build out their AI competency. Worse yet, having a strong partner will also have the negative effective of sucking the incentive out of your team to build their own AI competency.
4. **Acquire** – If you have the money, you can always purchase an existing AI company and fold it (or just the team) into your company. This is usually the most expensive way to achieve competency and has the additional risk that your newly-acquired AI experts may decide to leave the company after the acquisition.

There is one other way. A way that leverages the existing core competencies at your company (even if you don't have any in AI). A way that starts small and minimizes risk. And a way that will dramatically accelerate your path towards seeing revenue impact from AI-enhanced product and features. All within 18 months of your starting.

I call this approach the AI Accelerator.

What is an AI Accelerator?

An AI Accelerator is an internal team that is tasked with becoming knowledgeable about AI research, testing and trying out the latest AI tools, investigating datasets, utilizing a “Jobs To Be Done” methodology for new product, and building an AI core competency that infuses all of your employees and products.

Case Studies of Building AI Accelerators

Throughout my career I have worked at a wide variety of companies that have implemented different strategies for building an AI competency. These varied from very large companies creating new corporate divisions of hardware, software, and applications by purchasing existing software and hardware expertise from a university, to companies that created a less than ten-person team focused on just a few new AI features that improved an existing product line.

The case studies are introduced in the format of their challenges, how they formed their AI Accelerators, and the outcome of those initiatives. Finally, the case studies are labeled as either successes or failures based on the long-term impact of the products that were built with AI. In fairness, there were really no losers. All the case studies were huge successes in terms of accelerating the creation of an AI competency whether or not they enjoyed business success in the end.

Case Study 1 - FAIL: Texas Instruments - AI and Semiconductors

In the 1980s AI was the hot topic in technology (very similar to the levels of interest in AI that we see today). One of the main reasons for this interest was that the United States had great fear that “Japan Inc.” was pulling significantly ahead of the U.S. in the development of AI applications due to the launch of Japan’s Fifth Generation Computer Systems initiative. Japan’s Fifth Generation initiative was a ten-year project to build a massively parallel computer with logic programming. The ultimate goal was to create an “artificial intelligence”.

To further stoke the flames of AI interest, Edward Feigenbaum wrote a book in the U.S. in 1983 titled the “The Fifth Generation” that sent the U.S. government, especially the Defense Advanced Research Projects Agency (DARPA), into a mad scramble to respond. Industry also responded and a group of US companies and universities got together and founded MCC (Microelectronics and Computer Consortium), in Austin Texas. Texas Instruments (TI), also in Austin Texas, caught the AI bug and decided to enter the AI race. TI’s only problem was that they had no hardware or software expertise in AI.

To combat this problem, TI purchased a license for a special purpose AI computer technology (a LISP computer) from the Massachusetts Institute of Technology (MIT) and also all rights to the operating system software and began to build applications on top of that hardware and software. Leveraging their expertise in semiconductors and computer hardware, Texas Instruments was able to create the TI Explorer AI computer. It was a cheaper and more robust version of the very powerful LISP computer that was designed at MIT. They trained their own employees by embedding them at MIT and then hiring recent AI graduates from top tier schools like MIT (I was one of those hires).

Since Texas Instruments was a technology company with a high density of high performing engineers and computer scientists, they were able to release the new AI hardware, software and applications within a few years. The AI computers were successfully deployed at MIT and government research labs, as well as at private companies. They were successfully used for applications like a scheduling system for the Hubble telescope and some limited business applications. But overall, this AI Accelerator that TI built could be considered to be a “fail” because TI was not able to produce applications with their AI that were used beyond research labs and by 1990 the TI Explorer AI computer was phased out.

Case Study 2 - SUCCESS: Thinking Machines - AI and Supercomputers

Thinking Machines was the leader in parallel supercomputing and the application of AI to parallel supercomputers in the 1990s. Based in Cambridge Massachusetts, it was closely associated with MIT and Harvard and was founded by Danny Hillis and Sheryl Handler from MIT. The company had a large number of computer science researchers with impressive capabilities in parallel processing. There was also a small team of about 10 scientists who were focused on AI. Thinking Machines invested in this smaller group to prove that AI algorithms could run on a massively parallel computer. The AI research was applied to everything from news story search and retrieval, to protein folding, to retail coupon distribution, and even to the completion of an unfinished fugue by Bach.

AI was an important application for Thinking Machines because AI’s need for large amounts of data was a good match for the parallel processor architecture of the supercomputer. To build the AI group, the company hired a professor from MIT and he recruited computer science professionals with graduate degrees in AI, machine learning, and statistics. The team was successful in delivering research results which helped to sell the computers to academic and government researchers.

Eventually, pure research efforts became less valued within the company due to their own success in showing that parallel computer architectures were a good fit for AI algorithms. At that point the company needed to productize some of its

research. They took the most promising research and turned it into three new products:

1. Text retrieval for Dow Jones (searching for articles from the Wall Street Journal)
2. Retail product prediction for American Express
3. The Darwin AI workbench

The products sold to Dow Jones and American Express were highly successful but it was difficult to find other customers who would also purchase the products. The reason was that Thinking Machines' supercomputer (the Connection Machine), while powerful, was exotic and required learning a new style of programming that was not quickly accepted by the mainstream.

Thinking Machines' AI product (Darwin) was a knitting together of all the existing AI research projects that had been built bespoke for the Connection Machine computer. Though it was the first real AI workbench it suffered the same fate as the other commercial products and had limited acceptance in the market. Darwin was eventually sold to Oracle as one of two major assets of the company when it went into Chapter 11 bankruptcy after the supercomputer hardware business declined with the end of the cold war in the mid-1990s.

Thinking machines had the advantage of having its founder and chief scientist come directly out of MIT's AI lab. In fact, one of the founding thought leaders of AI, Marvin Minsky, was the thesis advisor for the company's founder. This made it easier to recruit quality scientists to the team and inculcated a culture of learning about AI for all members of the company. This included weekly lectures from visiting researchers which were open to all employees and also a reading group specifically focused on Artificial Intelligence called the Seminar on Natural and Artificial Computation (SNAC) that kept researchers abreast of the latest research in the field.

This case study has been labeled a "success" because Thinking Machines created a world-class (though small) AI team that was very productive in generating research and building products. Although they were relatively unsuccessful in selling the AI products due to the exotic nature of the computer hardware, the AI

product Darwin was considered to be one of the most valuable assets of the company and was eventually sold to Oracle to form the basis of Oracle's data science and AI products.

Case Study 3 - FAIL: Dun & Bradstreet - B2C and B2B customer data analysis

In the late 1990s Dun & Bradstreet was a multibillion-dollar conglomerate that included several leading companies that processed data for their industries:

- IMS processed data for the pharmaceutical industry to calculate sales and marketing efficacy.
- DBIS gathered data on payment history from companies to compile risk profiles on each company.
- Moody's Investors Service provided credit ratings and risk analysis of securities.
- Nielsen CPG gathered point of sale retail data that they used to predict product purchase patterns and successful marketing campaigns.
- AC Nielsen gathered and reported on television viewing habits of U.S. consumers.

Each of these subsidiary companies contained a gold mine of data that could be used to feed AI algorithms to do things like:

- Predict risk
- Detect fraud
- Predict the optimal marketing offer to a consumer
- Select the next best offer for a consumer
- Predict customer churn (non-renewal)
- Optimize the cadence of marketing interventions like coupons or postcards

While Dun & Bradstreet companies were expert and well-respected in their industries and application areas, and knew their data as well, they were not up to speed on the latest breakthroughs in Artificial Intelligence such as neural networks, memory-based learning, fuzzy logic, genetic algorithms, and decision trees. Regardless of this lack of expertise, they had tremendous need for, and enjoyed a unique opportunity, to apply AI to their markets and franchise.

To satisfy this need D&B decided to create a centralized AI team that would be shared amongst all the divisions: The Data Intelligence Group (DIG). This was a fitting name as many of the AI algorithms of the 1990s were productized under the moniker of 'Data Mining' (AI was in the nadir of its hype cycle and was still a bit of a dirty word as it had not lived up to the hype promised in the 1980s).

The Data Intelligence Group (DIG) visited with the teams at each of the divisions and built AI models for a number of important application areas. As AI experts, they worked alongside the industry experts and statisticians of each customer and shared with them the knowledge of AI.

The good news is that DIG was able to spread the knowledge of AI to each division and provide models that greatly improved their products. The bad news was that they were so successful that eventually enough knowledge had been transferred so that each corporate division began to build their own teams and eventually desired to control their AI expertise within the division rather than share it across the other divisions and rely on corporate to provide the resource. This AI accelerator case study is labeled as a "failure" solely because the actual team and group was not long lived nor successful as a business unit. On the other hand, the parent company and the division were successful in embracing and deploying AI technology to achieve comparative advantage and gain market share.

Case Study 4 - SUCCESS: Pilot Software - Business Intelligence / Mobile Telecom

Pilot Software was a \$50 million provider of business intelligence software in the 1990s. It had developed and released the technologies called multidimensional databases and OnLine Analytical Processing (OLAP). Their challenges stemmed from their successes - they had successfully developed and proselytized the new

OLAP technology and now had many competitors. They needed to differentiate themselves in the marketplace.

Additionally, they had the problem that they were positioned as a workbench that was sold to IT departments or to analysts, meaning their margins were small and price tags were limited. They wished to move towards utilizing OLAP to create more valuable applications and complete solutions that easily showed business value to the purchaser. If they could sell complete solutions rather than tools, they could sell products at a higher price and greater margin. They looked to include predictive AI into some of these applications as a differentiator and value add. Their problem was that they had no existing internal expertise in AI.

Pilot's challenges were:

- **No AI Core Competency** - They had engineers who were expert in building tools for PCs and world-class data engineers in relational databases and multidimensional databases but no expertise in AI.
- **Differentiator needed** - They were in a commodity market / red-ocean battle.
- **Viewed as commodity** - They were viewed as a tool that was required but did not provide clear added value.
- **Technology mismatch** - At that time most AI algorithms were run on special purpose AI machines (LISP machines) or supercomputers and deployed with optimized vector databases. Pilot's software ran on standard PCs with relational databases and multidimensional database architectures.

Pilot's visionary CEO, Erik Kim, understood the value of AI and set a mandate to begin to build an AI competency. To get started quickly, Pilot hired an existing AI Accelerator team that was already expert in AI and its application to business problems. The team initially consisted of the following roles:

- **AI accelerator team leader** - a visionary who was expert in AI as well as business applications

- **AI Academic** – a researcher who knew the latest and greatest AI algorithms and kept abreast of the latest successful applications of AI
- **AI engineer** – a computer scientist and math expert who could quickly write complex code for a variety of computer systems from parallel supercomputers to PCs
- **Statistician** – an expert who provided rigor to the new AI algorithms by comparing them to more tried-and-true statistical approaches as part of testing and quality assurance
- **Technical writer** – a team member with skills to create press releases, documentation, and marketing materials that could communicate the strange new technology of AI to senior management and prospective customers in a way that they would understand

This team undertook the hard problem of converting existing AI algorithms such as CART decision trees (Classification and Regression Trees) to run on the required hardware and with existing database architectures. To enable the team to move quickly, an existing implementation of CART was discovered and the source code was purchased.

Initially the product was an additional AI tool (named the Discovery Server) added to Pilot's existing OLAP workbench. Unfortunately, most of Pilot's customers did not understand how to use AI, and though it was successfully sold as an add-on, it was not highly used. It was, however, successful in providing a very successful differentiator in the marketplace for Pilot's other products. Though Pilot was years ahead of any of its competitors in having AI integrated with BI, it was still considered to be a tool with a relatively low price point and low profit margins.

To address this remaining challenge the decision was made to repackage the Discovery Server as focused, high-value applications in different vertical industries. To do this an approach was implemented based on the methodology espoused in the book "Crossing the Chasm". This approach focused on one valuable vertical and one high-value application at a time rather than a generic tool or workbench. To accommodate this approach additional members were added to the team:

- **Vertical Partner Champion** - A partnership was developed with Lightbridge which was the leader in providing consumer credit scores to mobile phone providers. Lightbridge dedicated a senior business person to develop a product with Pilot.
- **Industry and Application Expert** - Lightbridge also dedicated a marketing person who was expert in the industry and understood the way that users did their job and the way their customer's customers (the mobile phone users) behaved.
- **Consortium Leaders** - A consortium of three leading mobile telecom providers were recruited and each was paid to participate. They were required to attend meetings and provide feedback as the product was developed. They benefited from learning about AI and its application to their industry and also had early access and first mover advantage when the product was released.
- **Product Manager** - Pilot provided an official product manager who functioned as 'mini-CEO' and was responsible for revenue and costs of the product. They were also expert in the user, pricing, competitive analysis, and overall marketing and sales of the product.
- **Product Engineer** - Pilot also folded the AI effort into the mainstream product development cycle now that AI was better understood and it became just one more technology that could be used (rather than something strange and exotic). This role made sure that rigorous software development and QA techniques were used to provide a robust and high-quality product.

As the consortium was formed it became clear that the problem of predicting customer churn (consumers not renewing their annual mobile phone contract) was a good match for predictive AI and Pilot's Discovery server. This customer problem had these key attributes:

1. It was a hard enough problem that it required the power of AI to solve it.
2. It was a good match for existing AI algorithms (no new research required).

3. It had great value to the customer and great ROI for the customer even at a high price point.
4. It provided an opportunity for Pilot and Lightbridge to sell a high-priced, high-margin, solution-based product.

Because the AI tool had already been built for Pilot and because no new AI algorithms were required, the product was produced in under twelve months. Paradoxically it had far fewer features than the Discovery Server, was much less powerful, yet sold for 10x the price. This implementation of an AI Accelerator was a big success and employed many of the recommended best practices.

Case Study 5 - FAIL: Optas - AI and Pharmaceutical DTC Marketing

Optas was a company that provided Direct To Consumer (DTC) outbound marketing services to pharmaceutical brand managers. For instance, they managed a +1 million customer database of seasonal allergy sufferers for Schering-Plough's Claritin pharmaceutical product and determined which consumers should receive email or printed mail offers and reminders for their allergy symptoms.

Normally direct to consumer marketing would have a host of problems that would be a good match for predictive AI (e.g. churn prediction, next best product, co-selling) but in this case, the industry was only ready for basic segment level marketing rather than more sophisticated personalized marketing.

Optas already had employees with high levels of AI, but despite several attempts to create sophisticated predictive products, they could not be sold. There was simply no appetite for them from Optas's customers. The industry application was simply too new and the industry was not yet ready for AI.

Case Study 6 - SUCCESS: InvestorFlow - AI and Alternative Investments

In May of 2022 (before ChatGPT appeared and the current AI explosion began) the CEO of InvestorFlow, Todd Glasson, decided to investigate how to apply AI to the alternative investments industry. The problem was that InvestorFlow was a small company of less than 500 people with no existing AI expertise aside from those engineers who had taken a class or two in graduate school. It was also a

relatively novel idea (just two years ago) to apply generative AI to an industry which, up to that point, was a relatively human relationship-driven industry.

The alternative investments industry includes venture capital (which invests in high-risk, potentially high payoff startups) and private equity (which invests in more mature companies). The business model can be simplified into two tasks for venture capitalists and private equity General Partners (GPs):

1. Finding promising companies to invest in
2. Soliciting money from investors

Previous research in applying predictive AI had been done in the alternative investments industry. Many companies and research groups were seeking to solve the deal sourcing problem of discovering new companies to invest in and predicting which ones would provide the best return on investment. This is a valuable and appropriate use for AI (specifically for machine learning and prediction algorithms) but it is a very challenging problem to solve as most of the companies being invested in by venture capitalists have minimal historical or public information on which to base an investment decision (or for an AI algorithm to build a predictive model).

The challenges for InvestorFlow (and for the industry as a whole) were:

- No existing culture of using AI within the industry
- Lack of internal AI expertise
- No historical evidence that building an AI team would be successful
- A current focus on extremely difficult predictive problems with weak signal and limited data

What InvestorFlow did have was:

- Deep technical and engineering talent
- An existing industry-leading product that was critically valued by loyal users
- Deep knowledge of the industry and the users

- Pragmatic, product-driven approach to applying new technologies

InvestorFlow set off with a core set of best practices that matched their strengths and was able to become highly competent in AI in a short period of time and begin releasing product features powered by AI within 18 months of the start of their AI accelerator project. Some of the key best practices and philosophies of InvestorFlow were as follows:

- **Tick-tock model** - InvestorFlow utilized the general idea of the Tick-tock model introduced by Intel in 2007. This philosophy balanced periods of experimentation and re-architecting with periods of fine-tuning and optimizing applications. For InvestorFlow this meant periods of pure research into AI methods, balanced against engineering sprints where the results of that research would be used to create new product features.
- **JTBD** - The “Jobs To Be Done” framework, developed by Clayton Christensen at Harvard Business School, was implemented at InvestorFlow in order to keep the customer needs foremost in the thinking of the researchers and developers. Integration of key users and user experts and use of the JTBD methodology guaranteed that research was impactful and provided near-term value to the product and customers.
- **Re-use, Borrow, or Partner** - The field of AI has exploded recently and many firms in the financial services industry have already built large AI teams of more than 100 researchers. This is also true for tools providers such as Salesforce, Amazon, and Google. To stand a chance at being competitive in AI, a philosophy of partnering and reuse was emphasized as much as possible. Basic research was limited to areas that were not already covered by other larger players.
- **No Hype** - The CEO and AI leaders had experience with AI in the past and recognized that there is little value in producing press releases with a tag line of “AI inside”. Instead, they nurtured a philosophy of “Pragmatic AI” and sprinkling in AI where it would help the user with existing tasks.

InvestorFlow is perhaps one of the best examples of building an AI competency quickly and effectively. I believe that their focus on ‘no hype’ and real product improvements allowed them to produce one of the most successful case studies.

The key best practices for an AI accelerator

There are many ways to build an AI accelerator but these are some of the keys that have been successful in a variety of different companies and industries.

1. **Designate an AI leader** - This person does not necessarily have to have a strong academic background in AI, but should be a good evangelist, willing to learn and disseminate knowledge, and take on the responsibility for building the AI Accelerator team and pursuing product. Customer focus and strong business sense are essential.
2. **Don't neglect education** - Begin a program of instilling interest in AI with all parties in the company regardless of whether they are on the AI team or not. This can include creating a weekly reading group, an AI Slack channel where relevant articles can be posted, summarizing research through internal white papers, and presentations, and company-wide 'lunch and learns'.
3. **Have your product be your goal** - Don't allow the hype of AI to distort your goal to be just "implementing AI". AI is a marvelous breakthrough but it is just one more technology that will be mastered. Keep your team and company dedicated to figuring out where AI can best be used in your existing and new products.
4. **Be customer driven** - Focusing on product means focusing on your customer. Don't use a technology (including AI) that you don't need, and don't build any product unless you know the pains and the jobs of your customers who will buy it. Look into the "Jobs To Be Done" framework and build a consortium of beachhead 'pre-customers' / early adopters who pay to be part of the consortium.
5. **Be prospect driven** - Be wary of only talking to your existing customers who already love you. Make sure you have prospects in your consortium who know nothing about you but represent the full extent of your available market. Use the technology to build solutions which enable you to gain market share.

6. **Don't abandon good product methodology** – We recommend the Tick-tock method of balancing chaotic experimentation with disciplined and focused product development. When you have done your customer needs research and have the glimmer of an idea for a product, don't start building the product. Instead, consider starting at the end of the product process and force yourself to write a new product press release first. How will the launch of the product be described in a one-page press release? Get that right first, then write a two-page product sheet, then the marketing requirements, then the technical requirements. This will be quick and fun to pretend you are at the end and imagine what the product will look like. Starting with drafting the press release goes a long way toward avoiding dead-end features and products that don't match customer needs.
7. **Fail fast in your experimentation** – It is important to be rigorous in your experimentation, but also to be sure to set yourself up to succeed fast or fail fast in one week's time. No multi-year research ideas. Define a hypothesis that you want to test with measurable criteria (e.g. speed, accuracy, etc.) Make sure the experiment is something that you are likely to complete in a week and move on to another experiment if you don't have a measurable conclusion. Embrace the mantra of “fail fast / fail forward” by doing a quick *post-mortem* and write-up of the experiment to understand why it failed. It is perfectly OK to ‘fail’ at the experiment, but it is not OK not to learn from the experiment. This philosophy will also instill a culture of fearlessness and curiosity in your AI accelerator team.
8. **Embrace a “partner vs. build” mentality** – Your weekly experiment sprints should also include evaluating potential AI partners (e.g. tools, cloud providers, specialty data, free open-source libraries). Always stay in the mindset of: “don't build anything you can buy, borrow, or partner with to achieve—unless it is a critical competitive differentiator”. Always consider free, open-source AI libraries and new key data sources as possible ‘partner’ opportunities as well.
9. **Eschew the hype** – AI has had its ups and downs before, so be humble and don't overpromise. That being said, you should also probably have a

quiet confidence that this time is different and AI is going to fundamentally change your market and perhaps the world (it really already has begun). But the current AI wave will still abide by Amara's law: "we overestimate the impact of technology in the short-term and we underestimate the effect in the long-run".

A special note on the "Jobs To Be Done" methodology

Many folks are not that familiar with the "Jobs To Be Done" methodology but it is worth looking up a few of the Harvard Business Review articles and familiarizing yourself with the topic. The main idea is a humbling one: Customers don't want your product. Instead, they just want to get a job done and your product is a convenient way to accomplish that.

It is a pretty simple concept, but when you embrace it, it really changes your focus away from technology, away from product, away from features or benefits, and even away from the customer. It focuses your development team solely on the job that needs to get done. And this change in philosophy really changes everything about the way that you work and prioritize efforts.

So, you might then ask: Why is the "Jobs To Be Done" methodology particularly important for developing AI products?

The reason is that AI is a technology that is uniquely susceptible to the "build it and they will come..." mentality, which is a technology-driven philosophy and one that Steve Blank (the founder of the lean startup movement) calls "the bane of all product development."

You might then also ask: Why is the "build it and they will come" mentality particularly dangerous and alluring for AI? Here are two key reasons:

- Humans have an unarticulated, basic hope that AI can do everything all by itself (it is intelligent after all...right?), and if you just unleash the technology then good things will happen automatically (like a self-driving car).
- The belief that AI is magic, and anything you do will be an instant success.

Also, because AI simulates things that people can do, it makes sense that it could be applied to anything relating to a user interface. So, a developer might be tempted to sprinkle AI liberally everywhere within the user experience.

Just because AI can be added in doesn't mean that it should be, or that it will deliver value. It really needs to be measured and applied judiciously. Generative AI is powerful but also it can be quite creative and sometimes give the wrong answer. Knowing the job that your user needs to get done can help to guide you in just how much AI is sprinkled and where.

Who should be on the AI Accelerator team?

The most important decision in accelerating your AI initiatives is how to create the right, most productive team. Whom to recruit, and when? And how many? And critically important: how can you build an AI team without having to hire hundreds of people with PhDs in AI?

Here is a shopping list of the key personnel you need to have either full-time or at least involved with your AI Accelerator team. Almost all of them might already be working at your company.

- **Executive champion** - A C-level executive who knows in their gut that AI is going to be big but has no idea how to pull it off.
- **AI accelerator leader** - A veteran who has done this before, knows the AI technology, and has proven they can build product with it. This will probably be your one expensive external hire.
- **Product manager** - This role is responsible for the costs, revenue, and profitability of the products produced by the AI Accelerator. They are also ultimately responsible for competitive analysis and market understanding. They should have missionary sales skills.
- **Industry insider** - This role combines some technical understanding and interest in AI with deep understanding of the industry, user, and their jobs to be done. This person should have good selling skills, as they should also be responsible for creating and managing the profitability of the customer consortiums.

- **AI researcher** – This person can be found on your existing research teams. They are super smart with a degree in computer science or math and took a few AI courses in college. Most importantly, they have a desire to learn about AI and make it an important part of their future skillset.
- **Veteran engineering leader** – This is the person who makes sure the logistics of building a product happen according to existing corporate best practices. They know the company’s products well and have successfully released products and features before.
- **Data expert** – This is the person who will understand the complexity of such things as vector databases, graph databases, cloud databases, and new data sources that could be used. Because success in AI is heavily dependent on good data that is well managed, these are really skills that all technical team members should possess to some degree.
- **UX leader** – This is a normal requirement for all product teams, but is particularly important, since AI will have an outsized impact on the user interface and user experience.
- **Writer/editor** – This role is responsible for driving and editing the creation of blogs, press releases, and white papers. They play a critical role in building an understanding of AI within the company as well as with the company’s customers and prospects.
- **Marketing and sales** – Marketing and sales tend to get very excited about AI because it is cool and new; they can also be the worst offenders when it comes to overhyping any AI initiative. Stay humble, ignore the hype, and judiciously include marketing as you define your company positioning on AI. When you have product ideas, begin to work closely with a few key salespeople to help build your customer consortium. Your best salespeople will know their customers and the market well.
- **The rest of the company** – One of the goals of the AI Accelerator, besides building new product, is building a company-wide competence and interest in AI. So, consider all members of the company part of the team in some small way.

Where you could be four years from now...

The nice thing about deploying an AI Accelerator is that ... well ... it is an accelerator. It helps you to get wherever you want to go faster. For some companies, an AI Accelerator means building out large and deep AI-specific research and product teams. For others, it means building a credible AI competency with very small, cost-efficient teams, built from existing personnel.

Whichever way you do it, if you use the best practices detailed here to build an AI accelerator for your company, you will find that within a year you will be keeping up with the top competitors in your industry in regards to AI. Within two years, you'll find that you are one of the top companies in AI for your industry. And within four years, you will have quickly, cheaply, safely, and effectively built highly profitable and new revenue-generating products based on AI.

About the Author



Stephen Smith

Senior AI Consultant at AI Transformations

Stephen Smith is Chief Executive Officer of G7 Research, where he leads initiatives that translate advances in artificial intelligence into actionable insight for business and organizational decision-making. His work emphasizes the strategic and operational implications of AI, including applied analytics, AI-enabled innovation, and the governance considerations required for responsible deployment at scale. In his executive role, he guides multidisciplinary teams that produce research, thought leadership, and evidence-based frameworks for evaluating and implementing emerging AI capabilities. He contributes to scholarly and practitioner-facing discourse on the intersection of technology, management, and organizational performance, with interests spanning AI adoption, data-driven strategy, and risk-aware AI management. Smith is also engaged in professional knowledge-sharing through publications and presentations that connect academic rigor with real-world implementation challenges. Through his leadership at G7 Research, he advances collaborations that strengthen the bridge between research communities and industry stakeholders navigating rapid technological change.

[LinkedIn](#) · [Google Scholar](#) · [Website](#)

Copyright Notice

All articles are published in the Journal of Business and Artificial Intelligence under the Creative Commons 'CC BY' ("Gold Open Access") license, where authors retain the copyright of their articles. The author grants JBAI a license to publish the article under a Creative Commons 'CC BY' license, which allows the work to be freely accessed, shared, and used under certain conditions.

About the Journal

The Journal of Business and Artificial Intelligence (ISSN: 2995-5971) is the leading publication at the nexus of artificial intelligence (AI) and business practices. Our primary goal is to serve as a premier forum for the dissemination of practical, case-study-based insights into how AI can be effectively applied to various business problems. The journal focuses on a wide array of topics, including product development, market research, discovery, sales & marketing, compliance, and manufacturing & supply chain. By providing in-depth analyses and showcasing innovative applications of AI, we seek to guide businesses in harnessing AI's potential to optimize their operations and strategies.