CompTIA®

Building AI Strategy

Introduction

The history of innovation has a short list of defining moments, inventions that triggered massive disruption to the existing way of life and drove waves of progress for years to come. The wheel. The printing press. The engine. The internet. Each of these breakthroughs started a revolution that changed the shape of the world. Following the launch of ChatGPT in late 2022, many experts are predicting that AI will be the next inflection point in technological achievement.

Without a doubt, AI has tremendous potential. Modern tools are closer than ever to passing the Turing Test, producing results indistinguishable from human output. Images, text and video can be created with simple descriptions, and groundbreaking analysis is being done on enormous sets of data. Given that these outcomes are already coming in early stages of adoption, there's good reason to expect that later stages will have profound impact.

It's not pure upside, though. There are questions to be answered before AI meaningfully transforms business workflow. These questions center on not only implementation details as businesses integrate the technology into their workflow, but also existential dilemmas raised by countless sci-fi tropes. How will AI fit into existing architecture? What new vulnerabilities need to be considered? And, ultimately, what does it mean for humans and computers to co-exist?

Fully answering these questions will be the work of years, if not decades, and there will be many unforeseen obstacles along the way. However, CompTIA's early research in the field, coupled with extensive experience observing the technology industry, helps present some initial direction. AI's place in history will certainly lie somewhere between the nirvana projected by optimists and the apocalypse expected by doomsayers. Building a balanced strategy will help organizations be ready when the future arrives.

Understanding the Meaning of AI

The first step in any implementation strategy is grasping a definition of the thing being implemented. In the case of AI, there are two layers of definition. In the broadest sense, most people understand the basic concept of AI: computers that think like humans. Beyond that simple description, the term "AI" has meant many things throughout the different eras of computing.
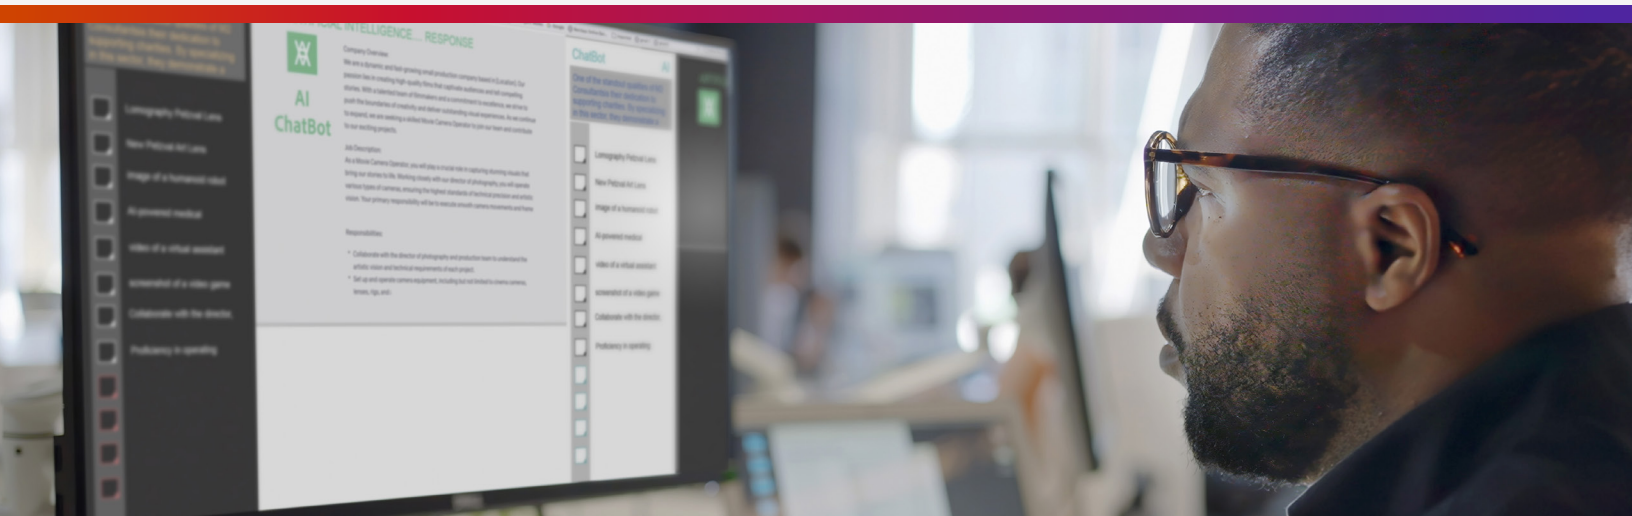
The previously mentioned Turing Test dates back to the 1950s, when Alan Turing realized the implications of computers representing information in binary form with ones and zeroes. If all information could be coded this way, and if decision-making processes could be coded as software algorithms, then machines could eventually mimic human intelligence.
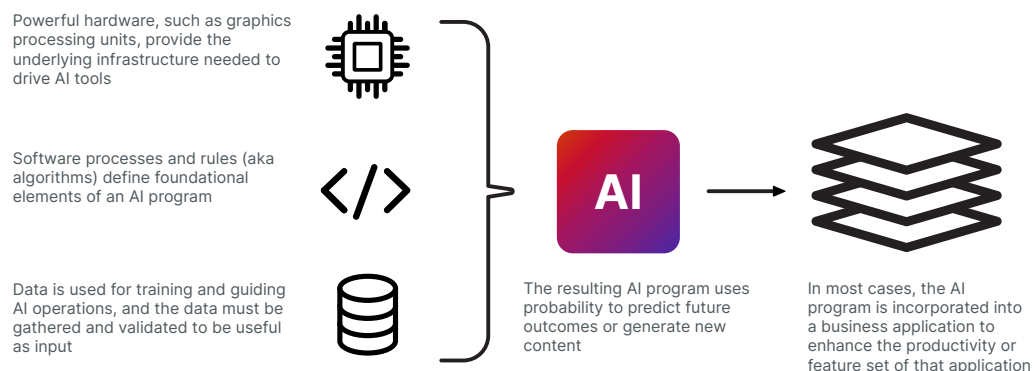
The evolution of AI has been defined by progress along these two vectors: development of the datasets and advancements in algorithms. One of the earliest AI efforts was a checkers-playing program written by Arthur Samuel in the mid-1950s. Here, the dataset was the rules of the game and the algorithms were some of the first machine learning practices, including a search tree of potential board positions and a memory of how those positions were tied to game results.

Samuel's program and other early efforts were limited in scope because of limitations on computing resources. As resources expanded, AI efforts grew more complicated. Margaret Masterman's semantic nets, Joseph Weizenbaum's ELIZA, and Richard Greenblatt's chess program were all early examples of advances in the field, even though resource availability still dictated specific use cases rather than general applications.

At various points, ceilings in compute resources caused delays in AI progress, sometimes called "AI winters." While progress has stalled, imagination has not. Science fiction stories and films have shaped public perception of what machines might be capable of once the ceiling is removed on computing and materials engineering.

The arrival of the internet signaled the beginning of virtually unlimited compute resources. With the world's information easily accessible, a program like IBM's Watson could successfully handle general trivia on the game show Jeopardy! At the same time, deep neural networks leveraging cloud computing gained the ability to perform highly advanced calculations, leading to the development of Google's AlphaGo to play the complex game of Go and Google Duplex to simulate natural conversations over the phone.

AI builds on several technology components to provide new function in business applications

Powerful hardware, such as graphics processing units, provide the underlying infrastructure needed to drive AI tools

Software processes and rules (aka algorithms) define foundational elements of an AI program

Data is used for training and guiding AI operations, and the data must be gathered and validated to be useful as input

AI

The resulting AI program uses probability to predict future outcomes or generate new content

In most cases, the AI program is incorporated into a business application to enhance the productivity or feature set of that application

Today, seventy years after the theory was proposed, the elements necessary for the Turing Test are largely in place. AI has broken into two main camps: predictive AI, which searches for patterns in historical data to predict future outcomes, and the latest breakthrough of generative AI (genAI), which uses transformer algorithms and large language models (LLMs) to produce output based on a deep knowledge of similar solutions.

Given the extensive history and the many different iterations, the exact definition of AI has become somewhat murky. The most practical definition is that AI is a discipline within software development, where the output relies on statistics and probability as opposed to a deterministic set of rules. Rather than being a standalone product, AI is best thought of as an enabling technology, bringing new features and functionality to both software applications and computing hardware. AI itself has many sub-domains, such as machine learning and natural language processing, which are leveraged to deliver either insightful predictions or unique content.

That definition helps explain some of the largest issues complicating AI adoption. The reliance on probability drives a need for human interaction; while the probability of a correct answer is often very high, the possibility of a wild error always exists. The hierarchy of an AI-enabled application complicates use cases; when companies claim to be using AI for business activities, are they using generative AI or neural networks? To what extent do they even know? At a high level, the details may not matter when organizations are weighing benefits and challenges. The picture changes slightly once it becomes time to implement a solution.

(Re)setting Expectations

As AI tools have grown more powerful, the prospects for extracting business value have grown as well. With any hype cycle, excitement builds as professionals and pundits imagine the possibilities; that excitement eventually fades and levels off as reality sets in. Hype cycles have become more complicated as consumer technology has become a focal point for innovation. The features found in consumer hardware and applications do not always translate seamlessly into enterprise use.

The recent wave of generative AI will likely follow this established pattern. The initial genAI offerings exploded in popularity for two reasons. First, the underlying technology could understand complex natural language inputs or ingest large amounts of data in order to have context for a query. Second, the tools produced instant results that delivered a relatively high level of quality and accuracy.

The output generated has been the primary driver of expectations. If AI can produce high quality content, how does that change (or replace) the output of existing workers? CompTIA's most recent data on AI usage shows a tight grouping of activities where companies are exploring AI, with content-related efforts such as chatbots, software development and content creation undoubtedly getting a big boost from generative AI.

Many corporate technology initiatives are incorporating AI in some way

Initiative	Percentage
Automation	67%
Data analysis	63%
Cybersecurity	61%
Chatbot	57%
Software development	57%
Content creation	57%
Financial modeling	54%

Source: CompTIA 2024 AI Progress Report │ n=511 North American tech and business professionals

The long-term benefits, though, might be tied more to the way the technology handles input. Even with the excitement around early genAI tools, the top applications for AI are consistent with CompTIA's earlier research around AI. Automation, data analysis and cybersecurity are tasks that do not typically have content-related output. Instead, they are related to daily operations. In these cases, AI can understand a wide variety of inputs related to the problem at hand, then provide various forms of assistance, such as direct automation of certain tasks, suggestions of patterns found in data, or predictions of cyber attacks.

The common thread through the planned uses of AI, similar to the common thread through all digital transformation efforts, is productivity. According to CompTIA's IT Industry Outlook 2024, the second-highest workforce priority (and the highest priority among companies with 500+ employees) is ensuring maximum productivity. While this objective may seem obvious, it is one that most organizations measure indirectly. Revenue growth or customer satisfaction are often the primary metrics, and productivity may be measured as a function of the cost required to meet those goals.

In this context, AI is viewed in the same light as most strategic technology investments. The traditional view of technology, which was highly tactical, treated the IT function as a cost center and drove toward providing required levels of technical capacity with constant or reduced investments. AI and other emerging technologies require a different mindset— being able to justify investments with expected returns. These types of calculations are new for most IT leaders, but those who can measure productivity gains following technology investment or integration will deliver the most value.

Looking to the future, the productivity impact of AI will be a complicated topic. Not only will AI make individual workers more productive, but in some cases AI tools will prove to be a lower-cost option for certain tasks. It is difficult to tell what the net effect will be. Historically, technological innovations have led to the elimination of certain job roles in the short run, but the long run has seen job creation. For all the concern about robot overlords, the safe bet is that AI will follow precedent.
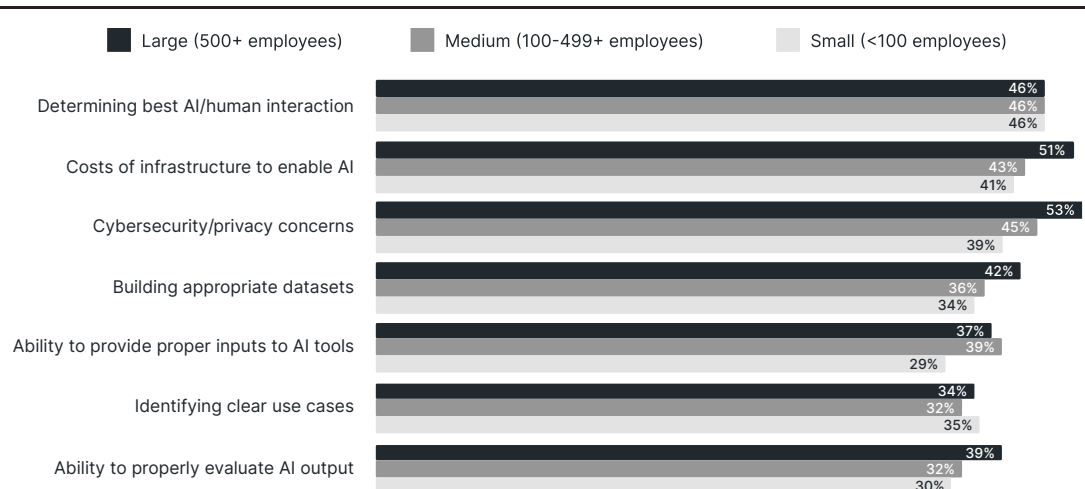
Obstacles to AI Success

Along with a basic reality check, there are other more discrete challenges that will get in the way of widespread AI adoption. These challenges come from many different directions, from the technical to the operational to the societal. Most organizations will not have to solve each challenge directly, but the nature of the solutions will be a key ingredient to AI implementation.
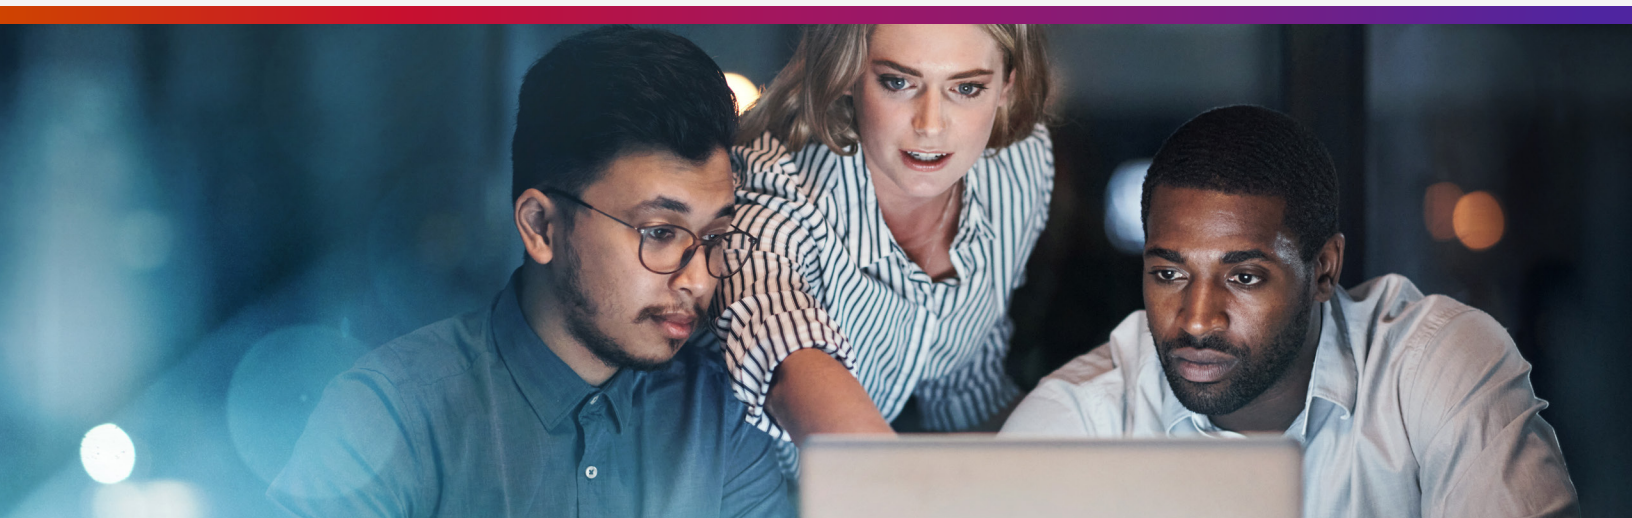
For the businesses surveyed in CompTIA's latest research, the most anticipated challenge is somewhat of a surprise. Rather than worrying about technical details or workflow transformation, companies are concerned about finding the best balance between their employees and new AI tools. Workforce reduction may be a part of the solution for some

Finding balance between humans and AI leads list of challenges across all company sizes

Legend: Large (500+ employees) | Medium (100-499+ employees) | Small (<100 employees)

Challenge	Large	Medium	Small
Determining best AI/human interaction	46%	46%	46%
Costs of infrastructure to enable AI	51%	43%	41%
Cybersecurity/privacy concerns	53%	45%	39%
Building appropriate datasets	42%	36%	34%
Ability to provide proper inputs to AI tools	37%	39%	29%
Identifying clear use cases	34%	32%	35%
Ability to properly evaluate AI output	39%	32%	30%

Source: CompTIA 2024 AI Progress Report | n=514 North American tech and business professionals

firms, but previous research along with the history of tech adoption suggests that reduction will not be the first lever that companies pull. Instead, there seems to be genuine interest in maximizing productivity, making the existing workforce more efficient through strategic AI applications and targeted upskilling.

Finding the human/AI balance will be a long-term endeavor, but the next set of challenges are more immediate. These technical issues will dictate the structure of AI within an organization. While the infrastructure costs of building and maintaining an LLM have been widely reported, most companies will not actually need to answer this question directly. By using cloud solutions or embedded functionality, many firms will enjoy AI benefits without developing proprietary systems.

Cybersecurity and data concerns are more pressing. It is assuring to see cybersecurity ranked high on the list of challenges. This suggests that businesses have learned from their technology initiatives over the past decade, when many firms leapt into new architecture or applications under the assumption that existing cybersecurity strategy would be sufficient, only to be unpleasantly surprised. Data is a different story. The level of concern here is perhaps too low given the importance of data for proper AI training. Previous hype cycles around Big Data and data analytics have shown that many firms have work to do in building foundational datasets and data management practices, and having clean data as a prerequisite for AI simply makes this challenge more pressing.

From an organizational perspective, the non-technical challenges are similar to other emerging tech adoption initiatives. Finding the proper use cases, especially for technology that explodes in the consumer space, is not always straightforward. The need to evaluate output hints at the need to modify workflow, since individuals may be presented with AI insights that could be revolutionary or could be dangerously inaccurate. Many promising technologies have fallen by the wayside as companies have struggled to change behavior. With AI, there will be plenty of individual use mirroring the consumer applications, but broad enterprise transformation will require workflow evolution.

Perhaps the most significant challenges are those around how AI will be governed. These challenges tend to lie outside of the purview of most organizations, but AI regulation can be added to the rapidly growing list of compliance guidelines for digital business. Deepfakes, copyright claims and biased algorithms are all problems being examined from a regulatory viewpoint. The decisions made by legislative bodies around the world will drive the development of safe AI solutions and the standards for content transparency.

Today's AI tools all still fall into the category of narrow AI, focused on a limited range of tasks. As we get closer and closer to general AI (also known as artificial general intelligence or AGI) and the ability to perform a wide range of functions, these universal challenges all point to the ethical question of control. Science fiction has conditioned us to be fearful of self-aware AI, but the stuff of stories could one day become reality. Businesses can build implementation strategies, and governments can set regulations, but ultimately the comfort level around technology is set by the people using the tools. Mobile devices and social media have already reshaped society, and time will tell how much further shaping will be done by AI.
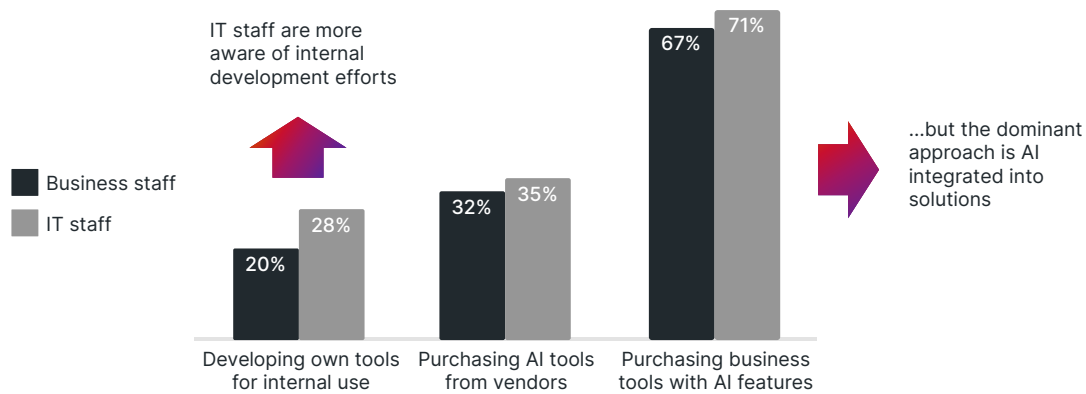
Turning AI Vision into Reality

For decades, IT was a product-centric function. Mainframes, servers, networks and endpoint devices were the critical components of an architecture that supported and accelerated business activity. The need to distribute and support products led to the creation of the IT channel ecosystem, and IT budgets were driven in large part by product procurement and maintenance.

This history still influences much of the thinking around enterprise technology. As emerging technology became a focal point at the end of the last decade, one of the difficulties for CIOs was positioning newer trends as components of a holistic solution rather than standalone solutions. Conversations that referenced concepts such as a "blockchain market" or "5G revenue" were grounded in a viewpoint of these technologies as products.

AI suffers from the same fate. Along with the label being applied to a wide range of software techniques, it is also being used as a shorthand for new applications. ChatGPT is called AI, but it is better described as a search engine or a content creation tool. The second description, especially in the specific incarnation, may be a novel product, but AI is still a component of that product rather than the product itself.

IT staff are more aware of internal development efforts

...but the dominant approach is AI integrated into solutions

Business staff
IT staff

	Developing own tools for internal use	Purchasing AI tools from vendors	Purchasing business tools with AI features
Business staff	20%	32%	67%
IT staff	28%	35%	71%

Source: CompTIA 2024 AI Progress Report | n=492 North American tech and business professionals

The distinction may seem minor, but business plans around implementation approaches highlight its importance. Few companies are planning to develop their own AI algorithms and tools internally, even though many firms house internal software and web development teams. Some of this reflects the cost implications, as mentioned previously. A larger driver is the complexity involved in developing core functions of business applications.

A slightly larger segment of companies expects to purchase "AI tools," by which the purchasers are likely imagining new product categories built on top of AI capabilities. It remains to be seen how viable these categories will be compared to existing business solutions. The majority of businesses expect to invest in these existing tools—such as customer relationship management, business productivity suites and HR systems—that have AI embedded to provide new features. These tools are already integral parts of corporate workflow, and AI will become a powerful new part of a complex solution stack.

Regardless of the implementation specifics, companies will have to establish or augment processes around evaluating new technology. This has been one of the main takeaways from the shift into emerging technologies: These new elements come with new issues around cybersecurity vulnerabilities, workflow management and infrastructure requirements. In order to properly understand these issues and accurately budget cost and effort for implementation, businesses have started forming teams to evaluate the pros and cons. These teams are often cross-functional, providing business perspectives on objectives along with technical perspectives on integration.

For AI, these evaluation efforts will need to address the many questions that come with the territory. First, what data will be needed to train the AI? This may require locating data across the organization and collecting it together. Second, what cybersecurity considerations come from the new tool? Along with potential vulnerabilities that could come from a new vendor, there are also privacy considerations depending on the data being used. Finally, how will the AI output be used? In some cases, this may be a straightforward answer. In many others, there will need to be decisions around evaluating the output and plugging it into a new and improved workflow.

The effort involved in evaluating and implementing AI points back to the primary challenge of optimizing human/AI interactions. Modern AI technology, as advanced as it may be, is still neither self-aware nor simple to implement. Complex solution stacks require thoughtful decisions in order to be modified and maintained. While the overall effect of AI on the workforce is unclear, there is no question that AI is driving demand for new skills.

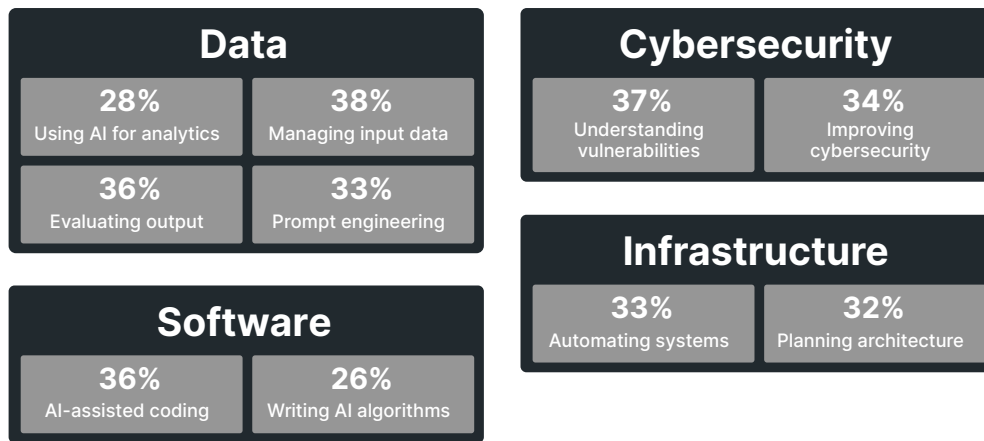# Growing Skills for an AI Future

The labor market was already in a state of turbulence when generative AI entered the picture. Historically low unemployment combined with ongoing efforts in digital transformation to create imbalance between supply and demand, and organizations had started shifting toward a skills-based approach to hiring and workforce development. AI, as with most new technologies, added questions around the creation of new skills or job roles that would thrive in the future.

Quality and quantity are both great unknowns when assessing the future potential of emerging skills. AI's parent discipline of software engineering provides a simple example. In 1900, the role of software developer did not exist. By the 1950s, companies could imagine software developers existing, but in very limited circumstances such as massive corporations or software vendors. By the 2000s, advances in open source, microservices and web development had driven software developers to be the dominant role within the technology workforce.

CompTIA's analysis of Lightcast job posting data found that postings for AI-related jobs and postings featuring AI-related skills climbed 16% in the six months between November 2023 and May 2024—a significant increase in a time period where overall postings for technology jobs remained steady. Reflecting the expectation that AI will become part of the overall technology stack, demand for AI-related skills is rising in every area of IT operations.

AI skill demands span all areas of IT operations

Data

28% Using AI for analytics	38% Managing input data
36% Evaluating output	33% Prompt engineering

Cybersecurity

37% Understanding vulnerabilities	34% Improving cybersecurity

Infrastructure

33% Automating systems	32% Planning architecture

Software

36% AI-assisted coding	26% Writing AI algorithms

Source: CompTIA 2024 AI Progress Report │ n=520 North American tech and business professionals
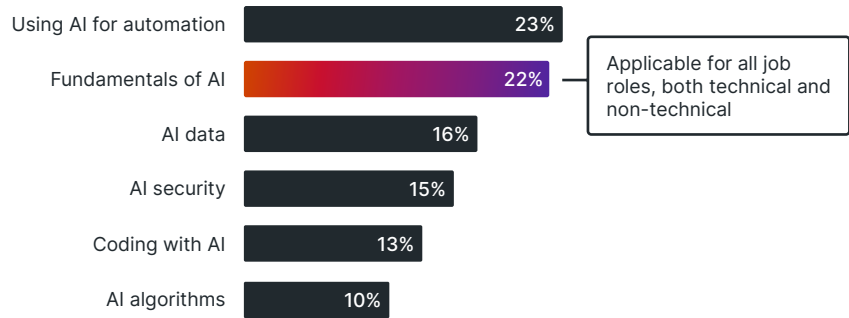
The field of data management and analytics currently generates the most AI interest. As cloud computing has democratized infrastructure and software, data has become ground zero for building competitive differentiation. The amount of churn in this rapidly changing discipline makes it an ideal candidate for AI-assisted skills, such as data analysis, output evaluation and proficiency in manipulating inputs.

There is also high interest around cybersecurity. From an evaluation standpoint, companies need to understand the new vulnerabilities being created as they build or incorporate AI tools. From an operational standpoint, companies plan to use AI to improve their existing cybersecurity activities. These activities can include routine cybersecurity tasks, analysis of cybersecurity incidents or creation of tests to proactively assess defenses.

While directly writing AI algorithms may be a niche activity for the time being, there is greater interest in using AI to assist with standard coding practices. This is one of the most common use cases that has emerged from the first wave of genAI products. Given that many companies struggle with certain aspects of their existing DevOps flow (such as version control or secure coding), the acceleration that comes from AI will definitely be something that requires updated skills.

Finally, the area of infrastructure, as well-established as it is, is still expected to develop new AI aptitude. The practice of systems operations (SysOps) has long been working toward greater automation, and AI will speed up those efforts while also offering new techniques around machine-driven decision making. Especially for those organizations that build and manage proprietary AI systems, there will be additional requirements around planning architecture to support intensive compute activity.

Training around AI will likely be incorporated into broader digital skill training

Using AI for automation	23%
Fundamentals of AI	22% — Applicable for all job roles, both technical and non-technical
AI data	16%
AI security	15%
Coding with AI	13%
AI algorithms	10%

Source: CompTIA 2024 AI Progress Report | n=514 North American tech and business professionals

The skill demands extend beyond core technology professionals, though. On the corporate list of priorities for workforce training, AI fundamentals ranks second, trailing only automation. Similarly, CompTIA's Job Seeker Trends published in January 2024 found that AI fundamentals ranks second on the list of perceived digital skills needs, suggesting the impending addition of AI acumen on resumes. From an HR perspective, CompTIA's Workforce and Learning Trends 2024 report found that 62% of HR professionals have positive views around the potential of AI. Clearly, AI will feature heavily in growing demands around both basic digital literacy and more advanced digital fluency.

The details around AI skills, not to mention the way they weave into a worker's overall skill set, will be determined by the market as businesses construct the most effective ways to structure job roles. Even as the market is churning, though, the words of Louis Pasteur ring true: Fortune favors the prepared mind. Building awareness and acumen around AI is one of the best preventive measures against job disruption. In the end, it may not be a machine that takes over someone's job; it may be another individual who has invested more in building skills.

Looking Ahead

Does AI signal the next stage of human achievement? Or, at a minimum, does it signal the next stage of business development? Technology trends come and go, and the long-term impact is difficult to predict, but it seems likely that AI will be a major milestone, either in this current iteration or a new version in the near future. As organizations continue expanding AI skill demands across a range of job postings, CompTIA will continue exploring the implications for both employers and employees.

Job posting data demonstrates the wide range of skills needed

41%
Data/prompts

26%
SysOps/Architecture

11%
Coding

15%
Cybersecurity

7%
Other
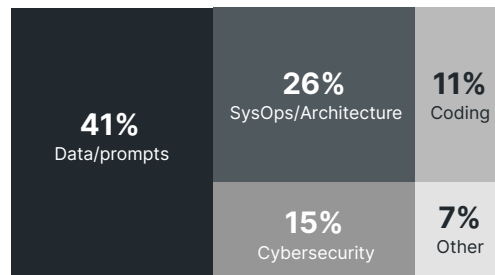
Source: CompTIA analysis of Lightcast job posting data

At the end of the day, no amount of debate over pros and cons will impede AI development. The best path forward is to understand the benefits, build safeguards against the challenges, and foster environments where innovation and skill-building can thrive.

CompTIA