

93/720-84
84/038-38
30/720-05
28/502-28

CODE_01: 78882 - 9189200 - 91878
CODE_02: 88472 - 72837451 - 83821
CODE_03: 88204 - 83839-103 - 95188
CODE_04: 84832 - 993-10003 - 84716
CODE_05: 96802 - 903-89510 - 17686

UNIT 01 - 84203892-8 953
UNIT 02 - 7863-87024 346
UNIT 03 - 0390233-871 426



Artificial Intelligence

Microsoft Practice
Development
Playbook



aka.ms/practiceplaybooks

About this Playbook

This playbook is intended for the business and technical leadership for new and existing Microsoft partners that are adding or expanding an Artificial Intelligence (AI)-focused practice to their business.

Objectives

The goal of this playbook is to help you accelerate or optimize your AI focused practice and understand how to define your practice strategy, hire and train resources, go to market, and optimize and grow your practice. We did not re-write the existing body of detailed guidance on how to perform any given recommendation; instead, we point you to resources that will help you.

For the business side, this playbook provides valuable resources for driving new revenue opportunities, developing strategies for marketing, selling, and lead capture, as well as building deeper and longer term engagements with your customers through potential new service offerings like managed services.

For the technical side, the playbook offers guidance on a number of topics that range from the technical skills your team will need, to resources that you can use to accelerate learning, as well as an explanation of some of the key opportunities for technical delivery to focus on as you get started and grow your practice.

How this playbook was made

This playbook is part of a series of guidance that was written by Microsoft Partner, [Solliance](#), while working in conjunction with Microsoft One Commercial Partner and 11 other successful partners who have volunteered time and information to provide input and best practices to share with the rest of the partner community.

To validate the guidance provided in these playbooks, we conducted a survey of 555 partners from around the world with MDC Research. In this survey, we gathered insights on a range of topics, including how partners hire, compensate and train resources; their business model, revenue and profitability; what practices and services they offer; and what skillsets they have in place to support their offers. The results of this survey are provided in-line with the guidance found within this playbook.

CONTRIBUTING PARTNERS AND DATA SCIENTISTS

[Applied Information Sciences](#)

[Pragmatic Works Inc.](#)

Bob Schmidt, Freelance Data Scientist

[UST Global Inc.](#)

[Crayon Group](#)

[VISEO](#)


[Health Populi](#)

Willy Marroquin, AI Researcher

[Interknowlogy](#)

[Wortell](#)

[KenSci](#)



Using the playbook effectively

Quickly read through the playbook to familiarize yourself with the layout and content. Each section includes an executive summary and key actions for that specific topic. Review these summaries first to decide which areas to focus on. Go over content several times, if needed, then share with your team.

TO GET THE MOST VALUE OUT OF THIS PLAYBOOK:

- Get your team together and discuss which pieces of the strategy each person is responsible for.
- Share the playbook with your sales, marketing, support, technical, and managed services teams.
- Leverage the resources available from Microsoft to help maximize your profitability.
- Share feedback on how we can improve this and other playbooks by emailing playbookfeedback@microsoft.com.

Table of Contents

About this Playbook	2	Hire & Train	82
Partner Practice Development Framework	5	Executive Summary	83
What is Artificial Intelligence?	6	Hire, build, and train your team	84
AI Opportunity	10	Job Descriptions for your Technical Team	90
Industry Opportunities	11	Recruiting Resources	100
Define Your Strategy	19	Training & Readiness	101
Executive Summary	20	Operationalize	115
Define Your Practice Focus.....	21	Executive Summary	116
Understanding the AI Practice	22	Implement a Process.....	117
The Microsoft Approach to AI	25	Claim Your Internal Use Benefits	121
Pre-Built AI using Cognitive Services	30	Define Customer Support Program and Process.....	126
Building Custom AI.....	37	Manage and Support an AI solution in Azure	130
Microsoft AI Platform Summary	44	Support Ticket Setup and Tracking	132
Define and Design the Solution Offer	45	Implement Intellectual Property Offerings.....	133
Understanding Project Based Services	46	Setup Social Offerings	134
Understanding Managed Services	56	Create Engagement Checklists & Templates	135
Accelerate your Managed Service Model.....	62	Go to Market & Close Deals	136
Understanding Intellectual Property.....	63	Executive Summary	137
Define Industry Specific Offerings	67	Marketing to the AI Buyer	138
Define Your Pricing Strategy	68	Engage Technical Pre-Sales in Sales Conversations	140
Calculate Your Azure Practice Costs	71	Architecture Design Session (ADS)	142
Identify Partnership Opportunities	73	Go-to-Market and Close Deals Guide	144
Define Engagement Process.....	75	Optimize & Grow	145
Identify Potential Customers	76	Executive Summary	146
Join the Microsoft Partner Network	77	Understanding Customer Lifetime Value	147
Stay Informed on AI Matters.....	79	Guide: Optimize and Grow.....	149
Identify Solution Marketplaces	80	AI Playbook Summary	150

Partner Practice Development Framework

The partner practice development framework defines how to take an AI practice from concept to growth in five stages. It is the foundation of this playbook, and each phase of the framework is covered in a dedicated chapter.



What is Artificial Intelligence?

Strictly speaking, Artificial Intelligence (AI) is intelligence displayed by machines.

The term “artificial” is intended to contrast this display of intelligence with natural intelligence – the form of intelligence that is displayed by humans and other animals. At its core, AI refers to scenarios where a machine mimics the cognitive functions associated with human minds. These cognitive functions include comprehending, expressing, perceiving, calculating, remembering, organizing, reasoning, imagining, creating, and problem solving.

AI IS NOT STATIC – IT LEARNS AND THINKS

Fundamental to all of these cognitive functions is *learning* – we humans learn to comprehend the text in books, to express our thoughts in speech, to reason both deductively and inductively, and solve problems. The natural intelligence we display would not be possible without the learning that preceded it. When we learn something, we often start by experiencing specific examples and then generalize the specific examples into something we can apply more broadly. For example, think of how you learned throughout school: you were provided examples of the subject matter, but it was up to you to generalize that knowledge so that you could answer those tricky questions on the final exam. Just memorizing the examples the teacher gave you probably didn't help you pass the exam, you had to take what you had learned about the subject and apply it in a new situation (the exam).

It is the learning component that makes AI different from historical approaches to building machines or programming software. AI is not explicitly programmed to respond a certain way, it learns to respond that way. Coupling such learning with the modern capabilities available to software programs, including Internet connectivity, the ability to store and process huge volumes of data quickly and without fatigue, and recall data perfectly, leads to machines that can complement and augment human capabilities.

AI AMPLIFIES HUMAN CAPABILITIES

Of the cognitive functions listed earlier, the detail oriented, indefatigable nature of the machine lends itself well to comprehension, expression, perception, calculation, recall, organization, and reasoning (specifically in terms of working out inferences and entailment in answer to a question). However, cognitive functions that fundamentally draw on creativity like creating, imagining, reasoning (specifically coming up with the right questions with which to start and drawing conclusions from results), and problem solving (specifically structuring the problem solving) are tasks for which humans are, and most likely will always be better. Machines might mimic these functions, but are not likely to duplicate them, and will need to lean on humans for that creative spark. This is why partners building AI solutions need to see AI not as something that is replacing human capability, but rather amplifying it with the strengths of the machine. Microsoft CEO, Satya Nadella, in his book, *Hit Refresh*, captures the spirit in which partners should view AI by comparing it to the evolution of aviation:

“Today we don't think of aviation as ‘artificial flight’—it's simply flight. In the same way, we shouldn't think of technological intelligence as artificial, but rather as intelligence that serves to augment human capabilities and capacities.”

Satya Nadella, *Hit Refresh*

AI centers around amplifying the unique cognitive ingenuity of humans (imagining, creating, reasoning, and problem solving) and marrying it with best traits from intelligent technology (comprehending, handling of extreme detail, calculation, memory and recall, and organizing).

The computer brings with it a computational speed that is getting close to that of the human brain, storage capacity and recall capabilities at a breadth, depth and accuracy that is difficult for any person, and an inexhaustible energy to continue working.

AI WILL ASSUME JOB TASKS, NOT NECESSARILY ELIMINATE THE JOB

For many, AI is associated with ever increasing forms of automation. It is important to see beyond this and look towards how this combination of human and computer cognitive functions yields incredible new capabilities, freeing humans from some tasks to handle the more creatively demanding aspects of the job. Taking an enterprise perspective, this arrangement is about enabling leadership to ask the big questions and then iterate quickly, letting AI deal with the high volumes of minutiae it was carefully designed to handle. These technological shifts may allow workers throughout the enterprise, including leadership, to refocus their attention on work that requires uniquely human skills. Relying on the ready access to data and the computational capabilities, AI can reduce the time and effort required to inform decisions.

In this context, we must consider the impact of AI innovations on jobs. Like so many hyper-impactful innovations that precede it (e.g., from the cotton gin to factory automation to the Internet), AI will make tasks more efficient, and with this new efficiency, certain jobs will need less human involvement or even none at all. In these cases, effort will be required to understand this impact and special care will be needed to help with redefining the job, retraining, and/or redeploying the human worker.

AI may allow companies to reinvest in other areas that will drive job creation and allow employees to shift focus in their current jobs.

AI IS DIFFERENT FROM MACHINE LEARNING AND DEEP LEARNING



The term AI, because of its buzzword status, is often liberally applied to solutions that do not mimic human cognitive functions. One of the most frequent areas of confusion we heard from partners was differentiating AI from machine learning and deep learning. It is not that the three are un-related. They are very related, but more in a parent-child sense.

Machine learning and deep learning represent *some* of the techniques and tools used in the construction of an AI solution. It is not uncommon for to leverage multiple machine learners, or to combine machine learners and deep learners to produce the resulting AI solution that mimics human cognitive functions. At this early stage of AI, the difference between the AI solution and a solution consisting of only machine learning and deep learning is subtle. For example, both machine learning and deep learning techniques can help you build solutions that predict when an elevator might need maintenance. AI can take this insight one step further by prescribing what should be done about it and even taking action. In the elevator maintenance scenario, AI might identify that, based on its previous experience, the next best action is to order the parts which will need to be replaced (because they take a week to arrive) and to schedule a maintenance technician to install them, after the parts arrive.

Ethical AI

It's remarkable how much technology has changed the way we live and work over the last decade or two. Digital technologies, powered by the cloud, have made us smarter and more productive, transforming how we communicate, learn, shop and play. And this is just the beginning. Advances in AI are giving rise to computing systems that can see, hear, learn and reason, creating new opportunities to improve education and healthcare, address poverty and achieve a more sustainable future.

But these rapid technology changes also raise complex questions about the impact they will have on other aspects of society: jobs, privacy, safety, inclusiveness and fairness. When AI augments human decision-making, how can we ensure that it treats everyone fairly, and is safe and reliable? How do we respect privacy? How can we ensure people remain accountable for systems that are becoming more intelligent and powerful?

To realize the full benefits of AI, it is important to find answers to these questions and create systems that people trust. Ultimately, for AI to be trustworthy, it should be "human-centered" – designed in a way that augments human ingenuity and capabilities – and that its development and deployment must be guided by ethical principles that are deeply rooted in timeless values.

PRINCIPLES OF TRUSTWORTHY AI

Microsoft believes that six principles should provide the foundation for the development and deployment of AI-powered solutions that will put humans at the center:

- **Fairness** – AI systems should treat all people fairly and not affect similarly situated groups in different ways. Build AI systems from a diverse pool of AI talent, using representative training data and analytical techniques that detect and eliminate bias. This will require the involvement of domain experts in the design process, and systematic evaluation of the data and models.
- **Reliability** – Customers need to trust that AI solutions will perform reliably and safely within a clear set of parameters and respond safely to unanticipated situations. This requires extensive testing of training data and models, a robust feedback mechanism, and processes for documenting and auditing performance and determining how and when an AI system seeks human input.
- **Privacy and Security** – Not unlike the other solutions you deploy, AI systems should be secure and respect existing privacy laws. Without such protections, users will not share the data needed to train the AI. AI systems should be transparent about data collection, use good controls and de-identification techniques, and have policies that facilitate access to the data the AI needs to operate effectively.
- **Inclusiveness** – To benefit everyone, AI systems should engage and empower people and use inclusive design practices to eliminate unintentional barriers. AI technologies must understand the context, needs and expectations of the people who use them, and address potential barriers that could unintentionally exclude people. AI can be a powerful tool to enhance opportunities for those with disabilities.
- **Transparency** – When AI systems help make decisions that impact people's lives, it's particularly important that people understand how those decisions were made. People should know how AI systems work and how they interact with data to make decisions. This makes it easier to identify and raise awareness of potential bias, errors and unintended outcomes.
- **Accountability** – Those who design and deploy AI systems must be accountable for how their systems operate and should periodically check whether their accountability norms are being adhered to and if they are working effectively.

GOVERNANCE FRAMEWORK

A governance model is key to shepherding an organization to a common framework for building and deploying AI solutions that adhere to the organization's ethical patterns and practices.

Internally, Microsoft has established the Aether Committee, a board of executives drawn from across every division of the company, to focus on proactive formulation of internal policies and how to respond to specific issues in a responsible way. Aether will ensure Microsoft's AI platform and experience efforts are deeply grounded within Microsoft's core values and principles and benefit the broader society. Among other steps, Microsoft is investing in strategies and tools for detecting and [addressing bias in AI systems](#).

While there is great opportunity in AI, ensuring we always act responsibly for customers and partners should be a hallmark of our work.

INDUSTRY PARTICIPATION

A continuing collaboration between government, business, civil society and academic researchers will be essential to shape the development and deployment of human-centered AI to be trustworthy. Ongoing dialogues among these communities will help to identify and prioritize issues of societal importance, enable further research and development of solutions and sharing of best practices as new issues emerge, and, where appropriate, shape policy that can more readily adapt to these rapidly evolving technologies.

Microsoft is a founding member of The Partnership on AI, a collaboration of industry leaders, academics, nonprofits and specialists to collectively develop best AI practices, advance public awareness, and provide an open platform for discussion and engagement around AI's impact on people and society.

Learn more about the industry discussions on this topic by visiting [The Partnership on AI](#).

HOW YOU CAN GET STARTED

Bias in AI will happen unless it's built from the start with inclusion in mind. The most critical step in creating inclusive AI is to recognize where and how bias infects the system.

[Read this guide written by Microsoft's Design team](#) that breaks down AI bias into distinct categories so product creators can identify issues early on, anticipate future problems, and make better decisions along the way. It allows teams to see clearly where their systems can go wrong, so they can identify bias and build experiences that deliver on the promise of AI for everyone.

Encourage your team to keep learning. [The AI School](#) provides best practice training on Microsoft's latest AI technologies, and this [6-week, self-paced course on edX.org](#) helps data professionals learn how to apply practical, ethical, and legal constructs and scenarios so that they can be good stewards of their critical role in society today, while achieving optimal results.

Visit the [Microsoft AI Blog](#) to stay current on the latest developments at Microsoft on this front.

FURTHER READING

For additional materials exploring the fundamentals of AI and digital transformation, we recommend reading the following books:

- *The Future Computed*, by Microsoft
- *Hit Refresh*, by Satya Nadella
- *Artificial Intelligence and Machine Learning for Business: A No-Nonsense Guide to Data Driven Technologies*, by Steven Finlay
- *The Mathematical Corporation: Where Intelligence + Human Ingenuity Achieve the Impossible*, by Josh Sullivan & Angela Zutavern

AI Opportunity

Today, practices differentiate by virtue of their use or non-use of AI in their service delivery.

In the near future, AI will be assumed, and practices will differentiate by their skill in the pragmatic application of AI. From advancing medical research, diagnoses, and treatments, to increasing farm yields, the wonders of AI continue to surprise us with new possibilities. If you're not already leveraging AI in your solutions today, chances are you will soon. Partners of all types, from system integrators to hosters, are finding their use of AI has become a key differentiator for their service offerings and a chance to re-engage customers with end-to-end systems that learn from data and experiences to deliver new insights, efficiencies, and innovations.

According to [IDC surveys](#), 67% of organizations globally have already adopted or plan to adopt AI. And many adopters have seen returns that meet or exceed expectations, leading many to increase spending on AI in the next two years. IDC sees the compound annual growth rate for AI spending near 50% in the U.S. and even higher in Asia/Pacific.

Fueled by the enormous storage and processing power available in the cloud, AI are now able to analyze data at cloud scale. The ability to make decisions based on probabilities and deliver solutions built on data is transforming software development. Partners can integrate technologies from several AI disciplines such as computer vision and human language technologies to create end-to-end systems that learn from data and experience. Applications that communicate with humans using language and act as automated agents, called bots, make possible creating experiences that are highly personalized and will differentiate one brand to another.

The human element has become increasingly important as language and vision models now allow computers to infer meaning and intent. This creates opportunities that go well beyond transforming businesses. AI solutions can enrich our lives and improve the quality of life for others. In healthcare, for example, experts are now able to predict the onset of conditions by building predictive models from scans or data from wearable devices, helping providers deliver timelier treatment at less cost.

aka.ms/practiceplaybooks

DRIVING REVENUE WITH AI

You have multiple opportunities to drive revenue selling AI-related services. To understand these opportunities, it is helpful to understand the high-level process for delivering AI solutions, since the two are related. The high-level process is as follows:

1. **Envision AI:** The first revenue opportunity is to help your clients envision the possibility of AI and what it could bring to their organizations – helping them see a roadmap where AI becomes more and more a part of their daily operations and the benefits this brings.
2. **Implement AI:** The second revenue opportunity is in helping your customers implement the AI solution by integrating existing, pre-built AI APIs or developing custom AI solutions.
3. **Deploy AI:** The final revenue opportunity comes in helping customers get their AI solution into production. If your services stop at helping them build the solution, what will their internal teams do when something goes wrong? AI solutions can be complicated, and a customer may not have the internal team capable of supporting them. Providing support for AI solutions creates a great opportunity for recurring revenue.

In speaking with partners, we found that generating revenue from AI does not necessarily mean selling services that address the entire process. Some partners focus exclusively on envisioning the AI opportunity with their customers, but don't help them build the AI solution. Others, help customers build the AI solution, but don't support the solution in production (e.g., they might partner with a 3rd party who monitors the AI solution in production). While no partner we interviewed for this playbook focused exclusively on supporting deployed AI solutions in production, all agreed a successful business could be built solely around this highly specialized type of production support.

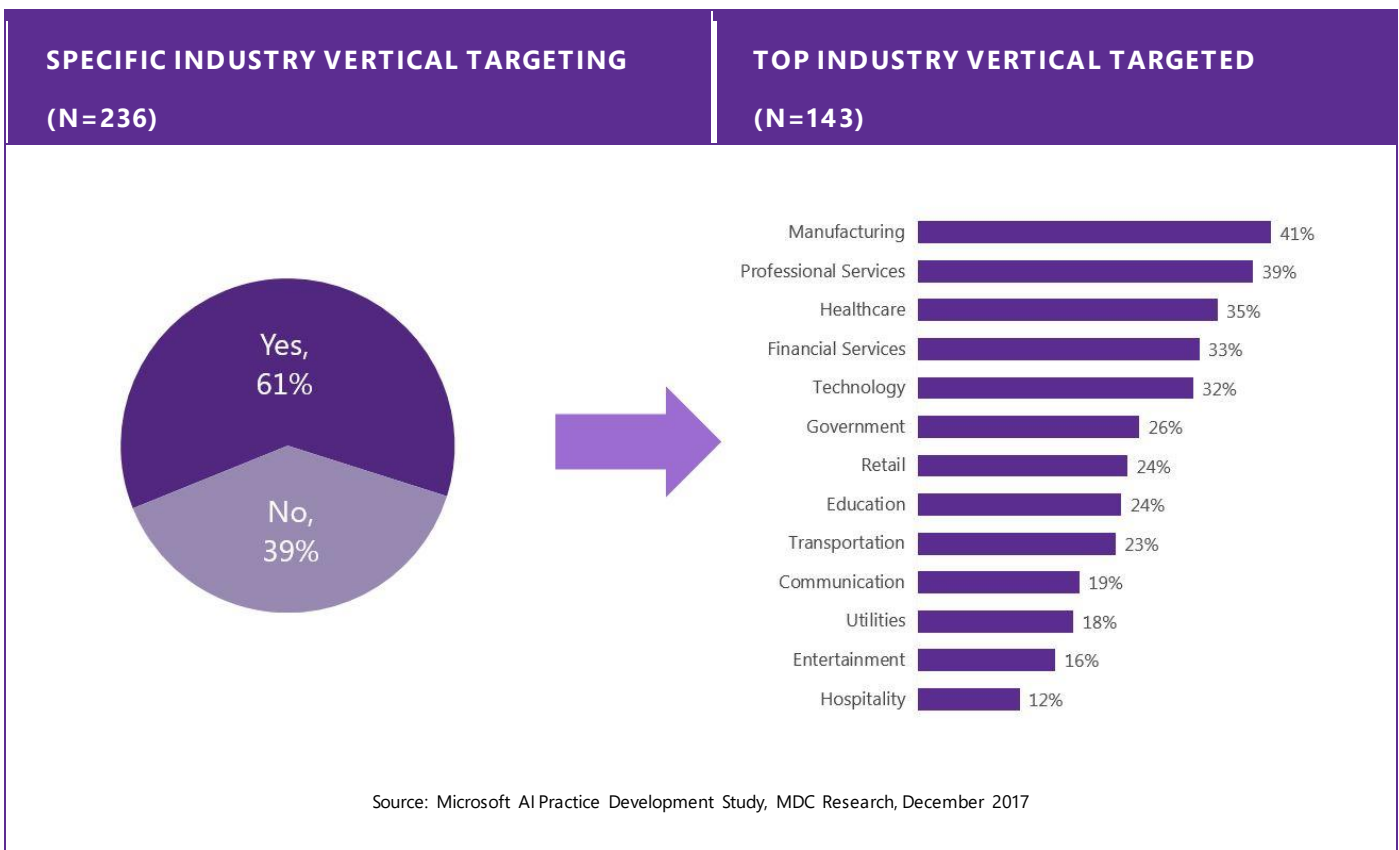
Finally, as AI becomes more and more integrated into the daily activities of a partner's customers, partners who offer AI services will be differentiated from those who do not. AI will not only create significant opportunities for partners, but also give them competitive advantages.



Industry Opportunities

The AI Opportunity is Universal

The opportunity for partners with an AI practice cuts across industry verticals. Irrespective if the industry is healthcare, finance, manufacturing, retail, government, or education, the AI opportunity is omnipresent and significant.





Partners have found opportunity in AI either by providing horizontal solutions that are applicable regardless of industry or by focusing their solutions and growing their domain expertise within specific industries. The following sections introduce some of the possible industry-specific solutions as motivation for the types of solutions partners may choose to deliver.

Almost Every
Industry Vertical

AI application and use
case

\$89.8 Billion

Revenue generated by
2025

A [2017 report from market research firm Tractica](#) forecasts that annual worldwide AI software revenue will increase from \$3.2 billion in 2016 to \$89.8 billion by 2025. In a related report, Tractica divides the field of AI into two broad buckets: those that enable human-like perception (e.g., object identification and detection, speech, language) and those enabling human-like analytical

capabilities (e.g., using big data approaches, detecting patterns and providing analysis). The use cases for AI across both buckets span nearly every conceivable industry.

On the business side, human resource departments are using performance analytics and machine learning to predict retention and development needs of employees and candidates.

Marketers are using real-time analytics and machine learning to customize and optimize customer experiences.

The examples go on and on, but in each case intelligent technology is being used to extend and augment human capabilities, freeing them up to be more strategic and productive.

And partners no longer have to employ a staff of data scientists to achieve such innovation. Microsoft's AI platform, Cognitive Services stack (including the Language Understanding Intelligent Service), provide the tools needed to build very capable AI-enhanced solutions. Along with Azure data services, Microsoft can help partners overcome another challenge to implementing useful AI solutions: the effort needed to collect and prepare an organization's data. Machine learning services can help automate the modeling of data for predictive analytics and visualization.

The Healthcare Opportunity

AI in healthcare is not about robots creating assembly-line healthcare. It is about systems that assist and support the wisdom and experience of well-trained clinicians in making better data-driven decisions and taking actions that best support the needs of those they serve. It does this by gathering and crunching massive amounts of data quickly and intelligently to identify patterns often overlooked or undiscovered in the traditional practice of care. The opportunity for AI in healthcare isn't just about making doctors and healthcare providers more efficient in their work; it's about making the lives of the patients better and saving lives.

The application of AI is nascent in the broad healthcare field, leaving a significant opportunity for partners.

AUTOMATING THE RECOGNITION OF DISEASE

Trained physicians can only scale to review and evaluate a limited volume of patients or patient images (X-rays, sonograms, etc.). In the beginning, the capabilities of AI to spot the disease will be used to extend the reach of the human physician beyond those who can make an office visit, and to enable evaluation of a higher volume of patient data than possible by the physician alone.

[Optolexia](#) built a dyslexia screening tool for young children. The tool uses a laptop, tablet, or desktop computer with an eye tracker mounted at the bottom. As a student reads text on the screen, the tool projects an infrared pattern toward the student's face and the eye tracker captures and analyzes reflections on the surface of the cornea. The data is sent to the machine learning engine in the cloud, which returns a numerical result that identifies the likelihood that the student has dyslexia.

Not all efforts in automating the recognition of disease require the implementation of custom machine learning models. Partners told us they are investigating ways that use the pre-built AI APIs like Microsoft Cognitive Services Custom Vision Service to train models by showing them smaller sets of examples of images with the disease (on the order of 50 or so images), and are reporting success in building highly accurate models that could one day automate the detection of rare diseases and diseases for which most physicians may not be trained or just may not be looking for – meaning the disease may go undetected until it is too late.

aka.ms/practiceplaybooks

Other resources are emerging that combine access to massive data sets of medical data along with pre-built AI APIs. One example of this is the [Microsoft Intelligent Network for Eyecare \(MINE\)](#) which provides APIs using AI to help reduce avoidable blindness. MINE is a consortium which will use geographically diverse and anonymized datasets of millions of patient records to build a global pool of knowledge and facilitate deep insights using Microsoft Machine Learning Technology. The insights derived from this data will help build the artificial intelligence to help in the elimination of avoidable blindness and scale the delivery of eyecare services across the planet.

EXPEDITING THE DISCOVERY AND USE OF THE LATEST RESEARCH

Systematic reviews are a required process that examines the latest research and establishes processes and guidelines around the application of research results. With the enormous volumes of research surfacing every day, it becomes very difficult for physicians to stay abreast of the research that is important to them and their patients. As [Cochrane](#) has demonstrated with its Project Transform and Evidence Pipeline, AI can be used to help automate the finding of relevant reports from studies and flag them for inclusion in systematic reviews. This analysis of medical documents at scale effectively helps the medical community find the needle in the ever-growing haystack.

RECOMMEND NEXT BEST ACTIONS FOR INDIVIDUAL CARE PLANS

Reaching an accurate diagnosis can sometimes be a challenge even for the most skilled providers. However, the mass digitization of patient data via EMR and EHR systems opens the door for AI and machine learning to help augment the diagnostic process to build effective individual care plans. Through analysis of the patient's digital health footprint, such as lab results, history, and reported symptoms, advanced analytics tools can surface potential diagnoses and recommend the next best actions for care.

As partner [Mazik Global identified with Pinnacle Hospital in Crown Point, Indiana](#), scheduling an operation in most hospitals today involves a dozen or more people, takes hours or days, involves a lot of duplicated work, and presents a frightening amount of room for error. Their aim? Using predictive analytics, Pinnacle doctors will be able to crunch data on thousands of similar patients to provide assistance in coming up with the best clinical pathway for that patient – and the expertise and other resources needed to treat them.

EMPOWER PATIENT COMMUNICATION

Without clear communication between patient and medical service provider, patients suffer. [Equadex](#) developed a way to facilitate communication between people on the autism spectrum who are nonverbal, or have language difficulties, with their parents, teachers, and medical providers. The traditional approach would rely on families and an attendant using a physical binder containing pictogram cards to communicate. Family members would line up the appropriate cards to construct a sentence and would show that to the child. This requires carrying a binder from room to room, creating opportunities for pictogram cards to be lost or misplaced.

Equadex™ created Helpicto™, an app for tablets and smartphones containing a database of pictograms made for mobile. The app can convert spoken text into series of images. Rather than build the AI technology from scratch, Equadex™ found its solution with Microsoft Cognitive Services APIs, and the Microsoft Azure cloud platform. Equadex™ hopes that in the future it can also improve communication for nonverbal children and adults with other conditions that make communicating verbally difficult, such as Alzheimer's disease.

The Financial Services Opportunity

Financial services and insurance services technology companies are disrupting the market leveraging data, advanced analytics and predictive models to offer better and more personalized products, prices and conditions. The application of AI in financial services leaves many opportunities including in areas like risk analytics, fraud prevention, and in prescribing the next best action for supporting a customer.

Partners can find opportunity providing financial services organizations with solutions that can help them engage their customers in new, interactive ways, empower their employees to innovate the customer experience, optimize their operations to drive efficiencies across their business and transform their products and services to become a trusted advisor.

ENABLING PERSONALIZED, REAL-TIME BANKING EXPERIENCES WITH CHAT BOTS

Bots and intelligent assistants are natural solutions for the first wave of customer contact. Much like how interactive voice response (IVR) transformed inbound service calls, AI can intercept and handle common, straightforward issues through chat and messaging services, so customers can quickly and independently resolve simple issues that would otherwise have required human intervention. With the chatbot, a customer can simply type in a question and the bot engages to surface the answer.

[Nedbank](#) believes bots are the next industry disruptor, and as one of the major banks in South Africa it is using bots within its virtual call center solution that can understand the context of customer questions, answering 80% of the questions asked at 10% of the cost of live agents. The application of AI powered bots enables the bank to scale out its virtual workforce quickly and cost-effectively. The bank can exploit bot technology to provide a one-two punch of competitive advantage: provide better client service even as it reduces the cost of providing that service. And enhanced client service at lower cost is key to another Nedbank strategy: expanding its individual investor business while maintaining its traditional base in institutional and broker-based financial services.

MODERNIZING SUPPORT

[Fiducia GAD IT AG](#) needed a solution that would help them address the IT-related support requests in its core business providing IT services for cooperatives and rural cooperative banks, private and special banks as well as commercial enterprises. They found low-hanging fruit in addressing questions regarding the status of a support ticket with the use of a chat bot, which made up for approximately 10% of all the support requests seen by their call center. The chatbot can access ticket data and generate answers on the current processing status, helping to alleviate the pressure on their live call center agents. Fiducia & GAD imagines a future where their bot can do far more beyond providing ticket status, such as searching user documentation stored at various locations and even presenting a personality that fits the company.

IDENTIFY THE NEXT BEST ACTION FOR THE CUSTOMER

Solutions for real-time insights incorporate AI tools such as sentiment analysis to help financial services organizations assess the likelihood of a deal closing or the level of a customer's loyalty. Personally-tailored encounters powered by machine learning recommendations can engage and delight customers with information and offers that are relevant to them. Additionally, financial services employees can generate cross-sell and upsell opportunities with a higher likelihood of acceptance, resulting in increased wallet share.



The Manufacturing Opportunity

Unlimited computing capacity in the cloud and real-time analytics capabilities enables manufacturers to access new insights and build systems of intelligence like never before. Forward-thinking manufacturers are looking to use these capabilities to optimize their supply chain and production operations, engage their customers in powerful new ways, transform their services and products, and empower their employees through customer insights. Partners should pay attention to the opportunities this creates.

PREDICT MAINTENANCE NEEDS BEFORE IT BECOMES NECESSITY

Elevator manufacturer [ThyssenKrupp](#) wanted to gain a competitive edge by focusing on what matters most to its customers in buildings the world over: reliability.

By connecting its elevators to the cloud, gathering data from its sensors and systems, and transforming that data into valuable business intelligence, ThyssenKrupp is vastly improving operations, and offering something its competitors do not: predictive and even preemptive maintenance.

The system contains an intelligent information loop: data from elevators is fed into dynamic predictive models, which continually update datasets. Now, the elevators can actually teach technicians how to fix them, with up to 400 error codes possible on any given elevator, which can significantly sharpen efficiency in the field, resulting in dramatically increased elevator uptime.

INTELLIGENTLY FILTERING THE SIGNAL FROM THE NOISE

About 20 years ago, [Rolls-Royce](#) went from manufacturing and selling engines to extending comprehensive maintenance services to the airlines that use its engines. The company's TotalCare® Services employ a "power by the hour" model in which customers pay based on engine flying hours. The responsibility for engine reliability and maintenance rests with Rolls-Royce, which analyzes engine data to manage customers' engine maintenance and maximize aircraft availability. This model has been highly successful for Rolls-Royce and has created relationships in which airline customers increasingly rely on the company to provide information that optimizes the costs and scheduling related to engine maintenance.

By looking at wider sets of operating data beyond their engines and using machine learning and analytics to spot subtle correlations, Rolls-Royce can optimize their models and provide insight that might improve a flight schedule or a maintenance plan and help reduce disruption for their customers. For example, by understanding the actual condition of a component versus the expected lifetime, Rolls-Royce can help their customer decide if maintenance on an aging component can be deferred to the next scheduled maintenance window versus requiring immediate maintenance. Moving to an approach based on a component's actual condition could potentially add up to tremendous savings across a fleet by minimizing the disruption and cost of maintenance.

AUTOMATICALLY CONTROL ENERGY GRID LOAD

[eSmart Systems](#) designed an automated demand response solution that collects data from virtually any type of meter or sensor. It then runs predictive models to forecast potential capacity problems and automatically control load to buildings or other infrastructure to prevent outages. The solution provides a short-term 24-hour forecast, a long-term monthly forecast, and a temperature forecast, and it offers a centralized way to monitor and manage the entire grid.



The Retail Opportunity

In the modern retail environment, consumers are well-informed and expect intuitive, engaging, and informative experiences when they shop. Retailers need solutions that can help them delight their customers with personalized experiences, empower their workforce to provide differentiated customer experiences, optimize their supply chain with intelligent operations, and transform their products and services. All of these are opportunities for partners to bring AI-powered solutions.

INNOVATE PRICING

Thanks to an innovative, dynamic pricing engine, the membership marketplace at [Jet.com](#) allows shoppers to watch prices of items in their basket change in real time, based on selections they make that affect the actual costs of that particular transaction – product warehouse location, items in a shipment, payment method, returns, and more.

The pricing engine will continually work out the most cost-effective way to fulfill an order from merchant locations closest to the consumer. The engine will also figure out which merchants can fulfill most cheaply by putting multiple requested items into one shipment. As a result of being smarter about fulfillment, Jet.com can cut 10% of the cost of a typical e-commerce transaction.

APPLY AI FOR BETTER CUSTOMER SERVICE

[Dixons Carphone](#) is a major European electronics and telecommunication retailer and service provider. The company decided to explore AI to create a conversational bot. Named Cami, the bot answers questions through the company's Currys brand website and Facebook Messenger to help customers and store employees research, find, and save products and check stock, which provides consistent online and in-store experiences. Cami accepts text-based input in the form of questions, and she also accepts pictures of products' in-store shelf labels to check stock status. Dixons Carphone programmed Cami with information from its online buying guide and store employee training materials to help guide customers to the right product – they did not have to create any new information. Instead, they were able to pull the existing content together in a different way and with Cami give customers the ability to search through it using natural language.

The Government Opportunity

A digital organization is one that is constantly striving to deepen the engagement between citizens and their government. Indeed, technology is the catalyst enabling the delivery of services and responsive that meet expectations. More than simply automating traditional processes, a digital organization is one that continuously strives to streamline operations across senathrvices. The integration of data and AI not only helps leaders interact and adapt, but also contributes to a lower cost of operations.

TRACK TRENDS THAT INFORM FUTURE PLANNING TO ACHIEVE DESIRED OUTCOMES

A modern government organization understands that the right insight is the key to successful outcomes. The good news is that AI can help. By using AI tools and machine learning to analyze organizational data, as well as public data such as news sites and social media, organizations can more effectively surface and track trends that inform future planning, allowing them to better identify areas of opportunity that would positively impact their citizens. The benefits of AI extend beyond the application against textual data, to aerial imagery provided by planes and satellites. For example, [land use classification models](#) can be used to track urbanization, deforestation, loss of wetlands, and other major environmental trends using periodically collected aerial imagery – all of which can feed into insights and trends into future planning around zoning, regulation, compliance, and enforcement.

LEVERAGE INTERNAL AND PUBLIC DATA TO MEASURE AND AUGMENT THE IMPACT OF GOVERNMENT INITIATIVES

In today's modern world, citizens use many tools to engage with their government. These tools create a digital footprint of citizen sentiment about everything from yesterday's announcement to tomorrow's big initiative. AI and advanced analytics can make that sentiment come to life and provide a clearer picture of what's resonating with your citizens. By aggregating engagement data and applying text analysis with AI, you can surface urgent issues and gauge satisfaction. Infuse machine learning to predict future sentiment and develop KPIs that drive effective evolution and change management.

A great example of public data that could enable powerful AI is [USA Facts](#), a website launched by retired Microsoft CEO, Steve Ballmer, that provides data to support a non-partisan look at government spending at the federal, state, and local levels.

CAPTURE, PRIORITIZE AND ROUTE SERVICE REQUESTS TO THE CORRECT EMPLOYEE AND IMPROVE RESPONSE TIMES

A busy government organization gets innumerable services requests on an annual basis. The ability to address such a volume of requests while maintaining high service levels can often be a strain on operations resources. AI tools can help automate this process by capturing incoming service requests and routing them to the correct employee in real-time. Machine learning can further impact service levels by adapting over time to surface insight around the speed to closure for specific types of requests, thus refining prioritization and improving response times.

LEVERAGE ADVANCED ANALYTICS TO EXPEDITE THE WORKFLOW PROCESS AND IDENTIFY THE NEXT BEST ACTION

Employees field countless requests from internal and external sources on a daily basis. Accelerate the case management workflow process with AI and advanced analytics. Instead of checklists or manual processes, employees can let AI tools do the heavy lifting by using knowledge chatbots to guide them through specific cases based on a pre-determined process workflow. This allows cases to be managed with consistency across the organization. Machine learning can also augment this process by identifying redundancies in process steps or areas of inefficiency within the workflow.



The Education Opportunity

AI presents a significant opportunity for partners in education. From enabling students with disabilities to succeed with AI as their always-on companion, helping them overcome their speech, reading or learning disabilities to helping teachers and districts optimize their efforts to enable student success.

PREDICTING STUDENT DROPOUT RISKS

The [Tacoma Public School district](#) wanted to answer the question – is it possible to predict whether students are at risk of dropping out of school? They began their journey as many AI initiatives do, by extracting the fundamental insights from the data they already have and shaping it to provide the descriptive analytics that told them what was happening. With a wealth of data in place, they were ready to take the next logical step: using data to predict which students were at risk of not completing their education. As they progressed with their model, they produced a model that was almost 90% accurate. This combination of descriptive and prescriptive analytics has helped Tacoma Public Schools improve graduation rates from 55% to 82.6%. Their efforts have not stopped there – they are moving towards implementing a real-time snapshot of student drop-out risk, enabling intervention and support to positively affect the student outcome at a much earlier time.

Microsoft
Partner
Network

Define Your Strategy

Artificial Intelligence

aka.ms/practiceplaybooks



Executive Summary

Now that you understand the opportunity before you in building an AI practice, the first step is to define the strategy you will use to build your practice. Like the foundation of the house, thinking through your strategy is critical to the long-term success of your practice, and it will be time well spent.

We begin by providing an overview of the maturity model for the AI practice, and touch on the opportunities present by industry vertical.

Then we will guide you through the process of defining your offer and its value proposition. This is a critical piece of your strategy – specifically, the definition of what you will sell and why customers will want to buy it. Along the way we will review the four cloud business models (reselling, project services, managed services, and intellectual property), their respective profitability, and how you can assess the profitability of your own practice. For project services, managed services, and intellectual property, we provide you guidance on what other successful partners are selling, as well as recommendations on what to include in your own project, managed services and intellectual property offerings.

We will help you drill into how to price your offer base, including what pricing strategy to use to drive adoption of your offer, and how to minimize your risk by establishing up-front fees and payment terms. The ultimate goal is to help you build a solid business plan that addresses your team, marketing, sales, and financial aspects.

Then we dive deeper into sales to help you define your pre-sales and post-sales engagement process, and how to compensate sales executives.

If you are not yet a Microsoft Partner, we will give you a walking tour of the Microsoft Partner Network, the programs you can leverage to grow your practice, how to earn competencies that yield additional benefits, and how to maximize the benefits you get from the program.

After that, we'll give you a head start in how to identify potential customers when starting your practice, as well as potential service offerings.

We'll conclude this section by helping you understand support – how to support your customers, Microsoft's support offerings, and the support-related benefits you get from establishing competencies in the Microsoft Partner Network.

Let's get on to defining your practice strategy.


Top 5 things to do

Here are the top 5 things you should absolutely do when defining the strategy for your practice.

- Understand the AI opportunity
- Define your practice focus
- Define business model
- Define and design the solution offer
- Develop your engagement process

Define Your Practice Focus

Through your AI practice, you can help your customers create solutions that amplify human ingenuity with intelligent technology. Partners who deliver AI solutions ultimately have practices delivering solutions that include one, two or all three of the following capabilities:



REASONING

Help your customers build solutions that learn and form conclusions with imperfect data.



UNDERSTANDING

Empower your customers with solutions that interpret the meaning of data, including text, voice, and images.



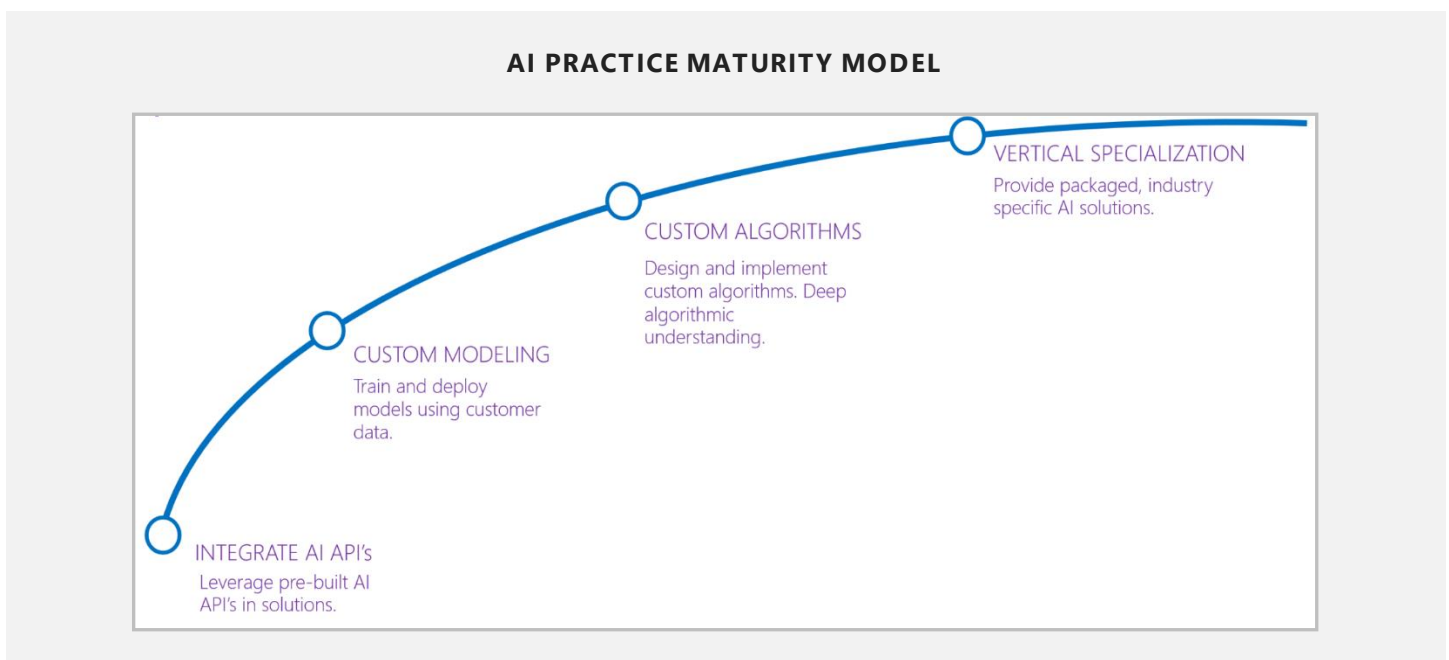
INTERACTING

Enable your customers with solutions that interact with people in natural ways.

Understanding the AI Practice

How do you start and grow your AI practice?

Reasoning, understanding, or interacting capabilities are what partners ultimately deliver in their practice, but how did they get to the point of delivering such capabilities in the first place? For some partners, their AI practice is an extension of their data, advanced analytics, or application development practice. It is important to understand that not all partners will necessarily have a 100% dedicated AI practice, even though they apply AI in the solutions they deliver. Along these same lines, most partners began their AI practices as generalist practices applying AI to the problem domain or industry of the opportunity present, without necessarily specializing to that domain or industry. The following diagram illustrates the AI maturity model and the sections that follow explain each phase in more detail.



INTEGRATE PRE-BUILT AI APIS

As partners grew and specialized their practice, we found that generally what evolved was the sophistication with which they implement AI capabilities. In our survey of 555 partners and one-on-one interviews, partners confirmed that the leveraging of pre-built AI APIs was the launch point for their AI practice. In this approach, partners build AI solutions by integrating various pre-built AI APIs in their application to deliver capabilities around reasoning, understanding, and interacting. This approach has the least skill demands on the practice, since it is predominantly a software development effort, and that is why it makes for a natural entry point into producing AI solutions.

CUSTOM MODELING

The next level of maturity we found partners demonstrating was in performing custom modeling, whereby they would train models using machine learning and deep learning techniques, using the customer's data to train the model. When building an AI solution, the core intelligence often is produced by applying machine learning or deep learning techniques to create models. Without getting into the details, models can be understood as capturing the relationship between input data and the historical outcome in such a way that the model can be applied in the future to inputs not seen and predict the outcome. The process whereby the model learns these relationships is called training.

These models effectively generalize from the patterns present in the customer’s data to support various forms of prediction – such as converting speech to text, understanding a user’s request, grouping similar data, predicting true/false or numeric outcomes, or synthesizing a human friendly answer. This level of practice maturity brings with it more demands on the skills present within the practice. In comparison to integrating pre-built AI APIs, partners performing custom modeling have to address many requirements:

- They need the historical data that contains the inputs and outcomes used to create the model.
- Because they now need to work with data, they need skills in data preparation and manipulation.
- They need to have an understanding of the machine learning or deep learning tools they are applying, and know how to evaluate if the model created is actually accurate enough to be useful.

CUSTOM ALGORITHMS

The most technically sophisticated form of AI capability creation occurs when partners may provide services that implement custom learning algorithms. In this case, the partner, in addition to performing custom modeling, is writing their own learning algorithms. Some partners interviewed with this level of sophistication are creating more than machine learning or deep learning libraries, they are building applications that self-learn, whereby the training is not guided by a human, but by the application itself as it writes and executes new programs to learn new capabilities. By the time the application has stabilized in its learning (e.g., it has “mastered the subject”) it may have written dozens of specialized learners, each helping it to learn from different data.

INDUSTRY SPECIALIZATION

The most mature partners we observed specialized in an industry vertical. That is, they provide packaged, industry-specific AI solutions. Partners at this level of maturity not only demonstrated technical sophistication in custom modeling (or in some cases custom algorithms), but deep expertise the domain of the industry. That deep domain expertise in turn yields a better informed AI.

Where should you begin?

There is significant value creation possible for partners who know the fundamentals of how to frame the customer problem in terms of AI, select the appropriate AI API, and integrate pre-built AI APIs into a customer’s solution. For the end customer, the end value delivered is indistinguishable from solutions built using more labor-intensive custom modeling or custom algorithms.

A curious side-effect of the large opportunity available in AI, was that partners who only integrated pre-built AI APIs indicated little need to evolve their practice past this phase in the near term, given the volume of opportunities they are currently realizing integrating AI APIs.

Business Transformation with AI

When defining your AI practice, you will deliver solutions that provide reasoning, understanding, or interacting capabilities, but in what form of are those capabilities sold to customers? The diagram illustrates three sales plays to get you started. Note that you should treat each of these as an evolutionary play that gets your practice in the door and expands your engagement with the customer as the sophistication of your AI practice and the AI solutions you deliver deepens.

CREATE INTELLIGENT AGENTS

Use conversational agents as low-friction entry points

MODERNIZE APPLICATIONS WITH AI

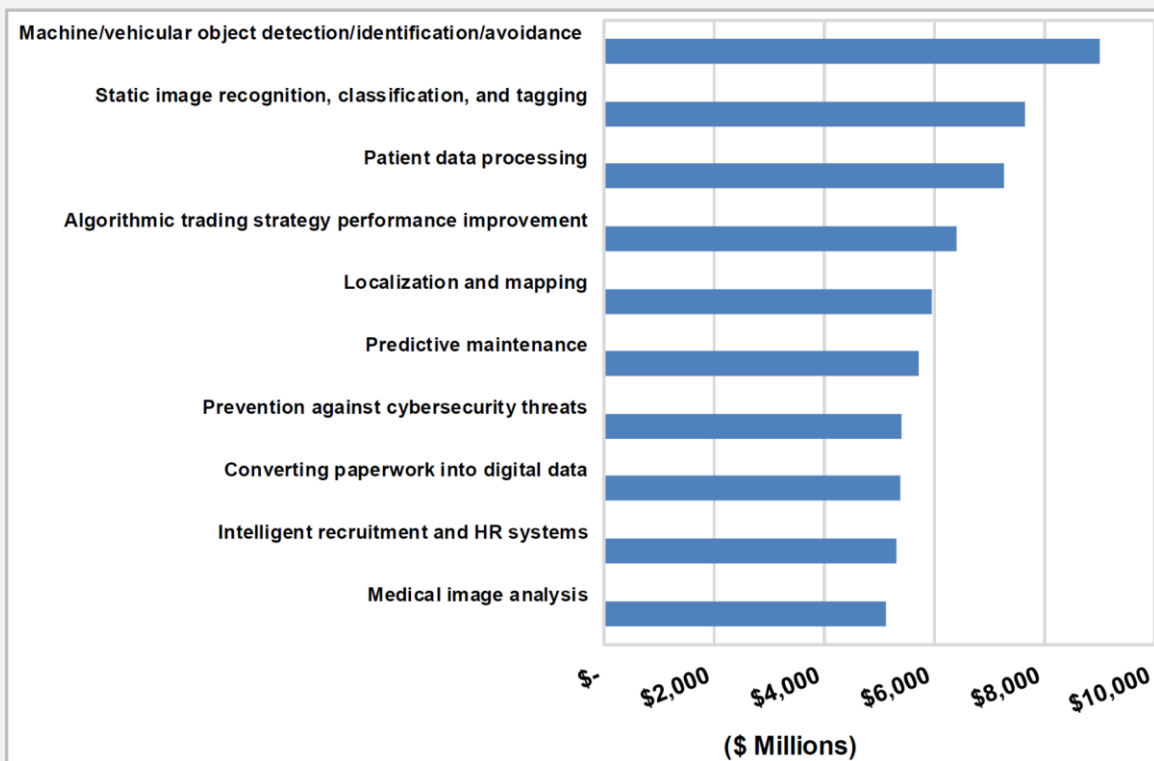
Modernize existing frontend and backend applications with infused AI

TRANSFORM BUSINESS PROCESSES WITH AI

Industry vertical approach by high priority scenarios in manufacturing, retail, and financial services

WHAT SPECIFIC USE CASES SHOULD YOU START WITH?

The maturity model presented describes the approach partners take to delivering AI, but what type of solutions are actually being delivered in industry? According to the [Tractica 2017 Artificial Intelligence Market Forecasts](#) report, the top 10 use cases for AI from a projected revenue perspective span across industries from healthcare (patient data processing and medical image analysis), financial services (algorithmic trading strategy performance improvement), manufacturing (machine/vehicular objection detection/identification/avoidance), retail (static image recognition, classification, and tagging), and government (prevention against cybersecurity threats). Consider the magnitude of the forecasted revenue numbers in the chart below when considering which AI use cases your practice should start with or further invest in.



Source: Tractica 2017 Artificial Intelligence Market Forecasts

What roles does your practice need?

The AI practice requires a unique combination of domain expertise, computer science expertise, and mathematical expertise. At the core of the practice is a data science team consisting of data scientists, data architects, data engineers, and developers.

The implementation team who builds the AI usually is a subset consisting of data scientists and developers. We explore these roles in greater detail, including their responsibilities, their skill and experience requirements, how these roles are trained and recruited in the Hire and Train chapter later in this book.

The Microsoft Approach to AI

A human-first approach that builds on the power of the Microsoft cloud.

The Microsoft approach to AI centers around enabling AI solutions that amplify human ingenuity and takes a principled approach that leads with ethics, accountability, and inclusive design to empower people and organizations and positively impact society. It is in this humans-first context, that Microsoft provides the technologies needed to build AI solutions in the form of its Microsoft AI platform.

WHY PARTNER WITH MICROSOFT FOR AI?

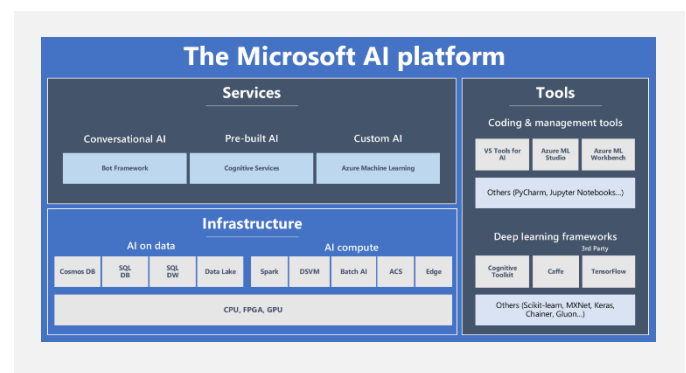
With more than 25 years of experience in AI research with a focus on product development, Microsoft has produced record-breaking advances in AI algorithms which they now offer to customers and partners in the Microsoft AI platform. Innovations used in products like Bing, Xbox, and Office are made available for partners to leverage in bringing their AI applications to life, from understanding natural language and speech, to facial recognition. Services like these appear in Cognitive Services, which are themselves easy to implement.

Additionally, the Microsoft AI platform provides a breadth of intelligence and knowledge of APIs so developers will be able to find the intelligence feature they need; and importantly, they all work on whatever language, framework, or platform developers choose. AI is built using a rapidly evolving open source ecosystem of tools, and the Microsoft AI platform does not force developers to abandon their favorite tools – instead, it meets them where they are and enables them to use their favorite tools in the construction of their AI solution.

It's important to realize that the Microsoft AI platform builds upon the strengths of the Microsoft cloud, notably:

- Microsoft is the only cloud provider that combines a Hyper Scale cloud offering, a truly hybrid platform, and an enterprise-level support for your cloud-based AI workloads with enterprise-level SLAs.
- Microsoft is considered by Gartner as a leader in six¹ of the nine cloud-related Magic Quadrants.
- Microsoft is industry leading with security, privacy and compliance. All of Microsoft's services are independently verified to meet legal and compliance requirements, are financially backed, and offer transparent information on their availability.

The Microsoft AI platform enables you to engage in creating AI solutions that leverage the scalability of the cloud, the flexibility to work with data on-premises, and provides numerous services that assist in building AI, infrastructure that is used train and run the AI solutions and tools that keep the AI implementation team (e.g., developers and data scientists) productive. In building your AI solution, you will need to select from the assortment of services, infrastructure and tools available in the Microsoft AI platform. In the sections that follow, we will explore the Microsoft AI platform from the perspective of the big decisions you need to make in your AI practice.



¹ Magic Quadrants: [Cloud Infrastructure as a Service](#); [Public Cloud Storage Services](#); [Access Management](#); [Server Virtualization Infrastructure](#); [Disaster Recovery as a Service](#); [Enterprise Application Platform as a Service](#)

Understanding How AI Models are Created

In the previous section, Understanding the AI Practice, we examined the relationship between AI and models, and defined models as capturing the relationship between input data and the historical outcome in such a way that the model can be applied in the future to inputs not seen and predict the outcome. We also identified that the process whereby the model learns these relationships is called training, and that we use a trained model to support prediction. In order to understand the value that the Microsoft AI platform provides its important to at least have a high level understanding of how these models are created.

The following diagram summarizes the high-level phases that lead to the creation of useful predictive model. For a detailed look at the model building process, see the section Applying the Microsoft Team Data Science Process later in this book.

1 Prepare Data

Prepare Data: In the Prepare Data phase, the data is collected from sources and prepared for use in training the model. During this phase, the data may be cleaned and de-duplicated, the contents of the data is understood and the data that is the most informative in predicting the outcome is selected. This work is often referred to as data wrangling. Typically the data wrangling is performed by either a data developer or data scientist who is writing programs to collect and prepare the data.

2 Build Model

Build Model: In the Build Model phase, a subset of the prepared data (which contains both the input and the outcome) is fed into a machine learning or deep learning algorithm to train the model. Then performance of the model is measured against another subset of the prepared data (referred to as the test or evaluation data set), and the model is evaluated on how well it performed in predicting the outcomes described in the test data set. As a simple example, one might ask *did the model predict the correct outcomes better than random guessing (e.g., 50% of the time)?*

Assuming the model performed adequately, then the model is saved to a file ready for deployment. The build

model phase is typically iterative as various approaches are tried to improve the performance of the model until the desired performance is achieved. The model building requires some understanding of the machine learning and deep learning algorithms used, and in some cases the math upon which they are built. As such, model building is usually performed by a data scientist who has the right background across computer science and mathematics.

3 Deploy Model

Deploy Model: In the deploy model phase, the created model file is typically copied to a location where it can be used by the AI application for making predictions. This step is typically performed by developers, or more specifically, DevOps engineers who are responsible for making sure the model is deployed correctly into production.

UNDERSTANDING THE EFFORT

In the model creation pipeline (prepare data, build model, deploy model), note that data preparation is a time-consuming phase requiring lots of detailed efforts to provide data that will yield a good model during training. It is often said that 80% of the time allocated to model creation is spent just in the Prepare Data phase alone. Also, consider that beyond the effort required in data preparation on existing data, that it is during this phase you often discover that you do not actually have the data you need and that you need to go collect it, which of course can add significantly to the time required to complete this phase.

If the Prepare Data phase typically takes about 80% of the total time, then the Build Model phase typically takes the next 19%. The Build Model phase does take a non-trivial amount of time, but more importantly the effort takes a certain expertise to complete successfully. Otherwise it is entirely possible to build models that are worse at predicting outcomes than a coin toss might.

As you have probably guessed, the Deploy Model phase takes at most the remaining 1% of the time. As we will see, the effort involved in model deployment is typically small involving either a file copy or the deployment of a web service.



Pre-built or custom AI?

One of the very first decisions you need to make based on your AI goals is whether you can utilize pre-built AI APIs provided by the Microsoft AI platform, or if you actually need to be building your own custom models against customer supplied data using infrastructure and tools provided by the Microsoft AI platform.

The pre-built AI option enables you with faster time to market because you do not need to actually build or train the AI, Microsoft has already done that across a range of scenarios. There is no complexity of having to collect and prepare the customer's data, identify the best data with which to train underlying models, select the appropriate machine learning or deep learning algorithms, and test that the selection works against real data not seen by your trained model. Instead, you simply integrate calls to REST APIs in the AI solution in order to benefit from the capabilities of the APIs. These pre-built AI APIs cover a range of scenarios that you might need in your AI application, and in most scenarios you can think of them as supporting the 80% cast, the general case – supporting the needs of the most common applications within those scenarios.

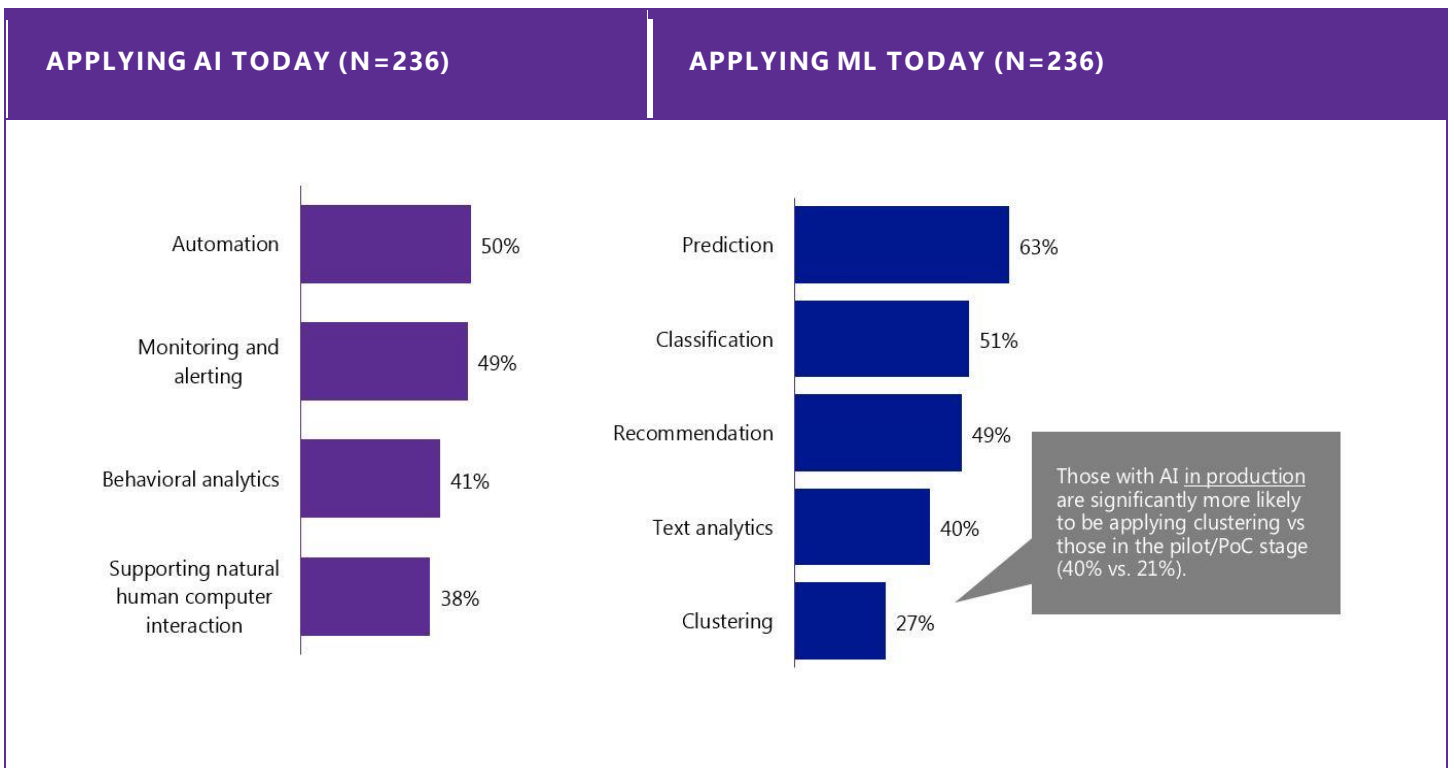
For example, one of the pre-built AI APIs in the Microsoft AI platform is the Speech API which converts human speech to text. As a pre-built API, you should expect that it understands common speech. If your AI solution needs to transcribe spoken messages between people talking about everyday things, in typical environments, the Speech API provides a quick and easy to implement solution. However, if your scenario is very domain specific, such as you want to transcribe a doctor's annotations to a medical chart (which might include a very unique vocabulary) or you need to recognize speech in unusually noisy environments, then the Speech API is likely not the right solution because it is too general purpose and might not transcribe as accurately.

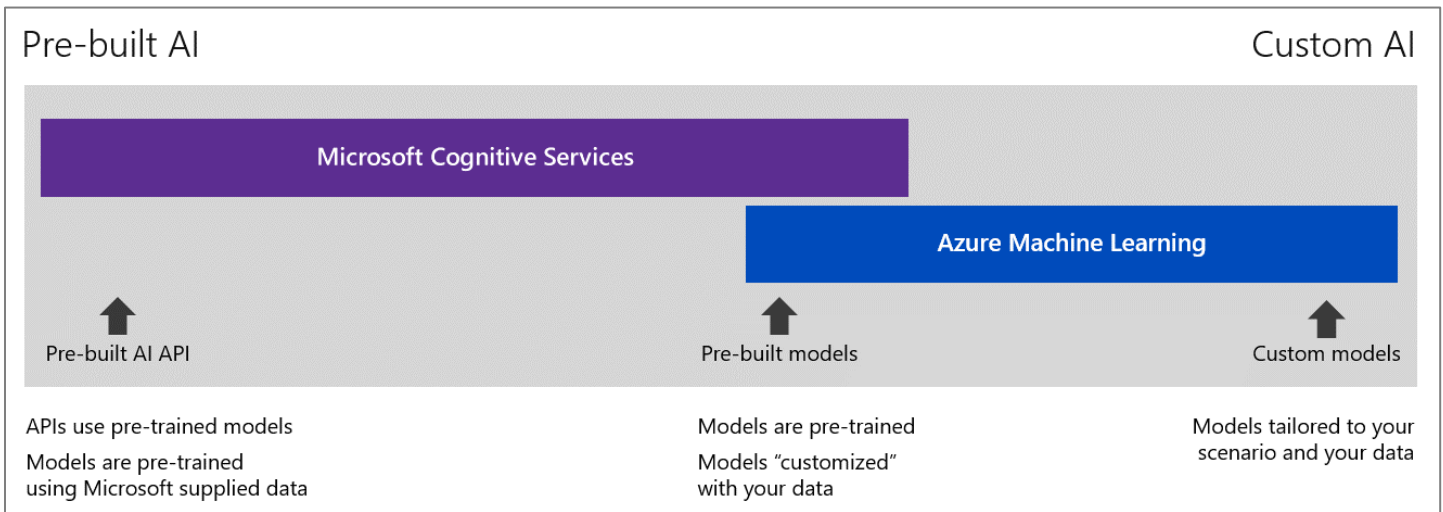
When it comes to designing your AI solution, it is tempting to want to use the latest technologies and to build a custom AI. Resist this temptation and always evaluate first if there is a pre-built AI that you can leverage in your solution. Then, in those cases where you have evaluated the pre-built AI options and have ruled them out, you can consider the options from the Microsoft AI platform that enables you to build custom AI. This approach will save you time and wasted development effort.

Comparing Approaches

Before diving into the technology choices available in the Microsoft AI platform that support either pre-built AI or custom AI implementation, it is useful to have a high-level understanding of how AI solutions come together, agnostic of the actual technology in the two different approaches.

In our survey of 236 partners with an AI practice, automation represents their top application of AI, whereas prediction is their top application for machine learning.





PRE-BUILT AI APPROACH

From a high level, there are two options for consuming pre-built AI. You can:

- **Use pre-built AI APIs.** In this case, you are integrating calls to AI APIs that are exposed as web services.
- **Use pre-trained models.** In some cases, models already exist to perform the type of prediction your AI needs to perform. All your AI application needs to do is load the models and invoke their predictive capability. Some pre-trained models are optimized to run on mobile devices, so they can bring AI capabilities to the device without draining its battery.

CUSTOM AI APPROACH

When it comes to custom AI, you need to choose where you will build your custom models and then how you will make the model available to application. The high-level options for doing this are:

- **Build the model on-premises.** Some models can be built on a laptop, others might need the resources of a datacenter. If your data lives on-premises, you can certainly perform the model building on-premises.
- **Build the model in the cloud.** If the data can be accessed from the cloud, then consider building your model in the cloud to take advantage of the significant scalability available for both data storage and processing.

Once you have a trained model, you have a few options for how you make its predictive capabilities available to applications:

- **Expose custom models as REST services.** In this scenario, you deploy the model so that it's predictive capability is accessible as a REST service.
- **Use custom models directly in the context of AI applications.** This approach is similar to the previous, except that you would save the model to a file after training, and then load the model from a file in your AI application when you want to apply its predictive capability.
- **Use custom models in the context of SQL database.** In this case, your trained model is stored within the table of a relational database, and its predictive capability is accessed by executing a SQL store procedure.

Now that you understand both approaches, the value of using a pre-built AI should be clear- if the pre-built AI provides the desired functionality, it saves you from having to perform ANY of the model creation pipeline.

In the sections that follow, we will explore the options for pre-built AI available from the Microsoft AI platform. Following that, we will tour the options from the Microsoft AI platform that help you build custom AI.

Pre-Built AI using Cognitive Services

Cognitive Services enable partners to infuse their apps, websites and bots with intelligent algorithms to see, hear, speak, understand and interpret user needs through natural methods of communication—accomplished via straight forward configuration and integration of REST APIs. The Cognitive Services models work out of the box, with no model training or testing required. This reduces the time required to get into production, the skills required to implement, and hence the cost to deploy.



An example of one of the Cognitive Services is shown in the figure that follows. On the left is a photograph of the cityscape. This photograph was uploaded to the Computer Vision Service using its REST API. On the right, you can see the output. At first glance this might seem like a lot of noise, but take a close look at the description. Did you notice the text which reads “a black and white photo of a city”? Remember, no human was involved in this process - it was a pre-built AI API looking at photograph it received and it came back with that text description. That is pretty incredible. That all it took to leverage this very advanced computer vision AI was a simple API call, that should strike you as amazing.






	<table border="1"> <thead> <tr> <th>FEATURE NAME:</th> <th>VALUE</th> </tr> </thead> <tbody> <tr> <td>Description</td> <td>{ "Tags": ["outdoor", "city", "building", "photo", "white", "large", "background", "black", "sitting", "tall", "skyscraper", "top", "water", "river", "park", "boat", "street", "field", "parked", "computer", "group"], "Captions": [{ "Text": "a black and white photo of a city", "Confidence": 0.958237767 }] }</td> </tr> <tr> <td>Tags</td> <td>[{ "Name": "sky", "Confidence": 0.998601139 }, { "Name": "outdoor", "Confidence": 0.9967468 }, { "Name": "city", "Confidence": 0.9352678 }, { "Name": "white", "Confidence": 0.728594661 }, { "Name": "skyscraper", "Confidence": 0.233053222 }] }</td> </tr> <tr> <td>Image format</td> <td>"Jpeg"</td> </tr> </tbody> </table>	FEATURE NAME:	VALUE	Description	{ "Tags": ["outdoor", "city", "building", "photo", "white", "large", "background", "black", "sitting", "tall", "skyscraper", "top", "water", "river", "park", "boat", "street", "field", "parked", "computer", "group"], "Captions": [{ "Text": "a black and white photo of a city", "Confidence": 0.958237767 }] }	Tags	[{ "Name": "sky", "Confidence": 0.998601139 }, { "Name": "outdoor", "Confidence": 0.9967468 }, { "Name": "city", "Confidence": 0.9352678 }, { "Name": "white", "Confidence": 0.728594661 }, { "Name": "skyscraper", "Confidence": 0.233053222 }] }	Image format	"Jpeg"
FEATURE NAME:	VALUE								
Description	{ "Tags": ["outdoor", "city", "building", "photo", "white", "large", "background", "black", "sitting", "tall", "skyscraper", "top", "water", "river", "park", "boat", "street", "field", "parked", "computer", "group"], "Captions": [{ "Text": "a black and white photo of a city", "Confidence": 0.958237767 }] }								
Tags	[{ "Name": "sky", "Confidence": 0.998601139 }, { "Name": "outdoor", "Confidence": 0.9967468 }, { "Name": "city", "Confidence": 0.9352678 }, { "Name": "white", "Confidence": 0.728594661 }, { "Name": "skyscraper", "Confidence": 0.233053222 }] }								
Image format	"Jpeg"								

Example results from the Computer Vision API

Cognitive Services includes a collection of APIs that provide functionality spanning the categories of vision, speech, language, knowledge and search, enabling the application of AI to text, speech, images, or video and search queries.



VISION

Vision makes it possible for apps and services to accurately identify and analyze content within images and videos.

	Service	Example applications
	<p>Computer Vision</p> <p>Distill actionable information from images</p>	<ul style="list-style-type: none"> Image classification Scene and activity recognition in images Celebrity and landmark recognition in images Optical character recognition (OCR) in images Handwriting recognition
	<p>Face</p> <p>Detect, identify, analyze, organize, and tag faces in photos</p>	<ul style="list-style-type: none"> Face detection in images Person Identification in images Emotion recognition in images Similar face recognition and grouping in images
	<p>Video Indexer</p> <p>Unlock video insights</p>	<ul style="list-style-type: none"> Face detection in video Object, scene, and activity detection in video Metadata, audio, and keyframe extraction and analysis
	<p>Custom Vision</p> <p>Easily customize your own state-of-the-art computer vision models for your unique use case</p>	<ul style="list-style-type: none"> Customizable image recognition
	<p>Content Moderator</p> <p>Automated image, text, and video moderation</p>	<ul style="list-style-type: none"> Explicit or offensive content moderation for images and videos Custom image and text lists to block or allow matching content Tools for including feedback from human moderators



SPEECH




Speech enables the integration of speech processing capabilities into any app or service.

	Service	Example applications
	<p>Speech Services</p> <p>Unified speech services for speech-to-text, text-to-speech and speech translation</p>	<ul style="list-style-type: none"> • Automatic speech recognition and speech transcription (speech-to-text) • Customizable speech recognition and speech transcription (speech-to-text) • Customizable speech models for unique vocabularies or accents • Automatic text-to-speech • Customizable voice fonts for text-to-speech • Real-time translation • Automated speech translation • Customizable translation
	<p>Speaker Recognition</p> <p>Use speech to identify and verify individual speakers</p>	<ul style="list-style-type: none"> • Speaker identification • Speaker verification

LANGUAGE


Language services ensure apps and services can understand the meaning of unstructured text or recognize the intent behind a speaker’s utterances.

	Service	Example applications
	<p>Text Analytics</p> <p>Easily evaluate sentiment and topics to understand what users want</p>	<ul style="list-style-type: none"> • Named Entity Recognition • Key phrase extraction • Text sentiment analysis
	<p>Bing Spell Check</p> <p>Detect and correct spelling mistakes in your app</p>	<ul style="list-style-type: none"> • Web-scale, multi-lingual spell checking • Contextual spell checking

	<p>Language Understanding</p> <p>Teach your apps to understand commands from your users</p>	<ul style="list-style-type: none"> Contextual language understanding
	<p>Translator Text</p> <p>Easily conduct machine translation with a simple REST API call</p>	<ul style="list-style-type: none"> Automatic language detection Automated text translation Customizable translation
	<p>Content Moderator</p> <p>Automated image, text, and video moderation</p>	<ul style="list-style-type: none"> Explicit or offensive content moderation for images and videos Custom image and text lists to block or allow matching content Tools for including feedback from human moderators

KNOWLEDGE


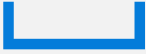
Leverage or create rich knowledge resources that can be integrated into apps and services.

	Service	Example applications
	<p>QnA Maker</p> <p>Distill information into conversational, easy-to-navigate answers</p>	<ul style="list-style-type: none"> QnA extraction from unstructured text Knowledge base creation from collections of Q&As Semantic matching for knowledge bases

SEARCH

Enable apps and services to harness the power of a Web-scale, ad-free search engine with Search. Use search services to find exactly what you're looking for across billions of web pages, images, videos, and news search results.

	Service	Example applications
	<p>Bing Web Search</p> <p>Get enhanced search details from billions of web documents</p>	<ul style="list-style-type: none"> • Ad-free web search • Safe web search • Location-aware web search
	<p>Bing Custom Search</p> <p>An easy-to-use, ad-free, commercial-grade search tool that lets you deliver the results you want</p>	<ul style="list-style-type: none"> • Custom search engine creation • Ad-free custom search results
	<p>Bing Video Search</p> <p>Search for videos and get comprehensive results</p>	<ul style="list-style-type: none"> • Ad-free video search • Video topic and trend identification
	<p>Bing Image Search</p> <p>Search for images and get comprehensive results</p>	<ul style="list-style-type: none"> • Ad-free image search
	<p>Bing Visual Search</p> <p>Get rich insights to help build compelling image applications on the device of your choice</p>	<ul style="list-style-type: none"> • Image identification and classification • Knowledge acquisition from images • Identification of similar images
	<p>Bing Entity Search</p> <p>Enrich your experiences by identifying and augmenting entity information from the web</p>	<ul style="list-style-type: none"> • Named entity recognition and classification • Knowledge acquisition for named entities

	<p>Bing News Search</p> <p>Search for news and get comprehensive results</p>	<ul style="list-style-type: none"> • Ad-free news search results • Trending topic identification
	<p>Bing Autosuggest</p> <p>Give your app intelligent autosuggest options for searches</p>	<ul style="list-style-type: none"> • Search query autocomplete

In the above services that provide pre-built models, Cognitive Services provides APIs that allow for limited customization of the model via training with the customer’s own data.

In our spectrum of pre-built AI and custom AI, these “custom” Cognitive Services fall a little more in between because you need to supply the training data. The custom services includes Custom Vision, Language Understanding, Speech Services, QnA Maker, and Bing Custom Search. Here are some examples:

- The Computer Vision service illustrated above is able to describe the contents of the image and identify some major landmarks, and you could train Custom Vision using sets of images and labels that you

provide to identify objects relevant to your scenario, such as the presence of a helicopter landing pad or HVAC equipment on the rooftop.

- While you could use the Speech Services without any customization to automate order taking at a fast food restaurant, the unpredictably noisy environments might not produce the best results since it is trained for less noisy environments. Instead, by use its customizable speech recognition to train voice recognition in noisy environments, such as the drive-thru of a fast food restaurant.

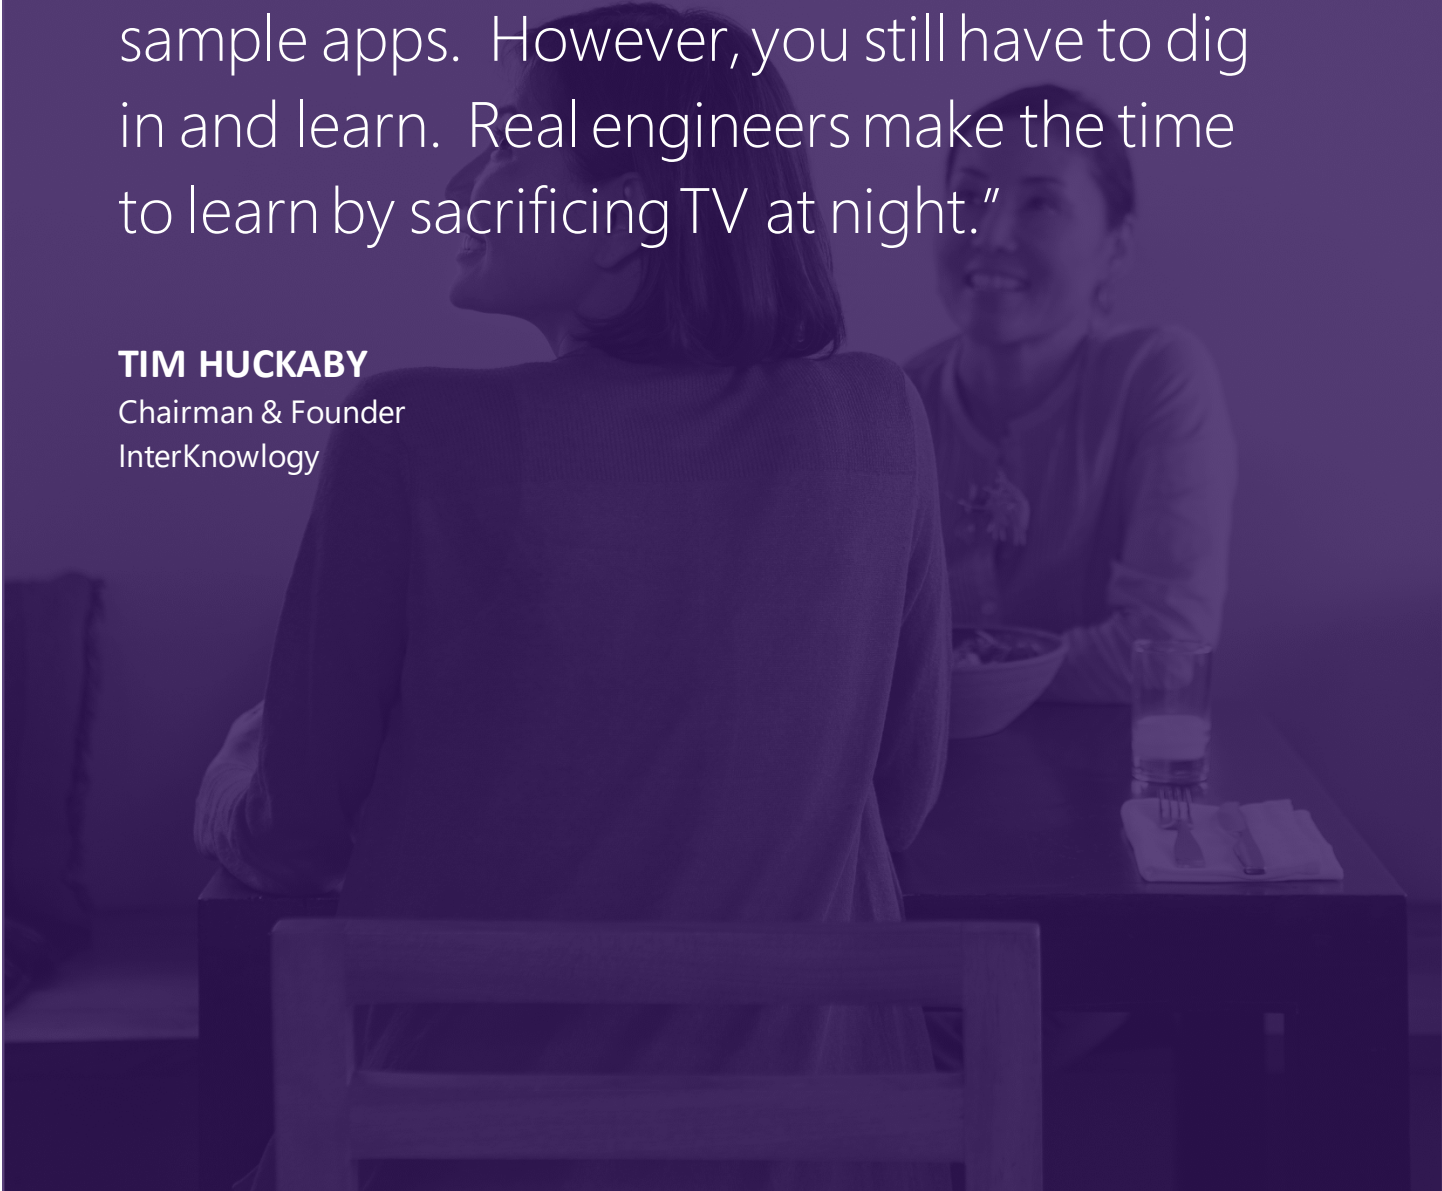
Additionally, some of the custom services like Custom Vision enable you to download the customized model for deployment within mobile apps or at the edge, enabling its prediction capability to be used locally and without an API call back to Azure.

“

“The Cognitive Services team at Microsoft provides all the guidance you could ever ask for to get smart on AI in terms of documentation, SDKs, sample code and sample apps. However, you still have to dig in and learn. Real engineers make the time to learn by sacrificing TV at night.”

TIM HUCKABY

Chairman & Founder
InterKnowlogy





Building Custom AI

The Microsoft AI platform offers a variety of ways to build custom AI and to deploy it into production for use by intelligent applications. Recall from the earlier discussion on the custom AI approach, when building a custom AI, you need to choose where you will build your custom models and then how you will make the model available to application. Let's revisit each in the context of the Microsoft AI platform.

MODEL BUILDING

With the Microsoft AI platform, you have the full range of choices as to where you perform your model building:

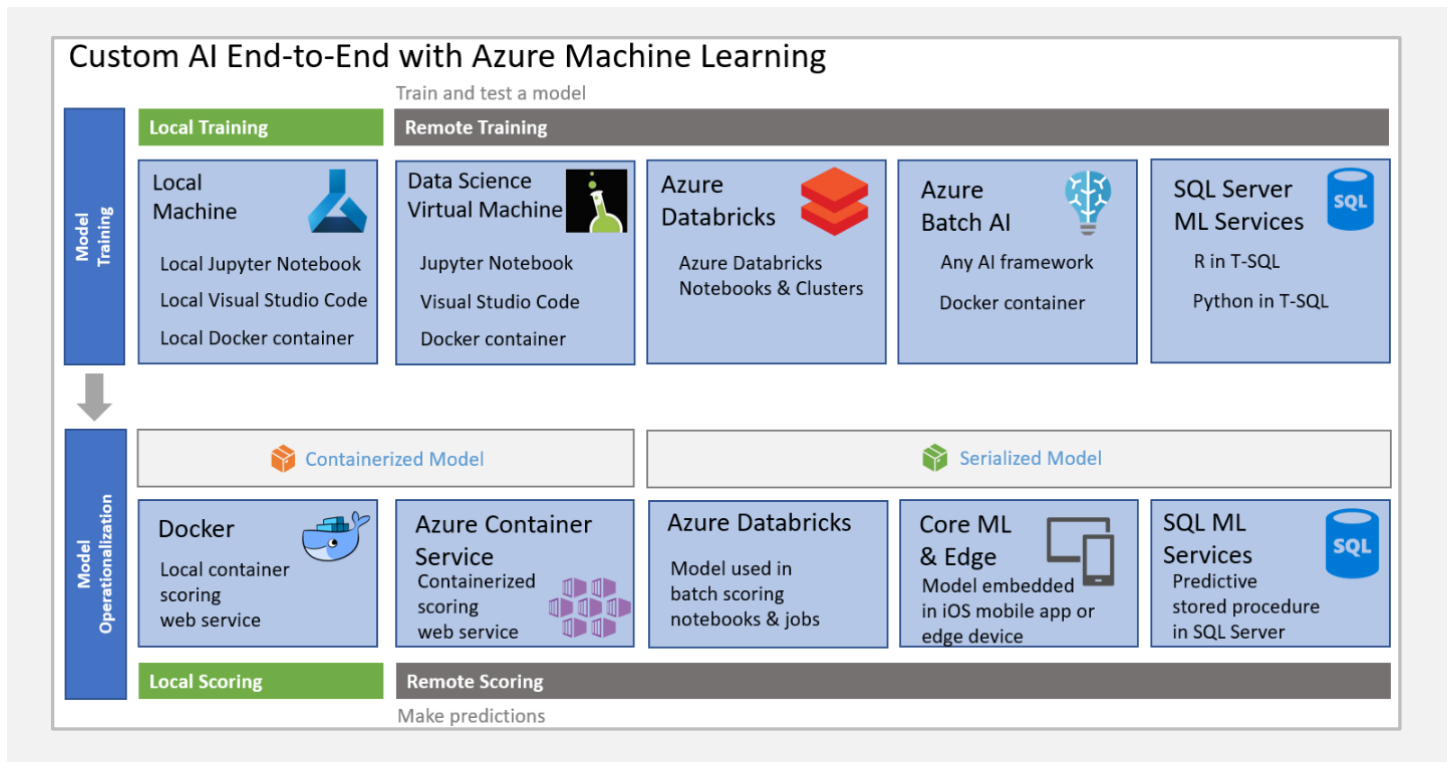
- **Build the model on-premises:** Your data scientists can perform model training on-premises in your local environment using Azure Machine Learning or SQL Machine Learning Services running within an on-premises instance of SQL Server 2016.
- **Build the model in the cloud:** Alternately, they can perform model training in the cloud by using the Data Science Virtual Machine, HDInsight, or by using SQL Machine Learning Services running in SQL Server 2016 in a VM or in Azure SQL Database.

MODEL DEPLOYMENT

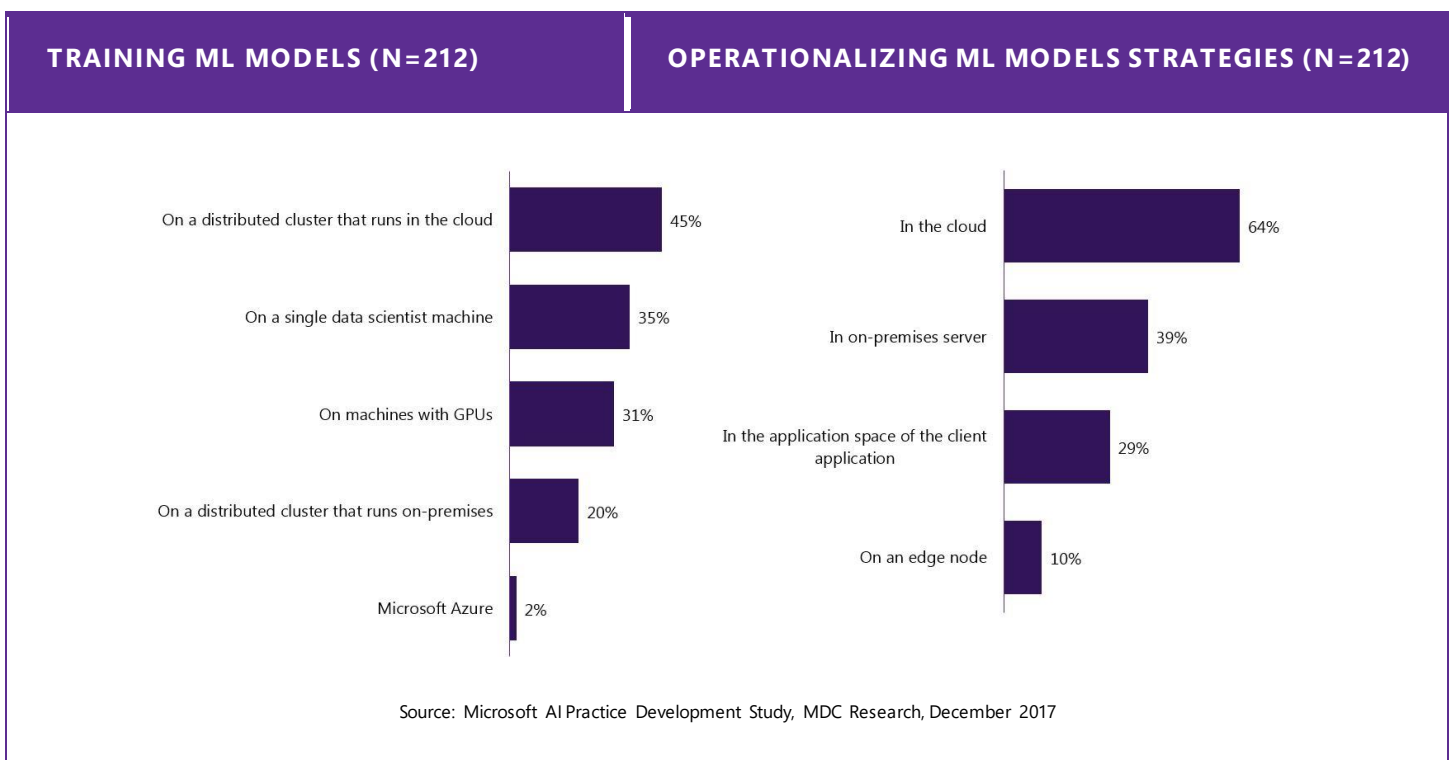
The goal of model training is to produce a model you can use for prediction, also known as scoring. With the Microsoft platform, you have great flexibility in how you can deploy this model and make it available to your applications.

- **Expose custom models as REST services.** You can use Azure Machine Learning to deploy your model as scalable web service that runs within Docker container. Your applications can then invoke a REST API for prediction. Containerized deployments enable operationalization to machines running Docker, to Azure Container Service clusters to IoT Edge devices.
- **Use custom models directly in the context of AI applications.** You can serialize your model using Apple Core ML to embed the model in an iOS app built using Xamarin, enabling the app to make predictions without any network communication.
- **Use custom models in the edge.** Once you have a model created, you also deploy the model to applications running on devices, such as in the Azure IoT edge.
- **Use custom models in the context of SQL database.** You can save the trained model within a table in SQL Server or Azure SQL Database, and enable applications to use it for prediction by invoking a stored procedure that uses the model for scoring.

The following chart summarizes the many ways the Microsoft AI platform allows you to build and deploy your custom AI models.



In the Microsoft AI Practice Development Study, almost half of the partners surveyed indicated they perform model training activities on a distributed cluster that runs in the cloud and deploy their models for use in scoring in the cloud as well.



In the sections that follow, we will explore the key services, compute options and tools from the Microsoft AI platform and how they help you with custom AI.

Build and Deploy Custom AI with Azure Machine Learning

Azure Machine Learning provides services and tools to support the full lifecycle of AI development. From data wrangling tools that are themselves intelligent in the Azure Machine Learning Workbench, to the broad support for experimentation where you build and train your models, to deployment and model management where you operationalize your models, version them and monitor their performance over time. At its core, Azure Machine Learning consists of two Azure services and four development components, each of which is covered in the following sections.

SERVICES

Azure Machine Learning Experimentation Service: handles the execution of machine learning experiments. The Experimentation Service constructs virtual environments to ensure that your script can be executed in isolation with reproducible results. Run history information is recorded in Azure and presented to you visually using the Workbench (covered later). With this history you can easily select the best model out of your experiment runs.

Azure Machine Learning Model Management Service: allows data scientists and DevOps teams to deploy predictive models into a wide variety of environments. Model versions and lineage are tracked from training runs to deployments. Models are stored, registered, and managed in the cloud. The Model Management Service is utilized via a command line interface installed along with the Workbench. You can containerize your model, scoring scripts and dependencies into Docker images.

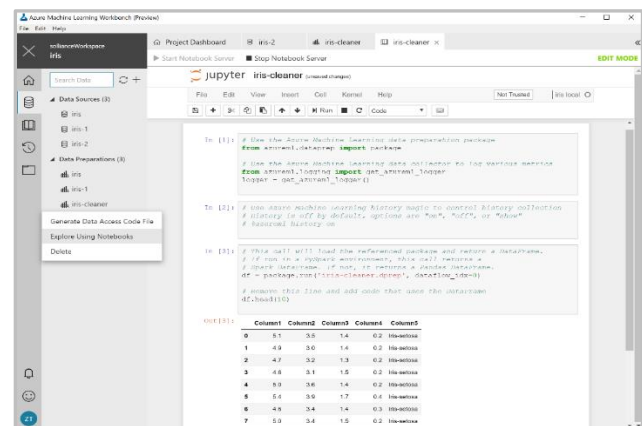
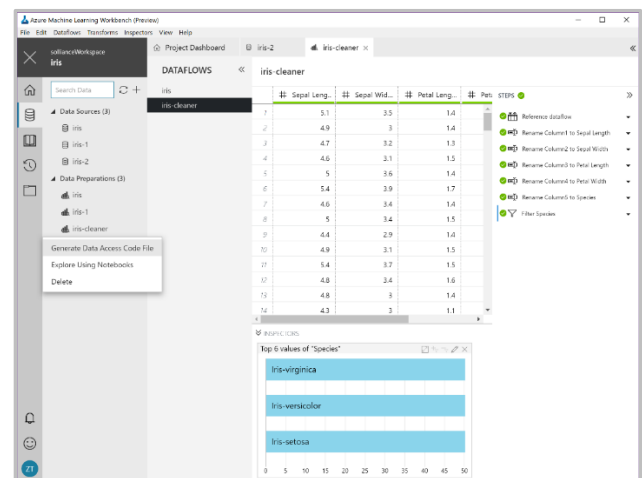
These images are registered in your own Docker registry hosted in Azure via the Azure Container Registry. Once available in a Docker registry, these containers can be deployed to Docker environment, ranging from local machines and on-premises servers, to IoT edge devices and to Kubernetes running in the Azure Container Service. Every time your model executes, telemetry is captured and processed by App Insights for subsequent analysis.

DEVELOPMENT TOOLS & COMPONENTS

Azure Machine Learning Workbench, CLI and SDK: a cross-platform companion desktop app used to wrangle data, build, deploy, manage and monitor models. It includes tooling to build and capture data preparation logic, Jupyter Notebooks for performing experimentation against both local and remote compute environments and tools for monitoring and managing experiment runs. The

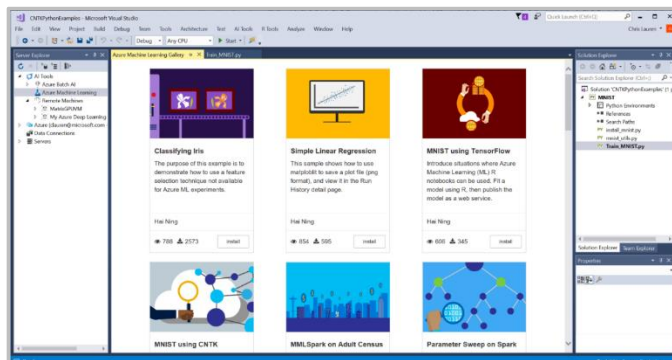
aka.ms/practiceplaybooks

Workbench supports integration with Python IDEs like Microsoft Visual Studio Code and JetBrains PyCharm and also includes command line tools used to manage your experiments and model deployments. The command line interface (CLI) enables you to manage experimentation and perform model deployments from the command line. The Python Software Developer Kit (SDK) enables you to do the same from any environment that can run Python, including Jupyter Notebooks, Azure Databricks notebooks and Azure Notebooks.



Microsoft Machine Learning Libraries for Apache Spark (MMLSpark Library): an open-source Spark package that provides deep learning and data science tools for Apache Spark. It integrates Spark Machine Learning Pipelines with the Microsoft Cognitive Toolkit and OpenCV library. It enables you to quickly create powerful, highly scalable predictive, and analytical models for large image and text datasets.

Visual Studio Tools for AI: Visual Studio Tools for AI is an extension for Visual Studio that supports deep learning frameworks including Microsoft Cognitive Toolkit (CNTK), Google TensorFlow, Theano, Keras, Caffe2 and more. You can use additional deep learning frameworks via the open architecture. Visual Studio Tools for AI leverages existing code support for Python, C/C++/C#, and supplies additional support for Cognitive Toolkit BrainScript.



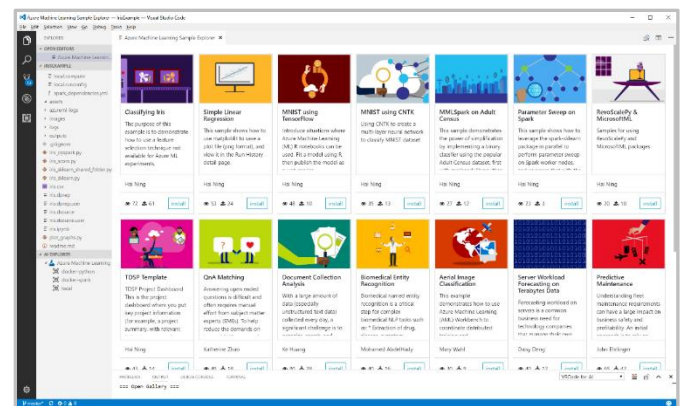
MODEL TRAINING AND DEPLOYMENT WITH AZURE MACHINE LEARNING

With Azure Machine Learning the phase where you work to train a model is referred to as experimentation. The Python code you write to train and evaluate your model performance makes up an experiment. Your experiment code runs in compute resources called compute environments. You can execute your experiments across a range of compute environments including:

- Local native
- Local Docker container
- Docker container on a remote VM
- Spark cluster in Azure for scale out training

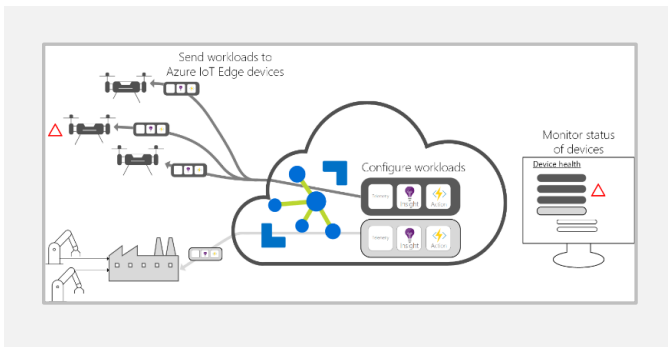
Once you are satisfied with your model, you use the command line tools included with the Workbench to prepare a Docker container that packages your model behind a REST web service. This container is registered in Azure Container Registry, and from there can be deployed to any environment that can run a Docker container.

Visual Studio Code Tools for AI: an extension in Visual Studio Code to build, test, and deploy Deep Learning and AI solutions. It provides a run history view displaying the performance of training runs and logged metrics, a gallery of bootstrap projects with the Microsoft Cognitive Toolkit, TensorFlow, and many other deep-learning frameworks and an explorer view for selecting compute targets for your scripts to execute.



Deploy Custom AI at the edge with Azure IoT Edge

IoT Edge enables hybrid cloud and the Internet of Things (IoT) solutions with a fully managed service that delivers cloud intelligence locally. You can seamlessly deploy and run AI, Azure services, and custom logic directly on cross-platform IoT devices – and manage it all centrally in the cloud with the security of Microsoft.



All machine learning models created with Azure Machine Learning are compatible with IoT edge and these models can be deployed to edge devices, enabling you to bring the capabilities of your Custom AI closer to the data source and without a roundtrip to the cloud.

This enables processing in areas with intermittent, low or no connectivity or in scenarios where there are restrictions preventing the data from leaving the premises. The Azure Machine Learning containers can talk to the IoT Edge runtime and be a part of pipelines on edge devices.

With Azure IoT Edge you can also run certain Azure services at the edge, including Azure Machine Learning, Azure Stream Analytics, Azure Functions and custom code. This can help to reduce IoT costs and simplify deployment, providing flexibility with AI machine learning models.

Build and Deploy Custom AI with SQL Machine Learning Services

SQL Server Machine Learning Services enables you to run, train, and deploy machine learning models using R or Python. You can use data located on-premises and in SQL Server databases or in Azure SQL Database, perform training within the database and deploy predictive capabilities in the form of T-SQL Stored Procedures.

```
CREATE PROCEDURE [dbo].[TrainTipPredictionModel]
AS
BEGIN
    DECLARE @inquiry nvarchar(max) = N'
    select tipped, passenger_count, trip_time_in_secs, trip_distance, direct_distance
    from nyc taxi_features
    tablesample (70 percent) repeatable (98052)
    '
    -- Insert the trained model into a database table
    INSERT INTO nyc_taxi_models
    EXEC sp_execute_external_script
    @language = N'R',
    @script = N'

    ## Create model
    logitObj <- rxLogit(tipped ~ passenger_count + trip_distance + trip_time_in_secs + direct_distance,
    data = InputDataSet)

    ## Serialize model and put it in data frame
    trained_model <- data.frame(model=as.raw(serialize(logitObj, NULL)));

    @input_data_1 = @inquiry,
    @output_data_1_name = N'trained_model'
    ;
    END
GO
```

Owing to this unique co-location of the machine learning models with tabular, SQL Machine Learning Services provides a simplified integration story for developers since they continue to work with the T-SQL Stored Procedures with which they are already familiar. Additionally, because they are implemented as T-SQL Stored Procedures, both access to the Stored Procedure and what the procedure can do against the data are all managed with SQL Server's robust security model.



Build and Deploy Custom AI with Spark and Azure Databricks

Apache Spark is an open-source parallel processing framework that supports in-memory processing to boost the performance of big-data analytic applications and supports a variety of workloads using the same programmatic framework including SQL based querying, stream processing, graph computation and machine learning. Azure Databricks is the premiere place to author, collaborate on, execute and manage Spark workloads. You can, for example, target data stored in Azure Storage blobs or in Azure Data Lake Store or streaming from Event Hubs for processing with Spark in Azure Databricks.

Spark MLlib and MMLSpark in Azure Databricks let you create models as part of Spark jobs that are executing on big data. Spark lets you easily transform and prepare data and then scale out model creation in a single job. Models created with Spark can be deployed, managed, and monitored through Azure Machine Learning Model Management. Training runs can be dispatched and managed with Azure Machine Learning Experimentation.

Use Spark when you need to scale out your data processing and create models as part of a data pipeline. You can author Spark jobs in Scala, Java, Python, R or SQL.

Build and Deploy Custom AI with Data Science Virtual Machine

The Data Science Virtual Machine (DSVM) is a customized VM image that runs in Azure and is built specifically for doing data science. It has many popular data science and other tools pre-installed and pre-configured to jump-start building intelligent applications for advanced analytics. It is available on Windows Server (Windows Server 2016 and Server 2012) and on Linux (Ubuntu 16.04 LTS and on OpenLogic 7.2).

Use the Data Science Virtual Machine when you need to run or host your jobs on a single node. Or if you need to remotely scale up your processing on a single machine. The Data Science Virtual Machine is supported as a target for both Azure Machine Learning Experimentation and Azure Machine Learning Model Management. Additionally, the Data Science Virtual Machine can be used to simulate IoT Edge devices, allowing you to easily test your Azure Machine Learning model deployments in the context of an IoT Edge device.



Build Custom AI with Azure Batch AI

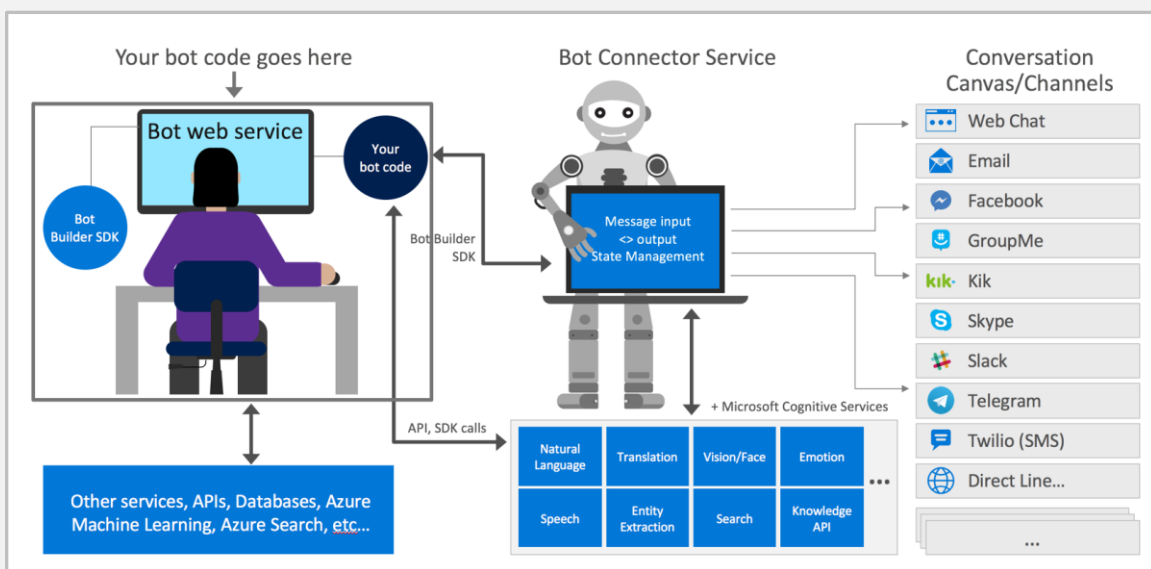
Azure Batch AI enables you to run your AI experiments in parallel, using any framework and perform model training at scale across a cluster of Virtual Machines with GPUs. Batch AI training enables you to scale out deep learning jobs across clustered GPUs, using frameworks such as Cognitive Toolkit, Caffe, Chainer, and TensorFlow.

The Azure Machine Learning Model Management service can be used to take models trained in Batch AI and then deploy, manage, and monitor them.



Add Interaction to your AI with Azure Bot Service & Bot Framework

Bots are fundamentally web services whose logic can invoke the capabilities of your AI, irrespective if it is leveraging pre-built AI from Cognitive Services or Custom AI that you have exposed as a web service. The Azure Bot Service provides an integrated environment purpose-built for bot development. It provides five bot templates you can choose from to quickly get started creating a bot. You can write a bot, connect, test, deploy, and manage it from your web browser with no separate editor or source control required. The bots you build are powered by the Microsoft Bot Framework and Azure Functions. The Bot Framework provides the software development kit (SDK) for implementing your bot in .NET or Node.js, and includes a library of dialogs and prompts that help you structure the interaction between the human user and your bot. The Bot Connector Service works behind the scenes and connects a bot to one or more channels and handles the message exchange between bot and channel. Once in place, your bot enables novel interaction with users that can occur across a variety of channels, including Email, Skype, text messaging and Slack.



Microsoft AI Platform Summary

Create AI solutions with options that include pre-built AI and custom AI.

	PRE-BUILT AI	Automate with computer vision. Use image-processing algorithms to smartly identify, caption and moderate your pictures.	<ul style="list-style-type: none"> • Computer Vision • Content Moderator • Custom Vision • Face • Video Indexer
		Enable natural language communication. Allow your apps to process natural language with pre-built scripts, evaluate sentiment and learn how to recognize what users want.	<ul style="list-style-type: none"> • Language Understanding • Text Analytics • Translator Text • Bing Spell Check • Content Moderator
		Empower speech. Convert spoken audio into text, use voice for verification, or add speaker recognition to your app.	<ul style="list-style-type: none"> • Speech Services • Speaker Recognition
		Leverage the world's knowledge. Map complex information and data in order to solve tasks such as intelligent recommendations and semantic search. Add web search to your apps and harness the ability to comb billions of webpages, images, videos, and news with a single API call.	<ul style="list-style-type: none"> • QnA Maker • Bing Autosuggest API • Bing Web Search API • Bing News Search API • Bing Image Search API • Bing Video Search API • Bing Custom Search API • Bing Entity Search API
+	CUSTOM AI	Prepare data. Understand the data, wrangle data, and identify features most relevant to model training.	<ul style="list-style-type: none"> • Azure Machine Learning - Data Prep • Azure Databricks • Azure Data Lake Analytics
		Training. Train model and evaluate model performance.	<ul style="list-style-type: none"> • Azure Machine Learning – Experimentation • Spark MLlib and MMLSpark
		Deploy trained model. Embed or integrate the trained model into a smart application.	<ul style="list-style-type: none"> • Azure Machine Learning – Model Management • Azure Bot Service
=	COMPREHENSIVE AI PLATFORM		

Define and Design the Solution Offer

With an understanding of the Microsoft AI platform in place, you must next understand the business models of the AI practice because not all revenue streams are equal.

THERE ARE FOUR WAYS TO MAKE MONEY SELLING CLOUD:

- Resale
- Project Services
- Managed Services
- Packaged IP

Partners that focus almost entirely on product revenue have the biggest barrier, and typically see margins in the range of 5–20%. This is because the margins for this revenue line are tied to vendor incentives. These partners are subject to changes in strategy and the desire to fund programs, and have the least control over their own destiny.

Project services typically drive a range of approximately 35% gross margin, but this has been under pressure for some time. This is a result of little differentiation in the channel, which has caused billable price points to hold

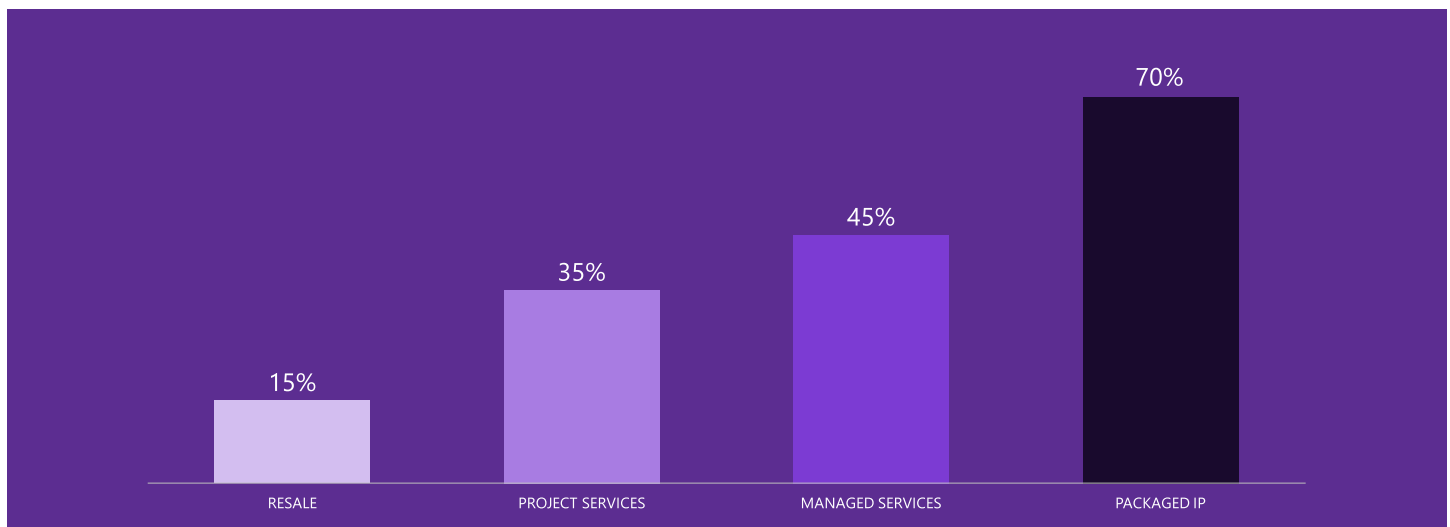
steady over the past five or more years, while increasing salary and benefit costs of consultants and inflation have eroded profitability.

As a result, aggressive and entrepreneurial members of the channel have adapted and gone after the higher margin opportunities of managed services, which generate on average 45% gross margin and packaged IP, which often exceeds 70%.

It is these partners who are setting themselves up to be rewarded. The mergers and acquisition space is quite active. The partners who have gravitated toward the recurring revenue lines and realizing healthy growth are being presented with much higher valuations. This can have a dramatic increase in the cash event of the company and overall shareholder value – far higher than what a traditional partner focused on product and billable services can realize.

A business plan is a critical asset that can help you envision and think through the details of your practice, identify gaps you will need to address, and explain the fundamentals of your practice to others. Leverage the [Cloud Practice – Develop a Business Plan guide](#) for details, profitability scenario overviews, business plan templates, and financial models.

Read on to understand what types of project services, managed services, and intellectual property you should be considering in your AI practice.



Understanding Project Based Services

When building an AI practice, most partners start with selling project-based services.

If you think about the AI practice maturity model introduced earlier, it makes sense that partners establish their AI practice using project-based services. With these early projects, partners augment their domain expertise (in the domain of the customer), assemble their processes and tooling and create a track record of successful deliveries they can leverage in closing future opportunities.

WHAT TYPES OF PROJECTS ARE PARTNERS SELLING?

In the Microsoft AI Practice Development Study, 555 partners that identified as having an AI practice were asked what project services they offer within their practice. The results are below. Observe that the top 5 project services sold were: predictive analytics, proof of concepts, data mining, diagnostic analytics, and data integration.

PROJECT BASED OFFERINGS

Predictive Analytics	47%	Custom ML Algorithm Development	16%
Proof of Concept	41%	Computer Vision	16%
Data Mining	36%	Model Performance Evaluation	16%
Diagnostic Analytics	32%	Model Algorithm Selection	15%
Data Integration	32%	Model Deployment/Operationalization	15%
Data Solution Implementation – SQL Based (e.g., SQL server, SQL Data Warehouse)	28%	Model Application Integration	15%
Strategy	27%	Data Preparation/Wrangling	14%
Chat Bots	25%	Mentoring	13%
Descriptive Analytics	24%	Model Tuning & Re-Training	13%
Assessment and Planning	23%	Data Solution Implementation – Open Source Big Data Based (e.g., Spark)	12%
Prescriptive Analytics	23%	Domain Specific Intelligent Application Development	12%
Visualization Development	23%	Data Integration Pipeline Design	12%
Training	22%	Data Feature Selection/Creation	12%
Data Solution Architecture Design	20%	Batch Scoring	12%
Text Mining	20%	Integrating 3rd Party Pre-Trained Models & APIs	11%
Natural Language Processing	20%	Data Integration Pipeline Implementation	10%
Deep Learning	19%	Agent Development	10%
Real-Time Scoring	18%	Deployed Model Performance Evaluation	10%
Data Quality Evaluation	17%	Device Embedded Models	8%
Model Creation/Training	17%	Ethics Advisory	5%

Source: Microsoft AI Practice Development Study, MDC Research, December 2017.

WHAT DO THE TOP 5 PROJECT SERVICES TELL US?

Let's unpack why the top 5 project services make sense and why you should be considering these for project services as well.

Predictive Analytics & Diagnostic Analytics: Analytic solutions built using machine learning tend to fall in one of four categories. They can be descriptive (describing what happens), diagnostic (explaining how/why it happened), predictive (predicting what will happen), and prescriptive (explaining what to do in response to a prediction). Industry-wide, most effort (and buzz) centers around developing predictive solutions, and partners are focused on the opportunity as well. Prescriptive solutions are less common because they depend on the prediction capabilities to exist, require significant domain expertise, and tends towards very high levels of complexity owing to all the forms of machine learning and deep learning they need in combination to produce their guidance. Partners seem to be cautiously approaching prescriptive analytics because this complexity makes it hard to deliver successful outcomes, and there is tremendous risk in overpromising the capabilities of AI (as the [failure of IBM's Watson](#) in healthcare famously demonstrated).

It is interesting to observe that diagnostic analytics appears in the top five type of project services that partners are delivering. Diagnostic analytics typically have less complexity than predictive analytics. Given the market buzz and excitement around predictive analytics, that diagnostic analytics is being commonly sold by partners is a testament to the value that is produced even in simpler forms of analytics- automating understanding by being able to explain why things occurred without requiring a human to perform the interpretation.

Data Mining & Integration: Good AI has good data as a pre-requisite. Without being able to collect and aggregate the data (data integration) or process and explore the data to understand the insights it contains (data mining) it is not possible to train the learning algorithms that are powering AI. That partners are selling data mining and integration is also indicative of how they started their AI practice – as an evolution of their data practice.

Proof of Concept: As we will discuss in an upcoming section, a Proof of Concept (PoC) is a critical tool in selling AI. The benefit of the PoC for the partner is that it allows them to focus on delivering the core AI value in smaller/shorter engagements, and for the customer it makes AI "real"—it enables them to gain confidence that they can realize the promise of AI in the context of their business.

WHAT OTHER PROJECT SERVICES SHOULD YOU CONSIDER IN YOUR AI PRACTICE?

In our partner interviews, partners emphasized three very specific services for AI that have a lot to do with the evolution of the sale. In chronological order they are:

- **Envisioning:** Envisioning sessions help the customer understand the opportunities for AI in their business requirements, while relying on the partner to keep the conversation grounded in what is possible versus what is hyperbole.
- **Proof of Concept:** Some partners indicated a preference to only provide Proof of Concepts as the first engagement, to help the customer get comfortable with the capabilities of both AI and the implementing partner and to enable the partner to understand the real situation with respect to the availability of quality data and the actual feasibility of solving the problems they are attempting to solve.
- **Pilot:** Other partners indicated a preference to sell pilot projects and effectively start all projects with the notion that the solution implemented would ultimately land in production.

We will explore each of these project services in the sections that follow.

Deliver an Envisioning Session

An envisioning session is a common agile methodology practice that is particularly important for AI engagements as it is up to the partner to help the customer realize the opportunities for AI in the customer’s problem statement.

The challenge is few customers will know if their problem represents a good opportunity to apply AI, or even if AI is the correct approach in their situation. As a partner, you need to have a discussion with your customer about their problem and be on the lookout for opportunities to build the solution using AI. You need to recommend the application of AI when it is appropriate, as well as discourage its application when AI would not be successfully applied in the near term. For example, the customer may want to build a predictive solution, but the labeled training data does not exist and would take years to collect. This type of high level requirements conversation is called an envisioning session.

The outcome of an envisioning session is a common vision with your customer on capturing what may be achieved, the very high-level mechanism by which it will be achieved and the potential value of achieving this vision.

The envisioning session is not intended to be an in-depth, big requirements up-front requirements gathering event that takes weeks or months to complete. Instead, an envisioning session is something that can be conducted in 1-3 days depending on the complexity of the business scenario, and how much explanation the partner team needs from the customer about their scenario.

During the envisioning session you will begin by identifying the customers desired future state. On your way to defining how the vision is achieved and the value it would bring to the business, you will discuss the problem statement, the business and technical capabilities desired as well those that are available and the existence of supporting assets (such as data sets).

The benefits of conducting an envisioning session are:

- You can answer fundamental business questions of what you will build, and a general sense of how.
- You will have improved productivity on the project having identified and thought through the critical business issues facing the project.
- You will have identified the overall business direction required by your architecture.



AT THE END OF THE ENVISIONING SESSION:

- The customer will have a vision of how your AI practice can help their organization realize their goals and drive business impact. They are excited by the value unlocked in the application of AI.
- You will have sufficient knowledge to write a proposal for an engagement.
- You will have likely identified additional opportunities to apply AI to the customer’s scenarios, and have thus already started identifying additional scope for future engagements.



HOW TO DELIVER AN ENVISIONING SESSION

To deliver an envisioning session correctly means to have a loosely structured conversation (ideally in front of a white board) where you brainstorm with the customer about each of the following core questions. It is entirely possible to iterate multiple times through these core questions, each time refining details. Remember this is a high-level conversation.

1. What is the customer’s desired future state?
2. What is the problem the customer is trying to solve?
3. What are the customer’s business capabilities with respect to the problem? For example, do they have experience in the domain for which they are trying to enter?
4. What are the customer’s technical capabilities with respect to the problem? For example, do they have technical resources who have worked with the data in the domain? If they had a solution in hand, could they operate it?
5. What are the supporting assets? For example, do they have the requisite historical data upon which to train predictive capabilities?
6. How will they know when the vision is achieved?
7. What is the value of achieving the vision?

WHAT AN ENVISIONING SESSION IS NOT

An envisioning session is not:

- An architecture design session (discussed later).
- A project planning effort.
- A technology selection effort.

If you find that your envisioning conversation with your customer is headed in either of these directions, pause. Evaluate if you have suitably addressed the core questions. If the core questions have been addressed, then you should formally conclude the envisioning session and allow time for your team to process the input and return with recommendations on next steps (which could be a proposal for a design, a PoC, or a pilot).

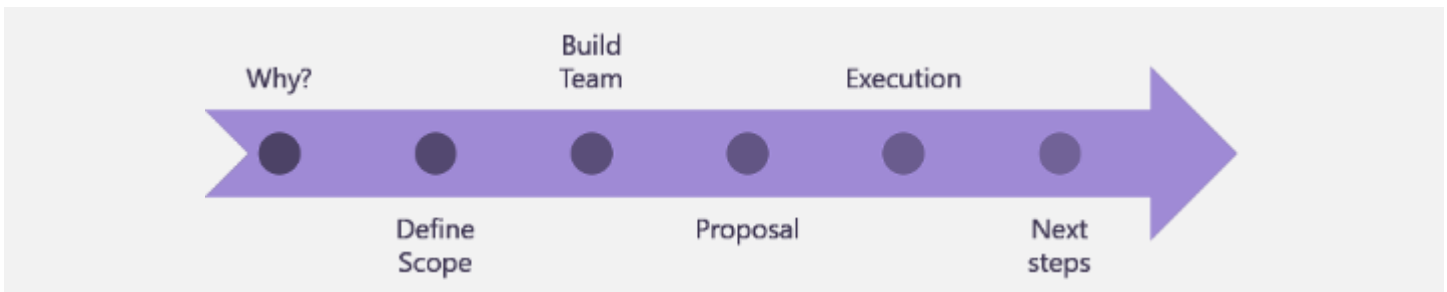
If not, then you should guide the conversation away from getting too much into the architecture, technology selection or other implementation details.



Deliver a Proof of Concept AI project

AI projects include significant risks – for example, the data required may not be available or the problem may not lend itself to accurate prediction. Reduce the risk of overpromising on AI capabilities by conducting a focused proof of concept that enables you to de-risk the riskiest elements of the larger AI solution, build trust with the customer and deliver working AI solutions in a shorter timeline.

PROOF OF CONCEPT HIGH LEVEL FLOW






Why Perform a Proof of Concept for AI?

Proof of concepts (PoC) serve several purposes. When it comes to the AI practice, a primary aim of the PoC is to substantiate that an AI solution can actually deliver on the vision capture during the envisioning session. The intent is to avoid skepticism on the part of client about the capabilities that can be realized and to detect situations over-promising on the capabilities of AI before any significant investment occurs.

The PoC also can serve as evidence that your practice can use for future engagements with the same customer or with new customers. Many times, the output of a PoC can be added to your practice's intellectual property list for demonstrations or used to accelerate future solutions. PoCs are one of the key tools when trying to displace the competition by rapidly showing value and hopefully a quick return on investment.




PoC Execution

Identify the technical resources needed for the PoC. This will include the technical implementation team (including developers and data scientists), as well as project management for tracking the progress of the engagement.

Beyond identification of resources, ensure all members of your team and your customer's technical team (if they are participating) are clear on responsibilities. Communication is critical, so ensure that the progress of the PoC is communicated to all stakeholders on a regular basis.


During the PoC execution, keep a watchful eye towards scope creep. The PoC should be hyper focused on proving that an AI solution can be delivered as envisioned by addressing only the core concerns originally identified during the scope definition. For example, most AI PoCs will want to attempt the solution first with pre-built AI, before going down the path of building out the custom AI. Avoiding the temptation to start with the more complex solution (such as building a custom AI) is a good strategy for keeping the minimalist scope required of a PoC.



Define Scope

A proper PoC is defined with a clear and concrete scope. Conduct an envisioning session or an architecture design session (ADS) to align business and technical requirements and set clear goals. This should include:

- Identify the AI workloads and features to demonstrate. Specifically identify and prioritize the aspects that are high risk.
- Identify the data that will be used to train any underlying AI models.
- Determine what you want to prove and which objections need to be overcome.
- Clearly demarcate responsibilities and set up organization.
- Set up subscriptions, define payment, and perform cost estimates of the PoC.
- Agree on the next step if success criteria are met.



Next Step

At the end of the PoC, create a report that explains the overall status of the PoC and any issues identified during the PoC. The report should elaborate on the pros and cons of the delivery and clearly explain the value prop of moving forward with a real implementation to the stakeholders along with expected production costs over time. Assuming the stakeholders agree to move forward, put a plan into place to deploy the PoC into production while ensuring that the PoC is designed for production usage.

For example, upon completing your AI PoC you might report on the success or failure of utilizing a pre-built AI, the quality and predictive strength of the customer supplied data used for training a custom AI, or the feasibility of realizing the customer's vision.



Examples of an AI PoC

EXAMPLE VISION	EXAMPLE POC
<p>We want to increase sales by getting shoppers to commit to a purchase they have been mulling over. If we can detect this behavior we'd like to provide just enough of a discount to encourage the purchase</p>	<p>Focus on the core aspect of the scenario (predicting a discount that leads to a purchase) and demonstrate the PoC in a notebook or other approach that minimizes efforts in building UI. The PoC scope would include analyzing the client's purchase history and building a custom AI identifying the right predictive algorithm to suggest a discount within a bounded range that has a high probability of sale. By conducting this PoC you will have verified if the customer has the necessary data to make such predictions.</p>
<p>We would like to add chat bot capabilities to our solution that could address commonly asked questions before needing to speak to a live operator.</p>	<p>Utilize a bare-bones chat bot UI and provide the customer an initial understanding of the experience by applying pre-built AI to help in turning assets like frequently asked questions into conversational AI. Avoid getting into custom AI in this PoC. This will help the customer get a sense for the type of interaction that is possible and how the computer to human handoff might work, without significant investment in developing a custom AI.</p>
<p>We want to monitor how happy people are during their chat-based interaction with our agents in real time. We are concerned that we can only do this after the fact, at which point it is too late.</p>	<p>Utilize pre-built AI to apply sentiment analysis to historical chat message flowing through a minimalist chat environment, which demonstrates various alerts appearing around the chat messages when sentiment thresholds are crossed. Focus on the speed of processing and the accuracy of the sentiment established. By utilizing pre-built AI first in the PoC you can identify if the vocabulary and patterns of speech are suitably understood by generic sentiment analysis before investing effort in building custom chat histories manually tagged with sentiment and used to train a custom AI.</p>

Deliver a Pilot for an AI project

A PoC might secure the project, but a pilot is what you will deliberately take to production.

Consider offering your services as a pilot project to your prospects. With a pilot project, the customer receives two important values. First, they get to better understand how the project goals will be successful, and second, they have a production-grade starting point for their larger efforts.

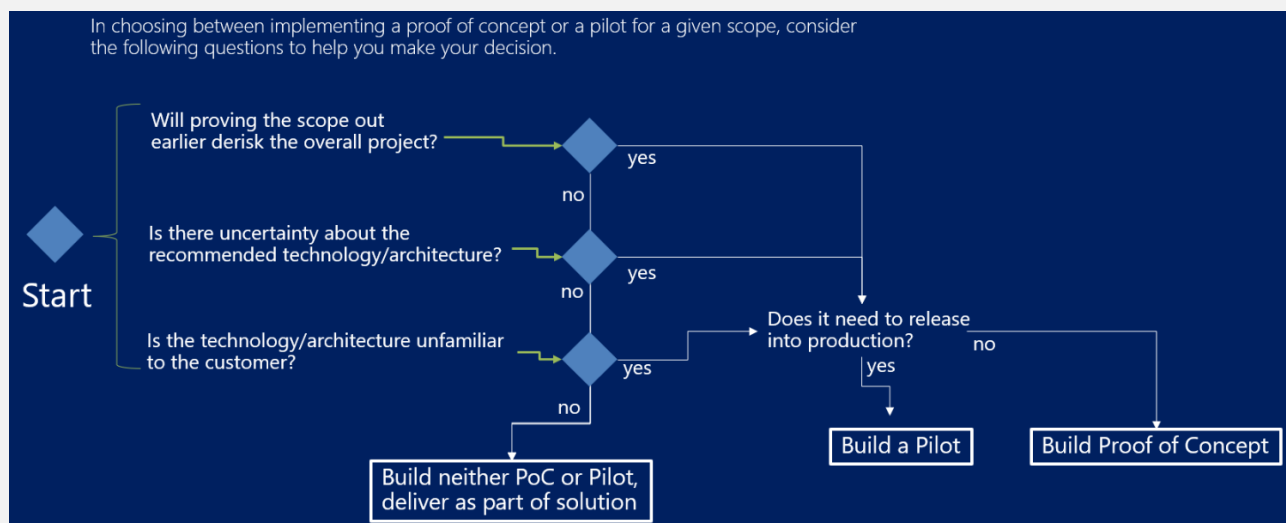
The AI partners interviewed for this playbook highlight another value of the pilot for the AI practice – it serves as a way for your implementation team (developers and data scientists) to stretch into new areas (learning the customer’s domain, learning to apply different predictive algorithms, working with unfamiliar data), while minimizing risk if something doesn’t work out quite right on the first attempt. This is because while the intent of a pilot is to ultimately take the solution into production, the initial pilot delivery is never rolled out directly into production, at best it usually affects a subset of the production environment which helps to minimize the impact of unexpected issues.

CHOOSING BETWEEN A PILOT OR POC

It is important to keep the distinction between Proof of Concept and Pilot clear – a PoC should never be considered for direct deployment into production, whereas a pilot should be constructed with a production release in mind.

For example, in an AI PoC you may not even touch the customer’s actual data set and instead use similar data from open data sets to show what the value of the predictive capabilities and to tangibly illustrate how they would be applied in the context of a running solution. In a pilot, however, you would want to start with the customer’s actual datasets because the goal is to end up with a model, that if successful, would flow into the production solution.

A pilot solution is a production-ready product whose influence is limited in scope (targeted rollout), customer base, or capacity. A well-executed pilot will give the customer a better understanding of how the project goals will be successful, while providing them with a production-grade starting point. Since a successful pilot will be scaled up to the final production solution, it is important to create the pilot following best practices.





Start the pilot design process with a clear plan on its initial scale, and develop a strategy to increase its scale as milestones are met.

For instance, if the intent of the pilot is to onboard a small portion of the customer base at first, identify the participants early and add your communication strategy to the project plan. Consider targeting users who are representative of the whole of your customer base. This should provide you with an accurate test of geographic, technological, and demographic factors.

A TYPICAL PILOT UNDERGOES THE FOLLOWING PHASES:

- Define scope – typically occurs during an architecture design session (ADS)
- Execute implementation – create, test, refine, repeat
- Conclude – post-mortem, validation, path to production plan

DURING THE PILOT CHECKLIST

The following checklist provides the core tasks you should complete when conducting a pilot.

- Go over established business and technical requirements from the detailed requirements gathering session (e.g., architecture design session)
- Determine which features will be included in the pilot that provide a minimum viable product
- Conduct a full-fledged design, mapping requirements to workloads and features of the pilot
- Establish team responsibilities and organization
- Perform cost estimates (e.g., for Azure services used like Cognitive Services, Machine Learning, etc.)
- Outline next steps after the success criteria is met

AFTER THE PILOT CHECKLIST

Throughout the pilot, you have collected valuable data from metrics and telemetry, and have compiled user feedback and taken this information to refine the pilot and prioritize features for general release. The next step is to propose the move to production and provide a cost estimate and delivery schedule to the customer. You should perform the following in concluding a pilot:

- Learn to identify the end of the pilot; when the proposed features are functional and validated, and when the business transformation begins
- Conduct a post-mortem of the pilot, talking to users, business managers, developers, and development managers. Evaluate lessons learned and refine feature list to include any features excluded from the pilot
- Upon agreeing to move to production phase, establish cost and delivery schedule based on established plan to transform to v1 product

“

A key element to a successful AI delivery is blurring the lines between customer and partner...in an open space have the customer come in and work directly with your team. When you become so integrated that it is hard to tell who is the customer and who is the partner, you know your AI will have the right balance of domain and technical expertise.

VINCENT THAVONEKHAM

Cloud Azure Strategy Manager
VISEO Group






Understanding Managed Services

With managed services, you can help your customers on a regular basis by offering white-glove services wrapped around your AI solution. Your offerings can span from planning, to enablement, and to day-to-day operations and support.

Managed services is not a new business model. For more than 20 years, large enterprises have relied on service providers to manage their IT assets on their behalf. Whether you call them an outsourcer or a managed IT provider, service providers have been managing their customers' workloads — either in their own data centers or those operated by their customers. Managed services create opportunities for partners building new lines of business to provide a white glove experience for the entire AI solution running on-premises, in the cloud or as a hybrid solution.

Managed services give you an alternative option to selling the time of your AI practice team for money (as you might in project services). Becoming a managed services provider (MSP) enables you to take the IP that is almost always created in the process of delivering an AI solution (such as insights into a domain, the data providing the greatest predictive capability, the algorithms which have performed best, techniques you invented to keep a model performant in production, etc.) and package the IP with services that the customer will pay for on a subscriptions basis. You can then sell that same set of IP plus services to other customers needing a similar solution without repeating all of the IP creation effort, at a profit.

 <p>PLANNING</p>	 <p>ENABLEMENT</p>	 <p>SUPPORT OPERATIONS</p>
<ul style="list-style-type: none"> • Help the customer envision scenarios where AI could amplify human ingenuity to achieve positive outcomes. Consider the opportunities that surface with reasoning, understanding and interacting. • Identify the “low-hanging fruit” opportunities that can be used to build confidence in your services and in the solution, such as those that utilize pre-built AI. • Build a roadmap that shows the path to the larger opportunities, building on the success of the smaller or more near-term opportunities. 	<ul style="list-style-type: none"> • Design and implement the AI solution with the customer. • Engage the domain expertise of the customer throughout the process to capture their insights and to validate the AI. 	<ul style="list-style-type: none"> • Offer further support while delivering on SLAs and uptime guarantees. • Operate and monitor the solution. • Monitor the performance of any machine learning models. • Re-train machine learning modes.

So what managed services can your AI practice offer? We'll examine that in the next section.

Managed Services for an AI Practice

The AI partners interviewed for this playbook suggested that when ideating about what to offer of your practice as managed services, that you might first consider the project services you are offering. With some creativity, the project services discussed earlier are all potential offerings for managed services.

According to the partners interviewed for this playbook, there is a **significant and unique opportunity for partners looking to offer AI as a managed service**. Moving from project services to managed services will help your AI practice create annuity income streams with higher professional services margins, increased customer loyalty and the increased revenue that naturally follows.

Why move from project services to managed services for AI services?

The low hanging fruit for most practices to offer managed services is to provide support for the solution delivered via project services. If you have an AI practice, you might think you do not want to be in the business of technical support. However, consider what happens when your customer takes your incredible predictive solution into production. Perhaps for an initial period, everything is working as intended and you do not hear from the customer at all. Then one day you get a frantic call from the customer because the underlying model is “not working” and neither the developers nor the IT professionals at the customer site know where to begin in troubleshooting the issue. This is the opportunity to provide AI support as a managed service.


Why? Given the shortage of data science capabilities, it is likely that AI solutions will be delivered to customers who do not themselves have internal data science capabilities. This means that for the customer’s long-term success, your practice is likely to be involved in perpetuity in supporting the solution in production. Would you rather scramble to assemble the team to support the customer’s panicked call once the team capable of resolving it has moved on to other projects, or would you prefer a controlled and organized response where you have already transitioned the solution knowledge to a managed services support team who is on standby and is ready to support the customer?

HOW TO BUILD AN AI MANAGED SERVICE

For even the most sophisticated software house, effectively measuring performance, handling re-training and patching, and staying ahead of the rapidly evolving AI landscape can be too difficult to manage without help. Partners can offer their services on a subscription basis to ensure the production AI solution continues to deliver the value and performance that got the customer excited at using AI in the first place.

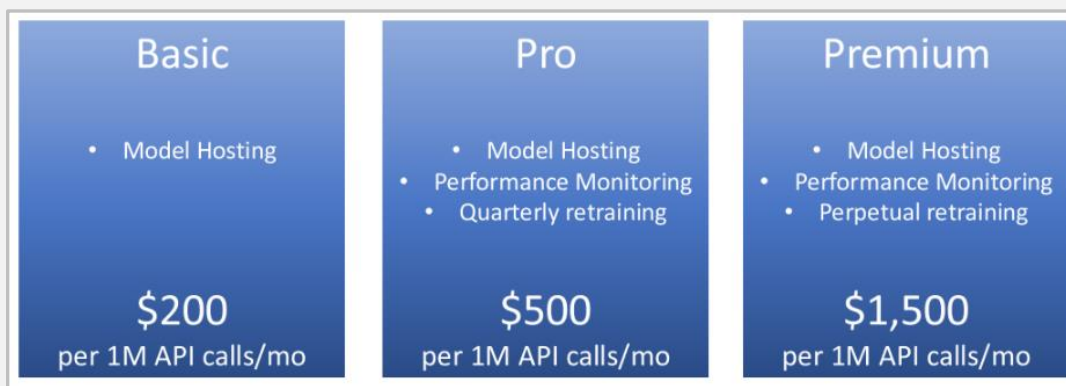
KEY CUSTOMER CHALLENGES

1. They lack the tools and expertise to effectively monitor the performance of an AI solution.
2. They are unable to identify, assess, and troubleshoot issues in production deployments.
3. They may not have the internal teams that are able to tune the model or retrain when the data semantics have shifted.



EXAMPLE OF AN AI MANAGED SERVICES OFFERING

Given these challenges, there is a clear opportunity for partners to package the hosting of the delivered AI solution along with support of the solution. For example, if your deployed solution uses Azure Machine Learning, the costs for Experimentation, Model Management and the Azure Container Service Cluster used collectively in delivering the predictive REST API are baked into a graduated monthly fee the customer pays, which may also include a limited set of data science services needed to maintain the model in production. In the following example, the quantitative value the customer gets is measured in units of millions of API calls, but qualitatively the customer is also securing access to partner data science resources to ensure the model’s performance remains at the desired levels. By offering the customer their model as managed API, the customer is completely able to avoid dealing with both infrastructural and data science challenges – instead, they can focus on harnessing the value created by integrating the API in their solution.



What are some concrete example of managed service offerings your AI practice could sell? In the Microsoft AI Practice Development Study, 555 partners that identified as having an AI practice were asked which managed services they offered within their practices. The results are below. Observe that the top 5 project services sold were: visualizations, dashboarding, report creation and maintenance, support, assessment and planning, analytics as a service, and troubleshooting.

MANAGED SERVICES SOLD BY AI PARTNERS

MANAGED SERVICES

Visualizations, Dashboards and Reports Creation/Maintenance	43%	Data Science as a Service	23%
Support	43%	Domain Specific Services	22%
Assessment and Planning	38%	Model Performance Monitoring	20%
Analytics as a Service (e.g., packaged APIs for ML models and agent interactions)	27%	Model Tuning & Re-Training	19%
Troubleshooting	26%	Model Hosting (e.g., as Predictive Web Service)	18%
		Online Training and Self-Paced Learning	13%

Source: Microsoft AI Practice Development Study, MDC Reserch, December 2017

WHAT DO THE TOP 5 MANAGED SERVICES TELL US?

Let's unpack why the top 5 managed services make sense and why you should be considering these for managed services as well.

Visualizations, dashboarding, report creation and maintenance: Without being able to demonstrate concretely the value that your AI solution is delivering on an on-going basis, it can be difficult for a customer to justify keeping up a subscription. Even when the AI solution is delivering the value they are paying for, customers get "accustomed" to their new environment enabled by your AI solution, unless you consistently remind them of the value it is delivering in the form of visualizations, dashboards and reports. Additionally, once your AI solution is embedded your customer's routine, customers will be quick to identify new visualizations, dashboards and reports because your solution will have enabled them to ask new questions. For many customers, this is a never-ending desire and you should meet this need by offering it in an on-going manner, as a managed service.

Support and troubleshooting: As will be discussed further in an upcoming section, support represents a very compelling managed service for an AI practice. While it may sound strange for an AI practice to offer support, the value of support is quickly understood when you include in your consideration the shortage of talent present in the market. If your customer does not currently have staff with AI experience, such as data scientists, then who will be there to help them monitor the performance of their predictive models, help them detect when models have become invalidated and need retraining, help them diagnose why the AI is no longer responding as expected or demonstrating an unexpected bias in its responses? Just as when customers looked to managed services partners to relieve them of the need to be experts in hosting and running data centers, so too will they be looking to AI partners to help them operate their AI in production.

Assessment and planning: Previously we introduced the notion of performing an envisioning session with customers to help them create a vision for a problem that could feasibly be addressed with AI. These are not exercises in "blue sky" thinking, they require partners skilled in AI to help temper brainstorming with the art of the possible. The need for you to help with the effort from vision to current state assessment to an actionable roadmap requires a similar expertise, and the customer is not likely to have this skill in house. As customers take on more initiatives that include AI in the solution, they will need a trusted advisor on an on-going basis to help them with tasks like assessing if the organization has the pre-requisite assets to address the vision (domain expertise, historical data, etc.), ensuring the plans can feasibly realize the vision or that plan once realized is in-line with ethics of the organization.

Analytics as a service: That analytics as a service appears as a top 5 offering for AI partners should remind you of the shortage of talent available in the marketplace. If you have the talent in-house to deliver advanced analytics solutions (e.g., using AI to reason about or understand the insights in data) and your customer does not have these resources, a good long term partnership could be in the making, either in providing access to your experts or the packaged solutions they have produced.

WHAT MANAGED SERVICES SHOULD YOU CONSIDER IN YOUR AI PRACTICE?

In our partner interviews, partners emphasized support and monitoring specifically as managed services to consider in an AI practice. We will explore each of these managed services in the sections that follow.

Support as a Managed Service in an AI practice

Support represents a unique managed services opportunity for AI partners.

Unlike the support you might expect to provide for a SaaS-based web application or web service in production, the support provided for an AI solution in production is different, as is the reason customers will want AI practices to include support as a managed service along with the delivered solution. At its core, the difference lies in the skillset required to support the uniquely AI and machine learning or deep learning heavy aspects of the solution, coupled with the reality that AI expertise is in short supply. This means that it is highly likely that your customers will be looking to you for help when something goes wrong with their solution in production, because they are not likely to have the resources to support the solution in house.

This is not to say that the support for an AI practice would support only the AI components, since no matter how well a cloud or hybrid solution is planned, provisioned, operated or monitored, problems will arise, and those problems will need to be remediated. It's your job as an MSP to offer support to your customers to deal with outages, breaches, inefficiencies, and disaster scenarios. AI simply adds additional items your customer will need support with, such as monitoring and maintaining model performance of the underlying machine learning or deep learning models, identifying when trends in the current data diverge from those in historical data and necessitate re-training of the model or diagnosing why an AI solution is seemingly displaying an undesirable bias.

MSPs need to consider the level of support that makes sense for their practice — in terms of resources and revenue — as well as what makes sense to the customers they serve.

KEY CUSTOMER CHALLENGES

- They lack the expertise and resources to troubleshoot problems.
- They are unable to determine the root cause of performance issues and glitches.
- They have no knowledge of how to remediate problems when they correctly identify them.
- They do not want to spend time and resources fixing problems.

KEY SERVICES FOR THIS OFFERING

- **Model Performance Support:** Provide support around monitoring and maintaining the performance of an AI solution and its underlying models in terms of the accuracy, currency or reasonableness of its output. Assist the customer specifically in diagnosing what has changed with respect to the data environment, the model or other factors and providing support in resolving the issue.
- **User Support:** Provide support for frequently asked questions, setup and usage, best practices, questions around billing and invoicing, break-fix support for developers, architecture design, and solution design support for architects.
- **System Support:** Provide customers with information on any service interruption, and relay expectations on when the system will be back online.
- **Product Support:** Provide support when the Microsoft product is not working as expected or the service stops working. Escalate to Microsoft when the issue cannot be resolved with existing documentation and/or training.
- **Extended Support Hours:** Many customers need the ability for 24/7 support, but cannot justify the overhead internally.
- **Account Management:** Offering an account manager that is responsible for reporting on service consumption and ultimately minimizing time to resolution is a service that can be offered at a premium.
- **Dedicated Support:** The value add of a dedicated support team cannot be overstated. Engineering resources that already know your customers' environment, including the business and technical reasons for how a solution was implemented can add a tremendous value over the lifetime of an agreement.

Monitoring AI Services

In the AI world, the tools and requirements have evolved, but the problem statement has not fundamentally changed. How do I monitor the health and performance of my infrastructure, inclusive of the models? There is no easy answer to this and customers expect their service providers to solve it for them. Most mid-market and enterprise organizations simply do not have the time, resources, or dedicated staff required to monitor every aspect of the solution, and this is where MSPs add the most value. While Azure offers many monitoring capabilities built within the platform to monitor the application related services, there is still a place for partners who (a) provide additional, deeper monitoring tooling that takes into account the health and performance of the AI, (b) triages the false positives from the real alerts, and (c) proactively acts upon the alerts before any measurable loss in performance.

KEY CUSTOMER CHALLENGES

- I don't have the time or resources to monitor all the components in my AI solution.
- I need a single pane of glass view that tells me how all my apps and AI models are performing, at any point in time.
- I find it challenging to diagnose the root cause of breakdowns, outages, or unexpected bias.
- How do I respond to so many alerts? How do I differentiate the false positives from the concerning ones?

KEY SERVICES FOR THIS OFFERING

The following table illustrates how a partner might construct a comprehensive AI solution monitoring offering, which includes AI performance monitoring.

SYSTEM HEALTH MONITORING	LOG ANALYTICS AND ALERTING	DATABASE MONITORING	APPLICATION PERFORMANCE MONITORING	AI PERFORMANCE MONITORING
<p>Complete monitoring of VMs, CPU utilization, memory usage, storage IOPs, and OS performance.</p> <p>Includes monitoring of application performance and operation health, and dashboards and reports on system health.</p>	<p>Every client, device, and user accessing a network produces data that is logged.</p> <p>Analyzing those logs can offer deep insight into performance, security, resource consumption, and other meaningful metrics.</p>	<p>A view into your customer's database that helps MSPs ensure high availability of database servers.</p> <p>The process involves keeping logs of size, connection time and users of databases, analyzing use trends, and leveraging data to proactively remediate issues.</p>	<p>End-to-end tracking of all aspects of an application (or webpage).</p> <p>App monitoring involves watching every part – from shopping carts to registration pages – of a customer's app(s) for performance issues in an effort to provide the best user experience possible.</p>	<p>Perpetual monitoring and evaluation of model performance in terms like accuracy, reasonableness and desired bias.</p> <p>Monitor to identify unexpected AI behaviors or situations where the AI was unable to provide a solution or is consistently responding with low confidence.</p>

Accelerate your Managed Service Model

Accelerate your Managed Service Model

The [Microsoft Cloud Solution Provider \(CSP\)](#) program enables partners to directly manage their entire Microsoft cloud customer lifecycle.

Partners in the CSP program utilize in-product tools to directly provision, manage, and support their customer subscriptions. Partners can easily package their own tools, products, and services, and combine them into one monthly or annual customer bill.



CSP DIRECT

The CSP Direct model is great for partners who have the infrastructure in place to do it all. If your business meets these requirements and you are ready to go, [enroll today](#).

CSP DIRECT REQUIREMENT CHECKLIST

- Services business model
- Customer support infrastructure
- Customer billing and invoicing capabilities
- Ability to scale

KEY SERVICES FOR THIS OFFERING

- You are the first point of contact for your customers' needs.
- You own and control the billing cycle.
- You sell integrated offers and services — one sales motion to drive services, attach, and upsell.
- You receive in-product tools to directly provision, manage, and support your customers.

CSP INDIRECT

Spend more time with your customers and provide specialized service offerings. If you're not interested in building an infrastructure to provide customer support and billing, [get connected with an Indirect Provider](#).

For additional details, review the [Azure Managed Services Playbook for CSP Partners](#).

Understanding Intellectual Property

The idea of developing “productized” intellectual property (IP) may sound daunting. But many partners find that they already had IP, it just wasn’t packaged that way. If you did something custom that was successful for one client, there may be more customers that would benefit from solving the same problem.

Review your most successful projects to see if there are repeatable elements that you can productize. Repeatable elements can be about your own industry or process best practices, or even focus on common customer pain points. Start small. Your IP can be a simple template or just a few lines of code that automates a function in a way your market typically needs. Productizing IP and creating repeatable processes has been a very successful strategy for many partners.

Some partners are achieving gross margins in excess of 70% by productizing IP and selling it to their customers on a recurring revenue basis.

Productizing IP helps you create stickiness with customers and opens up opportunities to sell your solutions through the partner channel. If you don’t want to create your own IP, you can also look to the partner ecosystem for incremental solutions that can be bundled with Microsoft’s offerings to round out your total solution. There are multiple opportunities for building intellectual property that can be used to expedite engagements, or even as an entire engagement. With the ability to create fully automated solutions, partners can challenge their creative side to offer up solutions that can save their customers money, as well as add a striking differentiator amongst peers.

Making money in the cloud usually partially requires you to retain IP to drive annuity. Annuity is a key strategic component to a cloud practice, and it is not different for practices focused on AI.

CREATING INTELLECTUAL PROPERTY IN YOUR AI PRACTICE

Broadly speaking, there are two forms of IP that partners create that can yield annuities:

- Provide the data and the platform to access, query and interact with the data
- Provide analytics apps and APIs atop a data platform

If, in the course of operating your AI practice, you have collected data assets, consider treating these as an important part of your intellectual property portfolio and think about how enabling controlled third-party access to your data sets might yield new annuities. If you do not own the data, consider the opportunity to build analytic applications and APIs that sit atop the data platform of another- be it the customer’s own data platform or that provided by a third party. This approach can be described as the app-ification of data with AI and the most common approach is to provide SaaS APIs integrated by others in delivering their solution.

CREATE SAAS APIs FOR YOUR GENERALIZED MODELS

For the predictive scenarios for which it applies, consider exposing your predictive services via REST APIs in the SaaS approach. In this approach, customers typically have access to a free tier to experiment with your predictive service, but then have to pay for use as their consumption crosses thresholds you specify. Consider placing your predictive web services built with Azure Machine Learning and hosted in Azure Container Service behind [Azure API Management](#) to monitor and meter third-party access to your intellectual property.

PACKAGE YOUR PROCESS

Another way partners are creating IP in AI practices is by packaging their assessments, documents, and processes into proprietary, reusable components that only they own and can deliver. For example, package a service around delivering envisioning sessions with customers that enable you to quickly get to the best possibilities quickly. Offer this service on an annualized basis as your customer’s and their data evolve, and never leave them without thinking about new opportunities to innovate with your help.

ENGAGE LEGAL COUNSEL

Key to partner success with IP is taking care with licenses, contracts and terms of use. To this end, partners should make sure to protect their IP by involving legal counsel early before any customer uses the new IP.

WHAT INTELLECTUAL PROPERTY SHOULD YOU CONSIDER DEVELOPING IN YOUR AI PRACTICE?

In the Microsoft AI Practice Development Study, 555 partners that identified as having an AI practice were asked which intellectual property offerings they provide within their practice. Observe that the top 5 project services sold were: pre-configured visualizations, dashboards and reports, proprietary algorithms, automated alerting and logging, analytics as a service and automated data migration and integration.



INTELLECTUAL PROPERTY OFFERINGS

Pre-configured Visualizations, Dashboards and Reports	32%
Proprietary Algorithms	29%
Automated Monitoring, Alerting and Logging	27%
Analytics as a Service (e.g., packaged APIs for ML models and agent interactions)	26%
Automated Data Migration & Integration	24%
Domain Specific Functionality	24%
Analytics Platform	24%
Custom Chat Bots	22%
Assessment Tooling	15%
Custom Agents	15%
Online Training and Self-Paced Learning	13%

Source: Microsoft Cloud Practice Development Study, MDC Research, June 2017.



Let's unpack why the top 5 IP offerings make sense and why you should be considering developing and selling these as well.

Pre-configured visualizations, dashboards and reports:

With each project, partners add new visualizations, dashboards and reports to their portfolio. By taking the extra steps of productizing and generalizing what would otherwise be specific to a customer, you are unlocking the ability to capture recurring and repeatable value from what was largely a one-time effort. As the portfolio grows, so does the menu from which customers can choose to purchase and in turn realize more benefit from their relationship with you.

Proprietary algorithms: Developing custom algorithms for custom AI is typically a very time and resource-intensive process. As partners supported in the study, these proprietary algorithms aren't just used once in a project for a single customer, which would leave a significant revenue opportunity on the table. Instead the custom algorithms are packaged up and offered for sale to other customers for use in their scenarios.

Automated monitoring, alerting and logging: according to the partners we interviewed, a comprehensive solution for monitoring the deployed AI solution in production that includes capabilities for monitoring the performance of the AI and underlying models is often something they build as a custom solution. It should come as no surprise, therefore, that some partners have taken the extra step to productive their suite of monitoring tools, processes and know-how into something they can sell, repeatedly.

Analytics as a service: we covered this option in the earlier discussion as a managed service. Once partners have created their custom AI (e.g., they have fine-tuned their models, or perfected their agent interactions), taking the extra step towards enabling third parties to use the AI capabilities these provide as a form of pre-built AI creates another revenue stream and adds value to the partner organization's IP portfolio.

Automated data migration and integration: Since data is such an important pre-requisite to any custom AI effort, for each project partners are likely often re-building similar sets of tools and processes for extracting the data, moving the data to the location where it can be analyzed, understanding the data and preparing the data. As the survey suggests, many partners have recognized the opportunity this creates- package this internal tooling and make it available for 3rd parties to leverage in dealing with their data, and create a new revenue opportunity that also financially support the innovation of the internal tooling.



Protect your AI IP with the Microsoft Azure IP Advantage

Microsoft’s [Azure IP Advantage program](#) represents the industry’s most comprehensive protection against intellectual property (IP) risks, particularly revolving around IP infringement. The Microsoft Azure IP Advantage program includes the following benefits:

- Best-in-industry intellectual property protection with uncapped indemnification coverage will now also cover any open source technology that powers Microsoft Azure services, such as Apache Spark used for machine learning in Azure HD Insight.
- Patent Pick: Makes 10,000 Microsoft patents available to customers that use Azure services for the sole purpose of enabling them to better defend themselves against patent lawsuits against their services that run on top of Azure. These patents are broadly representative of Microsoft’s overall patent portfolio and are the result of years of cutting-edge innovation by our best engineers around the world.
- Springing License: Microsoft is pledging to Azure customers that if Microsoft transfers patents in the future to non-practicing entities, they can never be asserted against them.

With these changes, Microsoft now offers our customers industry-leading protection against intellectual property risk in the cloud.

CONSIDERATIONS FOR YOUR AI IP

Fundamentally, in order to benefit from the Microsoft Azure IP Advantage program, some parts of your solution need to run in Azure, such that you meet minimal spend requirements. For example, if building your AI solution in a hybrid fashion you might perform some data wrangling and model building on-premises, but deploy the trained models to Azure using Azure Machine Learning services. As long as you meet the following requirements, you would be eligible for Microsoft Azure IP:

- For patent pick eligibility: you must (i) have an Azure usage of \$1,000 USD per month over the past three months; (ii) have not filed a patent infringement lawsuit against another Azure customer for their Azure workloads in the last 2 years; and (iii) show evidence of a current patent litigation that occurred after February 8, 2017. Legal transactional costs apply.
- For springing license eligibility: you must have an Azure usage of \$1,000 USD per month over the past three months.

Define Industry Specific Offerings

It is recommended that you specialize by industry, functional process, or technology.

In the earlier section on Industry Opportunities, we provided a range of examples of how partners have delivered successful AI solution by industry (across healthcare, financial services, manufacturing, retail, government, and education). Why do partners choose to verticalize their solution like this?

Irrespective of if you are building AI solutions or line-of-business applications, the reason that you verticalize boils down to maximizing your domain expertise. For non-AI solutions, partners can learn just enough of the domain to incorporate the required knowledge into the solution they deliver, and in reality they may not always need the domain expertise to deliver a viable solution. AI is different. The reason, as we will explore in more detail in the section on Hire and Train, has everything to do with the increased importance of domain expertise in the data science process that is employed by your AI implementation team when creating AI solutions. The following diagram summarizes the situation well:

AI solutions are built by teams having overlapping capabilities in computer science, math, AND domain expertise. Without the domain expertise, the AI solution might accurately predict an outcome, and its implementation could be extremely efficient from a computation standpoint, but the outcome itself could be meaningless or useless from the real-world standpoint of the domain.

As a simple example, consider the situation if you were building a dynamic pricing AI that automatically sets the price of a product in an e-commerce website to the price that encourages the user, whose browsing habit it observes, to purchase. What if the AI decided that price should be \$0? How did you know that would be an unacceptable price to recommend?

Answer: some knowledge about the domain would tell you that a price of \$0 would cause your customer to lose on every "sale" and that losing money is not an objective.

Now consider the importance of domain expertise to non-trivial examples: knowing how diseases are identified (in healthcare), how children learn and stay engaged (in education) or the human psychology behind the stock market transaction (financial service). As your practice delivers AI solutions for your customers, you will continue to enhance your domain expertise and create unique value, because you will have explored that domain from a perspective that is uncommon in the market. It is that value created from your domain expertise that you want to capture, refine and re-sell by specializing.

EXAMPLES OF THESE TYPES OF SPECIALIZATION:

- Industry specialization: manufacturing, banking, retail
- Functional process specialization: accounting, human resources, marketing campaign management
- Technology specialization: systems management, analytics, enterprise resource planning

Think about it this way; if there is lack of differentiation in the market owing to approaches like verticalization, then price becomes the primary differentiator between you and your competition. This can erode your margins and trap you in a business they can't afford to invest in as prices race to the bottom in order to win customers.

Establish your company as an expert in your selected areas. You can also focus on a specific technology or become known as an early adopter and technology leader. But the real value comes from IP or expertise in an industry, vertical, or business process. The combination of adding IP to an industry or business process expertise makes that advantage even more powerful.

Our research with partners suggests mastering one specialization before adding additional ones. We understand that it is easy to be distracted, by saying "yes" to every request, and by diversifying into too many offerings. But in the long run, it is better to say "no" to those projects that are outside of your focus. Partners have shown benefit from having a strict focus on one key solution and growing by expanding one industry at a time.

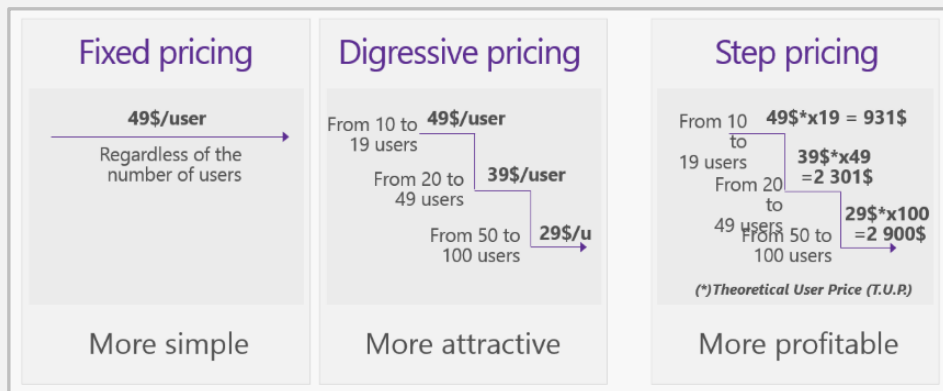
Define Your Pricing Strategy

Pricing your AI offer is no longer determined simply by cost plus margin.

Pricing an AI offering is different in today's marketplace. Increasingly it is about return on value (ROV) — the added benefits (e.g., better per-unit price, improved service characteristics) your customer gets by being a better customer of yours (e.g., buying contracts with longer durations, making upfront payments, etc.). Customers will only pay as much as the value they estimate they will get from the offering.

Yet, price is rarely mentioned on AI partner websites. The sales copy speaks of partnerships and the price is often revealed only after speaking to a sales representative. In an offer-based strategy, this will not work. In the business of cloud, price is always disclosed up front and the same is true if you are building managed services or intellectual property around your AI offers.

Because price is part of your value proposition and solution offer, your pricing is something you can be proud of and share. Remember you are in a race against yourself and the competition. Be proud and show your price early, removing any customer concerns. Pricing is now a result of the product, and it begins with your value proposition. To understand how to price your offers, let's review the pricing strategies to see the benefit they bring to your solution offer.



© 2016 Lemon Operations for Microsoft

STANDARD PRICING

Pricing is the consequence of the product and aligns to the accepted industry/application standard. Think of this as reference pricing; as in customers have seen similar products sold for this amount, so you price your offer so that it is similar. What's the standard price for a mobile phone app? \$0.99. If you charge more, you are breaking from the industry accepted, standard pricing. This is an old way to look at pricing. Buyers today will accept this model, but they do not prefer it and it provides minimal help in getting your offer purchased. Let's look at the other options that you should consider for your AI practice.

VIRTUOUS PRICING

Virtuous pricing is about using the price as a sales weapon. The goal of virtuous pricing is to create a virtuous sales cycle within your customers, where each sale encourages the next sale within the customer organization. It fosters product adoption and proliferation. Let's begin with a counter example of what is not virtuous pricing — a fixed price per user. For your AI practice, this could be pricing per client that consumes your AI web service. Here, you have a simple pricing structure (which is important), but there is nothing to encourage more aggressive purchasing by the customer.

Enter digressive pricing, which drops the per-unit price with the purchase of more units. Your customers get a discount per unit price the more they buy. This can help create a virtuous sales cycle within the customer because now the customer is looking for way to bring their cost per unit (e.g., user, client, etc.) down.

For example, assume one line of business has already purchased 19 users from you at \$49 per user for your AI web service. Now, there are discussions within another line of business within the same customer organization to purchase a similar product from a competitor or to purchase yours. Your existing customer is incentivized to lobby on your behalf because if the other line of business purchases your product, their cost per user will drop to \$39 per user. And the cycle can continue as each new group evaluates your solution offering.

There is a way to adjust digressive pricing slightly to make it significantly more profitable — step pricing. This method sets the price for each step as the top number of users in the range. Building on the example from digressive pricing, let's say that the customer purchased 15 users. They would pay for the equivalent of 19 users since that is the price for this range of units. Why is this more profitable? Because your customer is effectively paying you for the 4 users they are not using (yet) — which goes straight into your profits. What's more is you have amplified the virtuous sales cycle because the customer wants to get as close to the maximum number of users for the step as possible to get the lowest possible cost per unit within the step.



FLAT RATE PRICING

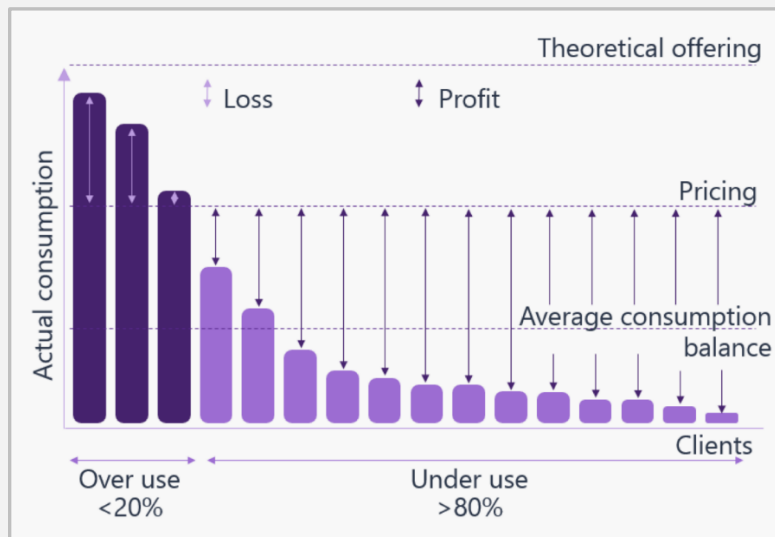
This is one of the most powerful business pricing strategies.

You have probably already experienced it, although you may not have realized. Flat rate pricing is leveraged by banks, insurance companies, etc., every time you pay premiums. While it can have varying levels of sophistication (banks and insurance firms use sophisticated versions of this based on significant work by their actuaries), the model can be described and implemented in a simple fashion.

The basic idea is that you provide a certain quantity of value for a set cost that all customers pay. Let's say your AI practice built a custom AI web service that helps predict a sales outcome. Some customers may come close to (or even exceed) using the full value of what they pay for - for example, they are the big box retailers who depend on your custom AI to customize the consumer's experience. The rest (the smaller chains and boutique online stores) are nowhere close (they have smaller traffic, and as such each individually is making fewer requests against your custom AI web service). A well-crafted model identifies the average consumption across all of your customers, and creates a situation where over 80% of the customers are using less than what they are paying for (and ideally less than the average consumption) and fewer than 20% are using more.

You set your price to be above the average consumption. By doing so, clients in the 80% who use less than what they pay for (the smaller chains and boutique online stores) generate your profit. The further they are below the average consumption, the more profit they generate.

For the 20% who use more than they pay for (the big box stores), you might take a loss on them individually. However, in the aggregate, the long tail represented by the 80% of customers who do not fully use what they pay for more than covers the cost of your heavy consumers, and these heavy consumers are likely to be your biggest champions. So, there are tangential benefits to supporting their cost. Pricing models built around flat rate pricing have shown between 1.5 and 3 times as much profit as traditional models.



© 2016 Lemon Operations for Microsoft

UPFRONT FEES

Another consideration of your pricing strategy is whether to charge your customers an upfront fee.

Reasons for doing so, including providing some working capital to get resources going in the early days of your practice, mitigate the risk that a customer abandons a project without any payment and ensure the customer is as invested in a project as you are. It can also serve to minimize financial impact to your practice when the customer has requested longer payment terms.

PAYMENT TERMS

The final consideration for your pricing strategy is the payment terms. This is defined as the duration of time between when you invoice the customer for services rendered and when you receive payment from the customer.

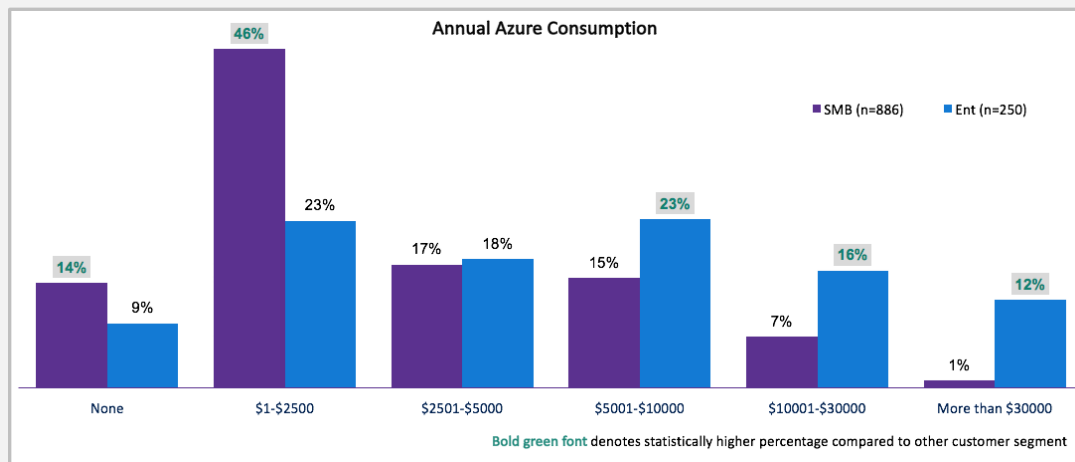
Payment terms are measured in days; for example, 10 days, 15 days, 30 days, or 90 days. These are usually expressed as NET 10, NET 15, NET 30, or NET 90 payment terms. In addition, you might consider offering the customer a discount for prompt payment on your shortest payment. For example, NET 2/10/30 is used to describe terms where a 2% discount is provided for payment received within 10 days of invoicing, otherwise the full invoice amount is due in 30 days. In interviews with partners we found that the most common payment terms used was NET 30, and that for SMB customers shorter payment terms were preferred.

Calculate Your Azure Practice Costs

Azure Practice Costs

Your practice relies on Azure services to deliver customer success, so understanding the Azure related expenses incurred in delivering a customer solution is critical. How do you calculate these Azure costs?

In our research of 1,136 Azure partners, we found that enterprise-focused partners spent a median of \$8,107 in Azure for research and development annually; whereas SMB-focused partners spent a median of \$1,933 annually.



Source: Microsoft Cloud Practice Development Study, MDC Research, November 2016.

Use the [Azure Pricing Calculator](#) to estimate Azure costs and be sure to check the resources list for pricing on the various services. You can build an estimate online and export it to Excel for further refinement and analysis. This tool will give you the retail rates (also known as the Pay-As-You-Go option) for Azure services, so treat it like the “high end” of your consumption estimate.

Become familiar with the discounted pricing and Azure credits:

- **Graduated Pricing:** Services like Azure Storage have tiered pricing based upon the volume used.
- **Enterprise Agreement:** By making a three-year monetary commitment, Azure services are available at a discount off retail rates. To learn more, see [Enterprise Agreements](#).
- **Azure Credits:** Microsoft Partners can receive Azure credits as a part of their benefit. For example, partners with the Silver Cloud Platform Competency receive \$350 USD per month in Azure credits; those with Gold Cloud Platform Competency receive \$600 USD per month in Azure credits.



It can be helpful to identify items which are used elastically versus items that have a fixed monthly cost. Significant savings can be achieved via elastic use of resources because you can turn them off (or pause them) when they are not in use.



FOR EXAMPLE:

- **Elastic:** Your HDInsight cluster is only used by data scientists during business hours for exploring big data sets stored in Azure Data Lake, preparing data sets for machine learning and training models against big data sets in a distributed fashion. So long as state such as Jupyter Notebooks are saved to the Data Lake prior to shutdown and Hive table metadata is saved to a SQL Database instance, the cluster can be deleted nightly and scheduled to re-deploy every morning. This can create significant savings when the cluster is not in use.
- **Fixed:** Your deployed custom AI webservice deployed using Azure Machine Learning needs to be always available for applications to benefit from its capabilities.

Finally, if you don't understand how much of a given resource you will use, consider building a scaled down proof-of-concept to get a first estimate.

Identify Partnership Opportunities

Partner to Partner

With an AI practice, not all partners will or even should do everything themselves. It is critical that partners learn to seek out other AI partners to fill gaps in their practice – from missing talent to expertise in building custom models.

It is tempting to want to do everything related to your business, but the fact remains you will never have enough time. This is a primary reason to seek out compatible partners that can help you to:

- Complete your solution
- Build credibility
- Partners for infrastructure
- Leverage joint marketing
- Add-on to sales
- Broaden your customer base


Finding the right type of partnerships is key to finding success in today's competitive market. Some partner combinations meld together well to create success. To help you find the right mix, here's an example of a partner recipe we know works:

ISV + channel-based MSP partners = international success

[Find out what's in it for the Partners and customer and the secret of success.](#)

Partner Across Domains

With an AI practice in particular, selling your AI solution to partners with expertise in other domains can be lucrative. A compelling example of this is in security. The security partner is expert in the domain of all things security and they likely have the historical data to back their experience. If you partnered with the security partner, you would gain domain expertise and data from a very complex and rapidly evolving domain. The security partner would gain the capability to leverage your AI know-how and capabilities to build intelligent systems with AI that could go beyond mainstream approaches.



Make connections with Dynasource.com

It's hard to beat the value of partner-to-partner conversations and networking. Luckily, there are communities to help you expand your network and make an even bigger impact on your business.

[Dynasource](#) is a global, two-sided marketplace that allows partners to connect with other firms that have complementary expertise and capabilities. If you have excess capacity, increase your billability and profitability by finding partners that can drive utilization of your staff. If you lack the capacity, Dynasource can help you connect with partners that have the expertise to enable and expand your solutions. If your customers are looking for a solution that is outside your particular expertise, you can use Dynasource to find an expert that can provide that solution. By connecting through Dynasource and working together, you can meet customer demand for needs across the Microsoft portfolio and grow your business.

TRANSFORM THROUGH COLLABORATION

After you've created a profile on Dynasource, you can search the Dynasource Microsoft Partner Community for qualified resources that can collaborate with your team on an opportunity. This allows you to transform your cloud business and expand your offerings at a pace that works for you.

INCREASE DEMAND FOR YOUR RESOURCES

Business is not always predictable but retaining quality staff is essential to your success. With Dynasource, you can make your resources and capabilities available to partners who can utilize them today, making your workforce agile and billable.



GET STARTED NOW

In order to join Dynasource, all you'll need is general information about your business. Once you create a profile, you can search for other resources and jobs, as well as create and post your own. You can control the availability of your resources and what level of information you would like to share about your capabilities. With the Dynasource premium membership, you can request connections with other members and resources.

Resources

- [Smart Partner-to-Partner Relationships](#)
- [A Recipe for Global Success](#)
- [Selling thru Partners](#)
- [The right ingredients for partner-to-partner success](#)

Define Engagement Process



Pre-Sales, Post-Sales, and Support

For your practice, you should define the technical effort required before the sale (pre-sales), after the sale (post-sales), and in support of the sale. You will need to decide on the technical pre-sales and post-sales requirements for your solution offer.

PRE-SALES

The technical effort required to make the sale involves:

- Discuss the customer requirements and address their objections.
- Develop technical pitch decks. Leverage the [Cloud Adoption Framework](#).
- Technical demo: This demo may be generic or may need customization to better meet the requirements of the customer. The goal of the technical demo is to inspire confidence in your ability to deliver the desired solution by demonstrating you have “already done something like it before.”

POST SALES

The technical effort required after the sale includes:

- Addressing follow-on customer concerns about the technology or implementation.
- Providing training to increase awareness of the solution that will be implemented.
- Providing a technical demo more customized for the customer to better understand their needs before moving on to the next phase of the project.
- Following up with the customer to ensure implementation is on track and meeting expectations.

For guidance with sales efforts, consider the learning paths available in the Microsoft Partner Network [Learning Portal](#).

SUPPORT

Define your customer support program and processes. This includes:

- Defining your support model
- Provisioning your support infrastructure
- Defining and implementing your escalation process
- Selecting and enabling your support options for Azure

Microsoft also provides support for technical presales and deployment services. See the section [Supporting your Customers](#) for more information on available resources and using Partner Advisory Hours.

Identify Potential Customers

Build your prospect hit list.

Your goal is to build the list of prospects that could potentially turn into customers. To accomplish this, create an awareness campaign to draw attention to your practice, highlight your service offerings, and use your success to earn additional business with your customers and the industry at large.



Use these awareness activities to help generate new customers:

WEBINARS AND PODCASTS

A great way to transfer knowledge, establish yourself as an expert, and pique the interest of potential customers.

REFERRALS

Ask for referrals in email and phone calls when talking with existing customers, partners, and vendors who might know someone who is ready for your services.

WHITE PAPERS

These are a great way to build credibility with decision makers. Technical staff often expect a white paper to help them understand underlying architecture and technology of your solutions.

NEWS ARTICLES

Leverage public relation efforts to drive publicity around your technology, things your company is doing in the market, and other topics of current interest.

SOCIAL MEDIA

Social media such as Twitter, LinkedIn, etc., is a place to build awareness, reputation, and customer satisfaction — and gain new customers.

REVISIT EXISTING CUSTOMERS

If AI represents a new practice within a going business concern, the easiest way to acquire new customers for your AI practice is to introduce the AI practice to your existing customers.

Join the Microsoft Partner Network

Partnering with Microsoft

The Microsoft Partner Network is the start of your journey. One of the first steps to partnering with Microsoft for your Azure practice is to join the Microsoft Partner Network if you are not already a member. As a partner, you will gain access to resources like training, whitepapers, and marketing material described in this playbook. It's also where you will set up your users to gain Microsoft Partner competencies and access to your partner benefits.

TO BECOME A MICROSOFT PARTNER

The Microsoft Partner Network provides [three types of memberships](#). Each type provides a set of benefits to help you grow your business. As you achieve your goals, participate in the program at the level that suits your unique needs, so you can access more benefits and develop your relationship with Microsoft and other Microsoft Partners.

- **Network Member:** Receive a set of no-cost introductory benefits to help you save time and money. Use our resources to help build your business as a new partner and discover your next step.
- **Microsoft Action Pack (MAP):** This affordable yearly subscription is for businesses looking to begin, build, and grow their Microsoft practice in the cloud-first, mobile-first world through a wide range of software and benefits.
- **Competency:** Get rewarded for your success with increased support, software, and training.

TAKE THE NEXT STEP WITH A COMPETENCY

As a competency partner, you can earn both gold and silver competencies in one or more areas. Earn a silver competency to help your business demonstrate its expertise or a gold competency to showcase your best-in-class capabilities within a Microsoft solution area. Later in this playbook we'll review the competencies relevant for launching a successful Microsoft Azure practice.

CLOUD ENABLEMENT DESK

The goal of the [Cloud Enablement Desk](#) is to assist partners in obtaining their first Silver Cloud Competency. Partners will be assigned a Cloud Program Specialist (CPS) for up to six months on their way to obtaining their first Silver Cloud Competency.

The Cloud Enablement Desk program requirements include:

- Partner must have a MPN ID.
- Partner must agree to and sign Conditions of Satisfaction that state partner is actively trying to achieve Silver Cloud Competency status and include the name of the primary contact person.
- Partner cannot have an existing Microsoft account management relationship.

Partner Programs for AI Partners

The following programs all build upon your membership in the Microsoft Partner Network and provide you with additional benefits and incentives as you prove your practice by earning competencies, passing assessments and winning customers in your practice focus area.



DATA + AI TRAINING PARTNERS

As a Data + AI Training Partner, you'll be showcased as a subject-matter expert to prospective learners around the globe, including on the [LearnAnalytics](#) portal. Benefits of membership include support with training curricula, content and best practices, access to the analytics portal and referrals from Microsoft. Visit the [Data + AI Training site to sign up](#).

BUSINESS INTELLIGENCE PARTNER PROGRAM

The Business Intelligence Partner program enables you to build your expertise and showcase your solutions to customers. It enables you to get listed on PowerBI.com and supports you in becoming a Solution Partner where you can be showcased on PowerBI.com, AppSource.com or get listed as a Solution Template partner on PowerBI.com. This program also provides funding programs you can tap into for delivering workshops or proof of concepts to your customers.

ADVANCED ANALYTICS PARTNER PROGRAM

The Advanced Analytics Partner program enables you to get listed as a partner who can deliver Advanced Analytic services as well as to get your solution showcased on the Cortana Intelligence Suite Solution Showcase.

CLOUD SOLUTION PROVIDER (CSP)

The [Cloud Solution Provider \(CSP\) program](#) rewards partners for driving cloud revenue through the CSP model, which enables you to drive deeper customer engagement with greater profitability as you get to combine high-margin service offerings with Microsoft cloud products and retain the ability to provide customer support and manage billing for your customer. You get to become your customers' single solution provider and trusted advisor by servicing all their cloud service needs.

P-SELLER PROGRAM

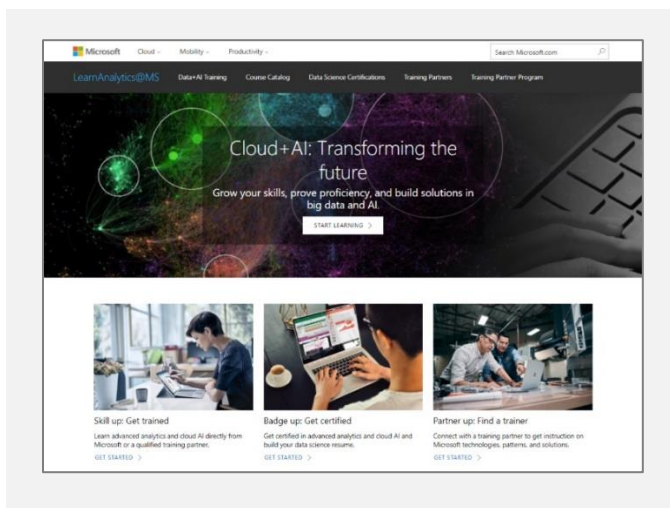
P-Sellers are Microsoft's "go to" partner resources across the customer lifecycle, and act as an extension of Microsoft in working with customers. As a Business Intelligence Solution Partner, you will receive sponsorship for your admission to the P-Seller program.

Stay Informed on AI Matters

There are three resources you should be very familiar with when defining your AI practice strategy. These are Learn Analytics, Azure AI Platform, and Microsoft AI portals

Learn Analytics

The [Learn Analytics](#) portal provides resources for you to use for training and certification of your AI practice team.

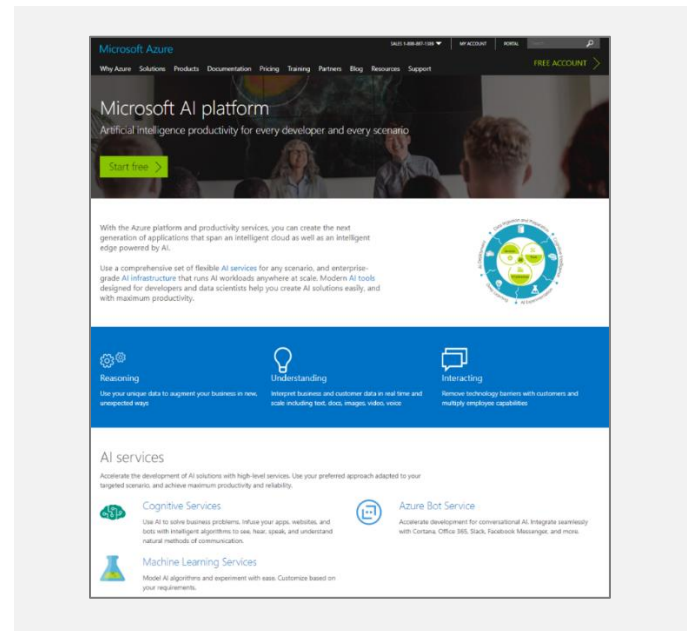


THE LEARN ANALYTICS SITE PROVIDES

- Training resources, including self-paced and upcoming live and in-person events.
- Curated certifications appropriate to the AI partner.
- References to qualified cloud and AI training from partners. Preferred training partners can train your teams and your customers' teams on Cloud AI technologies, accelerating the knowledge needed to develop AI applications.

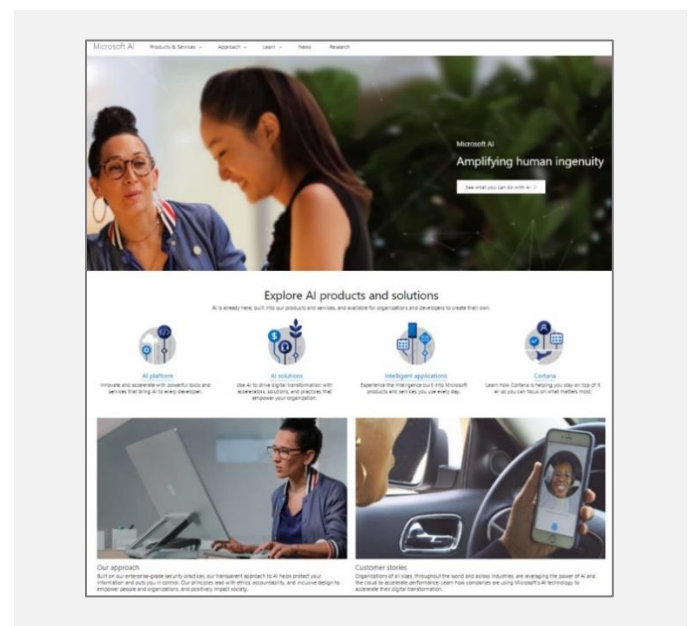
Microsoft AI platform

The [Microsoft AI platform](#) portal provides the guided perspective on utilizing Azure for delivering AI solutions. It provides the curated list of Azure services and Microsoft products most relevant to the practice.



Microsoft AI

The [Microsoft AI](#) portal provides the broadest lens on all the areas in which Microsoft is participating in AI, from the AI platform, to AI solutions and intelligent applications. It provides many resources such as customer stories, Microsoft AI products and services, Microsoft's philosophy on AI, the latest news and current research.



Identify Solution Marketplaces

Azure AI Gallery

The [Azure AI Gallery](#) enables developers and data scientists to share their analytics solutions and build industry reputation.

The gallery contains a variety of resources including:

MODELS

The gallery contains a wide range of models machine learning models that can be utilized in building projects and solutions. Many of these models are provided in the universal format for deep learning models called ONNX, making them immediately useable by solutions built using either Python or C# with WinML.

INDUSTRIES

Explore the Microsoft and community provided experiments and solutions by the industries including retail, manufacturing, banking and healthcare.

JUPYTER NOTEBOOKS

Jupyter Notebooks include code, data visualizations, and documentation in a single, interactive canvas. Notebooks in the gallery provide tutorials and detailed explanations of advanced machine learning techniques and solutions.

SOLUTIONS

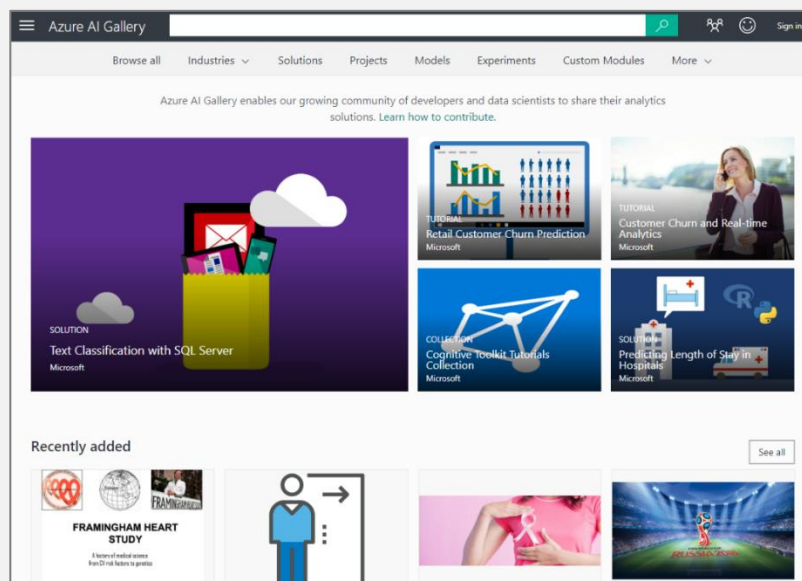
Quickly build Azur AI Solutions from preconfigured solutions, reference architectures, and design patterns. Make them your own with the included instructions or with a featured partner.

TUTORIALS

A number of tutorials are available to walk you through machine learning technologies and concepts, or to describe advanced methods for solving various machine learning problems.

COLLECTIONS

A collection allows you to group together experiments, APIs, and other gallery items that address a specific solution or concept.





Azure Marketplace

Azure Marketplace is an online store that enables you to offer your solutions to enterprises and Azure customers around the world. Within a single, unified platform, customers can easily search, purchase, and deploy your solutions on Azure with just a few clicks.

[Azure Marketplace](#) is the source for thousands of software applications and services certified by Microsoft to run on Azure. Azure Marketplace supports offers that include virtual machines, developer services, and solution templates.

Azure Marketplace gives your solutions exposure through the marketplace page and the listings integrated with the Azure Portal. For example, HDInsight Applications are integrated into the steps users take to deploy an HDInsight cluster (so users could layer your application atop their HDInsight cluster), but are also available via the Marketplace blade of the Azure Portal.

The process of getting your solution listed in the Azure Marketplace is referred to as getting Microsoft Azure Certified. This comes with benefits, many of which include select benefits from the MPN Silver Cloud Competency.

CERTIFY APPLICATIONS AND SERVICES

Solutions sold in the Azure Marketplace must be [Microsoft Azure Certified](#). This provides assurance to your customers that your offers have been tested for usability, readiness, and compatibility with Azure.

- Access broad-reaching Microsoft marketing channels, and receive co-marketing assistance and promotion outside of your Azure Marketplace listing.
- Leverage technical and business planning support to help you maximize your ROI.
- Utilize a self-service portal that contains ready-to-use marketing resources to enhance your communications and messaging.
- Take advantage of resources, such as the Sales Accelerator Toolkit and credits for display, and search advertising to help drive customer adoption.

Once you identify and select the marketplace to target for promoting or selling your services, and increase your visibility. View the [Integrate into a Marketplace guide](#) for details.

Microsoft
Partner
Network

Hire & Train

Artificial Intelligence

aka.ms/practiceplaybooks



Executive Summary



In the previous section, you evaluated the various services your business can pursue as you set up or build your AI practice. Now that you've identified some avenues of success, you may be wondering how to build and train your team.

First, we'll help you define the members of your team and the skills they should bring to the table. If you need to hire to fill gaps, we provide you with detailed job descriptions you can use, as well as ideas on where to look for resources, the factors you should look for in a candidate's skillset, and what you should expect to pay by role and region.

A big focus of this section is the critical piece of ensuring all of your practice resources are trained and continue to receive ongoing training. We cover not just the technical training, but also sales and marketing training.

Additionally, we'll give you details on the specific Microsoft certifications your technical resources should be working towards, both for their own professional development and to earn your organization Microsoft Partner Network competencies.

Top 5 things to do

You're crafting your gameplan to build your team, make sure you nail down these 5 tasks before you move to the next section.

- ✓ Define the members and roles required
- ✓ Identify capability and skills gaps
- ✓ Decide which skills to hire and train
- ✓ Hire to fill gaps in your team
- ✓ Train and certify your team



Hire, build, and train your team

All practices need a team to support the technical and business functions. Given the interdisciplinary nature of AI, practices delivering AI solutions need to be carefully structured to balance breadth of disciplines with depth of resources of a particular discipline.

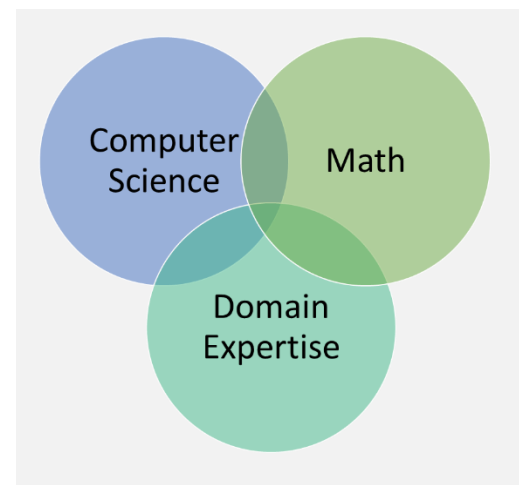
Domain Expertise

A team with data science capabilities contains an overlapping mixture of capabilities that includes computer science, mathematics and domain expertise.

In order to successfully deliver AI and ML based solutions, having strong expertise in the domain of the problem being addressed is absolutely critical to producing AI solutions that are provably correct, performant and well optimized. The domain experts may be permanent team members or customer stakeholders immersed with the rest of your solution team.

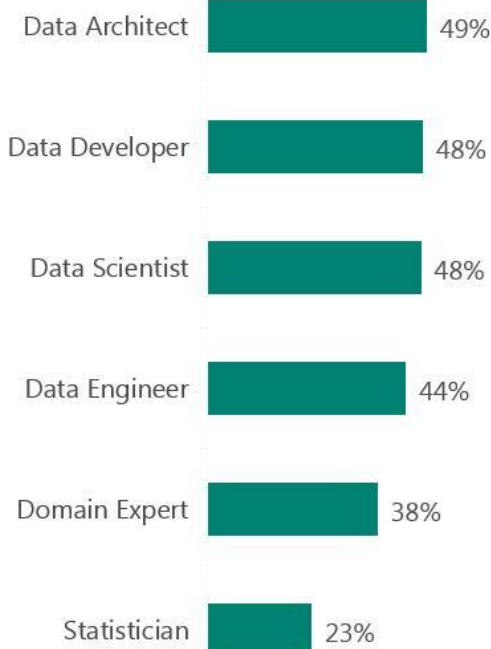
KEY ROLES IN THE AI PRACTICE

The following table identifies the key roles and summarizes the interactions between roles. In the sections that follow, these roles are described in greater depth.



According to our recent survey of AI partners, data architects, data developers, and data scientists are the top desired professionals in AI and ML services.

TOP AI OR ML PROFESSIONAL EMPLOYED (N=236)



Source: Microsoft AI Practice Development Study, MDC Research, December 2017


ROLE	DESCRIPTION
Chief Data Officer	A C-Level role, the CDO is a new and emerging role responsible for the enterprise wide data and information assets and has authority over the data scientists, architects, engineers and developers.
Data/AI Scientist	Models complex business problems, discovering new business insights and opportunities to apply AI, machine learning and deep learning.
Data Architect	Designs the AI solution, including how data is ingested, stored and processed, and works with the data scientist to identify where model training occurs, and how trained models are deployed for production use. Provides the roadmap that is implemented by the data engineers and data developers.
Data Engineers	Works with the data architect to provides the infrastructure to make data accessible to the data scientists.
Data Developers	Provides the code and automation to make data accessible to data scientists, assists in data collection, preparation and correlation and in enabling API access to deployed models.

Technical Roles (Architecture, Infrastructure, Development & Data Science)

These roles form the heart of your AI solution. Hiring the right people can turn your vision into reality.



The **Data Architect** drives customer initiatives leveraging data and analytics services to solve the biggest and most complex data challenges faced by enterprise customers. The DA is a technical, customer facing role, accountable for the end-to-end customer deployment and usage experience for Azure data services. DAs own the Azure technical customer engagement including: architectural design sessions, implementation projects and/or proofs of concept and pilots. The Data Architect is proficient in distributed computing principles and familiar with key architectures including Lambda and Kappa architectures, and has a broad experience designing solutions using a broad set of data stores (e.g., HDFS, Azure Data Lake Store, Azure Blob Storage, Azure SQL Data Warehouse, Apache HBase, Azure DocumentDB), messaging systems (e.g., Apache Kafka, Azure Event Hubs, Azure IoT Hub) and data processing engines (e.g., Apache Hadoop, Apache Spark, Azure Data Lake Analytics, Apache Storm, Azure HDInsight). The ideal candidate will have experience in customer facing roles and success leading deep technical architecture and design discussions with senior executives. Five plus years of experience with deep understanding of databases and analytics, including relational databases, data warehousing, big data, business intelligence and analytics. Five plus years of success in consultative/complex technical sales and deployment projects. Technical BS degree in Computer Science or Math background desirable.



The **Data/AI Scientist** is responsible for identifying the AI opportunities present in the customer's requirements, their data, and helping shape both the vision and the solution that deliver the desired insights, enable reasoning atop or understanding of the data by applying AI and machine learning or deep learning, or by creating natural interactions that enable the software solution to interact naturally with the human users, organically amplifying their capabilities. The Data Scientist is a technical, customer facing role, who along with the Data Engineer is accountable for the end-to-end data pipeline envisioning and development that starts with addressing issues of data acquisition and data sampling, data exploration and data quality assessment, data wrangling to massage the data so it is better suited to applying machine learning and deep learning, and visualizing or reporting on such data to make the insights available to the customer's business. The ideal candidate will have experience in customer facing roles and has a cross-disciplinary background consisting of statistics, probability, software development and ideally expertise in the problem domain. A technical BS degree in Computer Science or Math background is highly desirable. Three or more years customer facing experience desired.

The **Data Engineer** is responsible for helping to select, deploy and manage the systems and infrastructure required of a data processing pipeline in support of the customer requirements. Primary responsibilities revolve around DevOps and include implementing ETL (extract, transform and load) pipelines, monitoring/maintaining data pipeline performance. The Data Engineer is proficient in distributed computing principles and familiar with key architectures including Lambda and Kappa architectures, and has a broad experience across a set of data stores (e.g., HDFS, Azure Data Lake Store, Azure Blob Storage, Azure SQL Data Warehouse, Apache HBase, Azure DocumentDB), messaging systems (e.g., Apache Kafka, Azure Event Hubs, Azure IoT Hub) and data processing engines (e.g., Apache Hadoop, Apache Spark, Azure Data Lake Analytics, Apache Storm, Azure HDInsight). The ideal candidate has three or more years of experience working on solutions that collect, process, store and analyze huge volume of data, fast moving data or data that has significant schema variability.

The **Data Developer** can be seen as the development peer to the Data Engineer. Whereas the Data Engineer is focused on infrastructure and DevOps, the Data Developer is more focused on supporting the data collection, integration, cleansing and implementation of the data pipeline supporting the AI solution. The Data Developer may be a customer facing role, but the primary responsibilities include implementing ETL (extract, transform and load) pipelines, monitoring/maintaining data pipeline performance and implementing big data or advanced analytics solutions. The Data Developer is proficient in distributed computing principles and familiar with key architectures including Lambda and Kappa architectures, and has a broad experience across a set of data stores (e.g., HDFS, Azure Data Lake Store, Azure Blob Storage, Azure SQL Data Warehouse, Apache HBase, Azure DocumentDB), messaging systems (e.g., Apache Kafka, Azure Event Hubs, Azure IoT Hub) and data processing engines (e.g., Apache Hadoop, Apache Spark, Azure Data Lake Analytics, Apache Storm, Azure HDInsight). The ideal candidate has three or more years' experience working on solutions that collect, process, store and analyze huge volume of data, fast moving data or data that has significant schema variability.

The **Cloud Architect (CA)** drives customer initiatives in collaboration with customers. The CA is a technical, customer-facing role that is accountable for the end-to-end customer cloud deployment experience. CAs own the technical customer engagement, including architectural design sessions, specific implementation projects and/or proofs of concepts. The ideal candidate will have experience in customer-facing roles and success leading deep technical architecture and application design discussions with senior customer executives to drive cloud deployment. A computer science or related engineering degree is required.

The **Senior Developer** has a history of designing, owning and shipping software, as well as excellent communication and collaboration skills. With a focus on cloud-based application development, the candidate must have demonstrable experience architecting and deploying applications to cloud platforms, the ability to effectively integrate disparate services as needed, and decide when to implement IaaS, SaaS, and PaaS components. As a mentor to junior developers, the senior software developer should have a solid understanding of the software development cycle, from architecture to testing. They should have a passion for quality and be a creative thinker. A senior developer will write secure, reliable, scalable, and maintainable code, and then effectively debug it, test it and support it live. This person should also be comfortable owning a feature and making decisions independently. Another aspect of a senior software developer, is that they can effectively gather customer requirements, and ask clarifying questions when needed. This person must be able to translate these requirements to actionable tasks they will perform, or delegate to members of the team.

The **Developer** enjoys the challenge of building applications that solve today's business needs. This person must be willing to keep up to date with the fast-moving cloud services landscape to remain an effective member of the development team. A software developer should work equally well on a team or independently, given a set of project requirements or tasks. This requires the developer to possess excellent communication and collaboration skills. The developer should understand the aspects of the software development cycle, from architecture to testing. This person will design, build, and maintain efficient, reusable, and reliable code.

Given the natural evolution of skills as developers gradually take on more AI development responsibilities, their AI experience lies on spectrum. Microsoft has identified the following three developer personas, from least experienced with AI to most experienced with AI: **Pre-emerging AI Developer**, **Emerging AI Developer**, and **Professional AI Developer**.

Leadership Roles

Consider the following management positions if your development effort will involve eight or more technical staff. In smaller teams, senior-level employees (e.g., lead data scientist, lead developer) sometimes take on management duties along with their other responsibilities, removing the need for dedicated managers.

The **Chief Data Officer** (CDO) drives the definition of balancing data governance, protection with data discovery and analytics. This role establishes the organization's data analytics platform strategy, selection of appropriate technologies and focuses on strategic and timely talent acquisition. An important responsibility of the CDO is in creating a learning culture within the organization by providing for and fostering an environment for learning.

The **Data Protection Officer** assesses and advises across the company group for data protection and privacy matters related to security. This role is a subject matter expert in the handling of personal data, and ensures there are policy and compliance processes to comply with local data protection legislation. Expert knowledge of global and national data protection law and practices, as well as the General Data Protection Regulation (GDPR) is a requirement, as well as the ability to fulfil the tasks referred to in Article 39 of the GPDR. Experience in conducting data privacy compliance, reviews, and audits is beneficial.

The **Product Manager** (or Product Management team) establishes and sustains the business case for the project and plays a key role in identifying and setting priorities across the target audience. This includes ensuring that business expectations are clearly articulated and understood by the project team, and that the functional specifications respond to business priorities.

Product Management owns the vision statement for the project. The vision statement is an informal document that communicates the expectations and assumptions on which the project is based.

Product Management is also responsible for high-level project communications such as business projections, project costing, and contract negotiation. Product Management communicates the high-level milestones to the target audience and other team members.

The **Program Manager** or Program Management team "owns" the specification for an application's features and functionality and coordinates the day-to-day communication required to develop and deliver the application effectively and consistently within organizational standards.

Program Management has a key communication and coordination role. With input from other team leads, Program Management assists Product Management in articulating the vision for the project. Using this vision, Program Management drafts the initial version of the functional specification and is considered the keeper of the functional specification. Program Management is responsible for all activities associated with analysis, specification, and architecture. Program Management is also responsible for defining how the project will interoperate with external standards, maintaining external technical coordination and communication, and managing the master schedule.

Support Resources

A lot of effort goes on behind the scenes, or in positions that involve post-sales customer engagement. To ensure long-term success of your projects, consider hiring some of these support roles.

In an AI practice the support roles are generally similar to those for any other software solution, with one exception: monitoring model performance. In this case, your data scientist team will need to be involved to assist in crafting the dashboards that your customer support team monitors, and may need to be available to assist in escalations when a non-performing model has been detected.

The **Customer Success Manager** is passionate about engaging your customers and helping them expand their use cases. They have excellent relational skills and can create win/win environments for all parties they work with. In their day-to-day responsibilities, they own the overall relationship with assigned clients by increasing adoption and ensuring retention and satisfaction. They make a large impact on your enterprise security business by establishing a trusted and strategic advisor relationship with each assigned client, driving continued value of your products and services. The Customer Success Manager will help drive sales by working to identify or develop upsell opportunities. Additionally, they will advocate customer needs and issues cross-departmentally and program manage account escalations. Qualifications include prior experience in customer success or equivalent history of increasing customer satisfaction, adoption, and retention.

The **Quality Assurance (QA) / Test Technician** is thorough and detail-oriented and should work well with established processes. The primary goal of this role is to help avoid defects in your final product or solution. This person will be involved throughout the development process and use their intuition to problem solve and identify technical, procedural, and usability concerns. They must take meticulous notes, be organized about recording process steps, and work well with others since they will be coordinating with technical and management teams to ensure that the correct measures are put into place to align the final product with the initial goal.

The **User Support Specialist** assists customers who are having technical issues with your product, or who need help realizing the full benefit of your solution to help them deliver their cloud-based workloads. They will likely be in a position to help customers navigate the operational challenges of cloud computing, so thoroughly training them on both your product — and the infrastructure on which it is built — is paramount to their success, and ultimately, your customers' satisfaction. Qualifications include technical support experience and great communication and interpersonal skills (soft skills). Experience with cloud technologies is a major plus.



Job Descriptions for your Technical Team

The following tables provides detailed job descriptions you can utilize to hire the key technical resources. All technical skills, non-technical skills, certifications, and technologies listed are potential items a candidate should have, but no candidate will have all the items listed.

Data Scientist

A Data Scientist is responsible for identifying the insight opportunities present in the customer's data and helping shape the data pipeline that deliver the insights by applying advanced analytics (e.g., machine learning) in collaboration with the customer. The Data Scientist is a technical, customer facing role, who along with the Big Data Engineer is accountable for the end-to-end data pipeline envisioning and development that starts with addressing issues of data acquisition and data sampling, data exploration and data quality assessment, data wrangling to massage the data so it is better suited to applying advanced analytics, and visualizing or reporting on such data to make the insights available to the customer's business. The ideal candidate will have experience in customer facing roles and has a cross-disciplinary background consisting of statistics and software development. A technical BS degree in Computer Science or Math background is highly desirable. Three or more years of customer facing experience desired.

Technical Skills

Advanced analytics, analysis services (tabular, multi-dimensional), backpropagation, bagging, boosting, Bayes, big data, business intelligence, classification, clustering, cloud data analytics, data architecting, data cleansing, data migration (cross platform / upgrade), data modeling (physical and logical), data movement, data potency, data transformation, data warehouse design, database architecture, database design, decision trees, descriptive analytics, forests, genetic programming, image processing, inverse deduction, machine learning, neural networks, predicative analytics, prescriptive analytics, recommendation, regression, rules, support vector machines, statistics, text mining.

Non-Technical Skills

Consultative requirements gathering, collaboration, stakeholder management, relationship management, technical oversight, technical recommendations, problem solving, risk management, architecture design session, program management, proof of concept design, technical demonstration.

Technologies

AWS Machine Learning, AWS Lex, AWS QuickSight, AWS RDS, AWS Rekognition, AWS Redshift, AWS S3, Azure Data Catalog, Azure Data Factory, Azure Data Lake, Azure Storage, Azure DocumentDB, Azure SQL Data Warehouse, Azure Import/Export, Azure HDInsight, Azure Machine Learning, Azure SQL Database, Azure Search, Azure Cognitive Services, Cassandra, Cloudera, Cognos, Cortana Intelligence, Data warehouse, Database, DB2, Excel, Flume, IBM Bluemix, Hadoop, HBase, Hive, Hortonworks, Machine Learning, MapR, Microsoft R, ML, MongoDB, MariaDB, MySQL, NoSQL, Oracle, Oracle Exadata, Oracle SOA, Pig, Power BI, PostgreSQL, Python, Qlik Tech, Revolution R, SAP HANA, SAS, Spark, SQL Server Analysis Services, SQL Server, SQL Server IaaS, SQL Server Integration Services, Sybase, Tableau, TSQL.

Programming/Scripting Languages: R, Scala, Python, DMX, DAX, MDX, SQL, T-SQL, Java
Platforms: Linux (Red Hat, Ubuntu, Debian, etc.), Windows.

Certifications

MCSA in Machine Learning, having passed the following exams: [70-773: Analyzing Big Data with Microsoft R](#), [70-774: Perform Cloud Data Science with Azure Machine Learning](#)

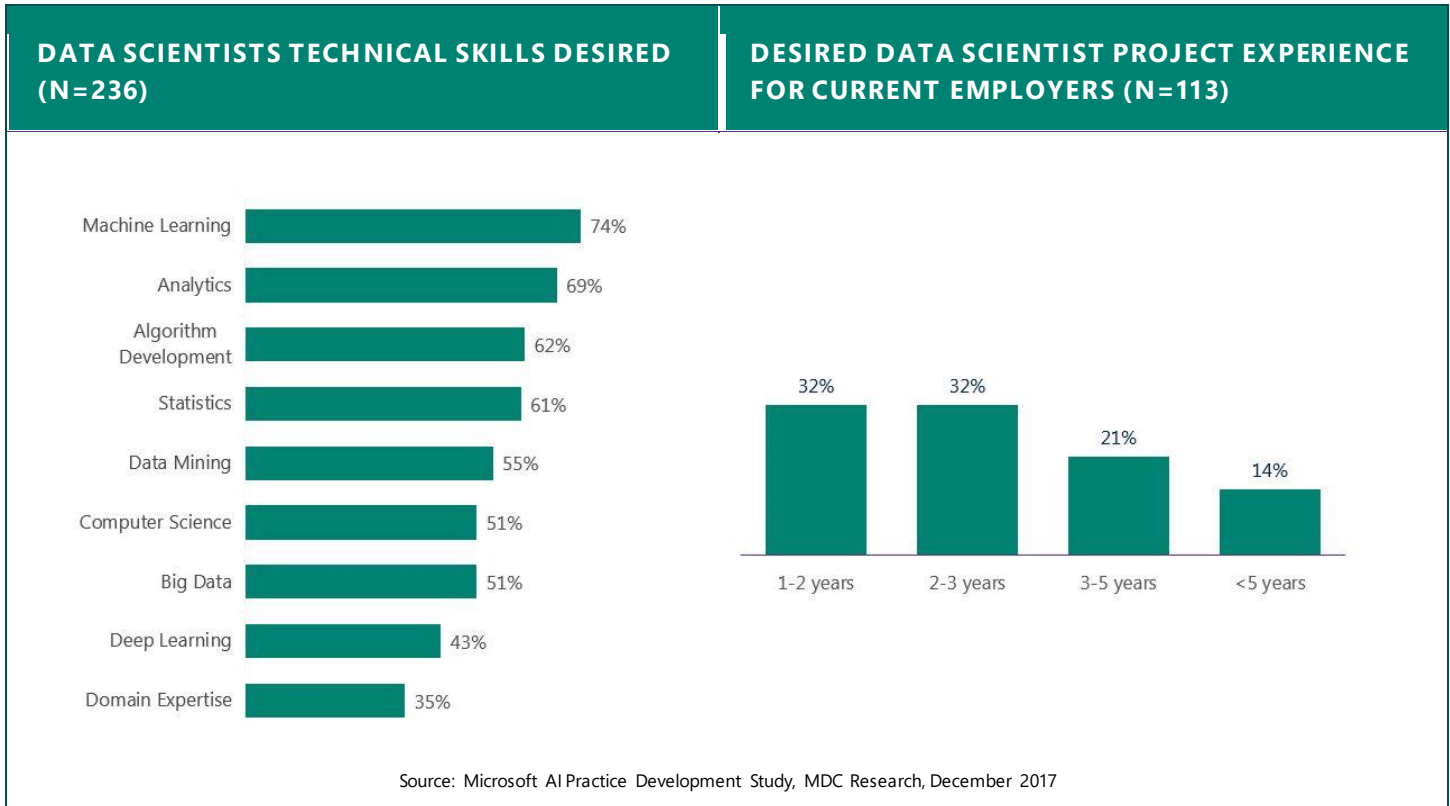
Other certifications include: Master or PhD in Data Science, Statistics or Probability from accredited universities, Certified Analytics Professional (CAP), Certification of Professional Achievement in Data Sciences, Cloudera Certified Professional: Data Scientist (CCP:DS), edX Verified Certificate in Data Science Curriculum, EMC Data Science Associate, MCSE Business Intelligence, MCSE Data Management and Analytics, Revolution R Enterprise Professional, SAS Certified Data Scientist.

Project Experience Types/Qualities

Advanced analytics (including machine learning), automating data munging, building visualizations, machine learning modeling, distributed training, training on large data sets, operationalizing models, significant challenges.



In the Microsoft AI Practice Development Study, AI partners indicated that machine learning, analytics, and algorithm development are the top skills they desire to see in a data scientist. However, their expectations for the years of experience a data scientist should have were relatively short, with the over half looking for less than 3 years.



Data Architect

A Data Architect drives customer initiatives leveraging Azure data and analytics services (e.g., ranging from SQL Server to SQL Database and SQL Data Warehouse to Cortana Intelligence Suite) to solve the biggest and most complex data challenges faced by enterprise customers. The DA is a technical, customer facing role, accountable for the end-to-end customer deployment and usage experience for Azure data services. DAs own the Azure technical customer engagement including: architectural design sessions, implementation projects and/or proofs of concept and pilots. The Data Solution Architect is proficient in distributed computing principles and familiar with key architectures including Lambda and Kappa architectures, and has a broad experience designing solutions using a broad set of data stores (e.g., HDFS, Azure Data Lake Store, Azure Blob Storage, Azure SQL Data Warehouse, Apache HBase, Azure DocumentDB), messaging systems (e.g., Apache Kafka, Azure Event Hubs, Azure IoT Hub) and data processing engines (e.g., Apache Hadoop, Apache Spark, Azure Data Lake Analytics, Apache Storm, Azure HDInsight). The ideal candidate will have experience in customer facing roles and success leading deep technical architecture and design discussions with senior executives.

Five plus years of experience with deep understanding of databases and analytics, including relational databases, data warehousing, big data, business intelligence and analytics.

Five plus years of success in consultative/complex technical sales and deployment projects. Technical BS degree in Computer Science or Math background desirable.

Technical Skills

Advanced analytics, analysis services (tabular, multi-dimensional), application architecture, application development, application lifecycle management (ALM), big data, business intelligence, capacity planning, cloud archival, cloud data analytics, cloud disaster recovery, cloud storage, cloud systems management, cloud systems operations, cloud transformation, compliance (PCI, HIPPA, etc.), data architecting, data cleansing, data migration (cross platform / upgrade), data modeling (physical and logical), data movement, data potency, data transformation, data warehouse design, database and server virtualization, database architecture, database design, database lifecycle management, database management, database sharding, database tuning, diagnostics, distributed application design, distributed application development, distributed database design, event sourcing, HADR / replication, health checks, identity and security, information architecture, information management, in-memory database architecture, IoT, Kappa architectures, Lambda architectures, MapReduce, master data management, mission critical DB design and architecture, modern applications, monitoring, performance tuning, polyglot resiliency, predicative analytics, reporting services design and deployment, resiliency (clustering, etc.), scalability (up and out, high performance), security architecture, security compliance, technical migration upgrades, technology architecture.

Non-Technical Skills	<p>Consultative sales, collaboration, stakeholder management, relationship management, technical oversight, technical recommendations, problem solving, risk management, architecture design session, program management, proof of concept design, technical demonstration.</p>
Technologies	<p>AML, Analytics, APS, AWS, AWS Athena, AWS Glue, AWS QuickSight, AWS RDS, AWS VM, AWS Redshift, AWS S3, Azure Active Directory, Azure Data Catalog, Azure Data Factory, Azure Data Lake, Azure Storage, Azure DocumentDB, Azure SQL Data Warehouse, Azure Import/Export, Azure HDInsight, Azure Machine Learning, Azure SQL Database, Azure Search, Azure Stream Analytics, Azure Event Hubs, Azure IoT Hub, Azure Web Jobs, Azure Cognitive Services, Birst, BizTalk, Business Objects, Cassandra, Cloudera, Cognos, Cortana Intelligence, CouchDB, Data warehouse, Database, DB2, Docker, Excel, Flume, IBM Bluemix, Google App Engine, Greenplum, Hadoop, HANA, HBase, Hive, Hortonworks, IBM, IBM Teradata, Informix, IoT Solutions, Java, Kafka, Machine Learning, MapR, MapReduce, ML, MongoDB, MariaDB, MySQL, Netezza, Networking, NoSQL, Oracle, Oracle Exadata, Oracle SOA, PBI, Pig, Power BI, PostgreSQL, Python, Qlik Tech, REST, Revolution R, Samza, SAP HANA, SAS, Security, Spark, , SQL Server, SQL Server IaaS, SQL Server Integration Services, Storage, Storm, Streaming, Sybase, Tableau, Teradata, TSQL, Virtualization, webMethods, WebSphere, YARN.</p> <p>Programming/Scripting Languages: C#, DMX, DAX, MDX, SQL, T-SQL, Java, Scala, Python, PowerShell, R, Ruby.</p> <p>Platforms: Linux (Red Hat, Ubuntu, Debian, etc.), Windows.</p>
Certifications	<p>MCSE Business Intelligence, MCSA Cloud Platform Solutions Associate, MCSA Linux on Azure Solutions Associate, MCSE Cloud Platform and Infrastructure, MCSE Data Management and Analytics, AWS Certified Solution Architect, AWS Certified Developer, AWS Certified Developer, Big Data, Certified Analytics Professional, Certificate in Engineering Excellence Big Data Analytics and Optimization (CPEE), Cloudera Certified Developer, Cloudera Certified Specialist, Data Warehousing, IBM Certified Data Architect/Engineer, Mining Massive Datasets, Graduate Certificate (Stanford), Oracle, Salesforce.com, SAP, SAS Certified Big Data Professional.</p>
Project Experience Types/Qualities	<p>Advanced analytics (including machine learning), database modernization, coordinate and execute pilots, prototypes or proof of concepts, provide validation on specific scenarios, document and share technical best practices, further customer investment, hybrid solutions on premises or in the cloud, industry-visible, large project relative to size of customer, lift and shift, migrations and upgrades (SQL, etc.), on-premises to cloud, production environment, projects where data is born in the cloud, cross-platform SQL Server migration, size of project team (complexity), significant challenges, IOT – Connected Devices, IOT- Command and Control, IOT- Data Ingestion, batch analytics, interactive analytics, real-time/streaming analytics.</p>



Data Engineer

A Data Engineer is responsible for helping to select, deploy and manage the systems and infrastructure required of a data processing pipeline in support of the customer requirements. The Data Engineer is not typically a customer facing role. The primary responsibilities include implementing ETL (extract, transform and load) pipelines, monitoring/maintaining data pipeline performance. The Big Data Engineer is proficient in distributed computing principles and familiar with key architectures including Lambda and Kappa architectures, and has a broad experience across a set of data stores (e.g., HDFS, Azure Data Lake Store, Azure Blob Storage, Azure SQL Data Warehouse, Apache HBase, Azure DocumentDB), messaging systems (e.g., Apache Kafka, Azure Event Hubs, Azure IoT Hub) and data processing engines (e.g., Apache Hadoop, Apache Spark, Azure Data Lake Analytics, Apache Storm, Azure HDInsight). The ideal candidate has three or more years' experience deploying, monitoring and managing solutions that collect, process, store and analyze huge volume of data, fast moving data or data that has significant schema variability.

Technical Skills

DevOps and DataOps, Advanced analytics, analysis services (tabular, multi-dimensional), application architecture, application lifecycle management (ALM), big data, business intelligence, capacity planning, cloud archival, cloud data analytics, cloud disaster recovery, cloud storage, cloud systems management, cloud systems operations, cloud transformation, compliance (PCI, HIPPA, etc.), data architecting, data cleansing, data migration (cross platform / upgrade), data movement, data potency, data transformation, database and server virtualization, database architecture, database design, database lifecycle management, database management, database sharding, database tuning, diagnostics, distributed application design, distributed database design, event sourcing, HADR / replication, health checks, identity and security, information architecture, information management, in-memory database architecture, IoT, Kappa architectures, Lambda architectures, MapReduce, master data management, mission critical database design and architecture, modern applications, monitoring, performance tuning, polyglot resiliency, predicative analytic pipelines, reporting services design and deployment, resiliency (clustering, etc.), scalability (up and out, high performance), security architecture, security compliance, technical migration upgrades, technology architecture.

Non-Technical Skills	Technical recommendations, problem solving, risk management, proof of concept design, technical demonstration, consultative requirements clarification and issue troubleshooting.
Technologies	<p>Azure Machine Learning, Analytics, APS, AWS, AWS Athena, AWS Glue, AWS QuickSight, AWS RDS, AWS VM, AWS Redshift, AWS S3, Azure Active Directory, Azure Data Catalog, Azure Data Factory, Azure Data Lake, Azure Storage, Azure DocumentDB, Azure SQL Data Warehouse, Azure Import/Export, Azure HDInsight, Azure SQL Database, Azure Search, Azure Stream Analytics, Azure Event Hubs, Azure IoT Hub, Azure Web Jobs, Azure Cognitive Services, Birst, Business Objects, Cassandra, Cloudera, Cognos, Cortana Intelligence, CouchDB, Data warehouse, Database, DB2, Docker, Excel, Flume, IBM Bluemix, Google App Engine, Greenplum, Hadoop, HANA, HBase, Hive, Hortonworks, IBM, IBM Teradata, Informix, IoT Solutions, Java, Kafka, Machine Learning, MapR, MapReduce, ML, MongoDB, MariaDB, MySQL, Netezza, Networking, NoSQL, Oracle, Oracle Exadata, Oracle SOA, PBI, Pig, Power BI, PostgreSQL, Python, Qlik Tech, REST, Revolution R, Samza, SAP HANA, SAS, Security, Spark, SQL Server, SQL Server IaaS, SQL Server Integration Services, Storage, Storm, Streaming, Sybase, Tableau, Teradata, TSQL, Virtualization, webMethods, WebSphere, YARN.</p> <p>Programming/Scripting Languages: C#, DMX, DAX, MDX, SQL, T-SQL, Java, Scala, Python, PowerShell, R, Ruby.</p> <p>Platforms: Linux (Red Hat, Ubuntu, Debian, etc.), Windows.</p>
Certifications	MCSE Business Intelligence, MCSA Cloud Platform Solutions Associate, MCSA Linux on Azure Solutions Associate, MCSE Cloud Platform and Infrastructure, MCSE Data Management and Analytics, AWS Certified Solution Architect, Big Data, Certified Analytics Professional, Certificate in Engineering Excellence Big Data Analytics and Optimization (CPEE), Cloudera Certified Data Engineer, Data Warehousing, IBM Certified Data Architect/Engineer, Mining Massive Datasets Graduate Certificate (Stanford), Oracle, Salesforce.com, SAP, SAS Certified Big Data Professional.
Project Experience Types/Qualities	Advanced analytics (including machine learning), automated data pipelines, database modernization, further customer investment, hybrid solutions on premises or in the cloud, industry-visible, large project relative to size of customer, lift and shift, migrations and upgrades (SQL, etc.), on-premises to cloud, production environment, projects where data is born in the cloud, cross-platform SQL Server migration, size of project team (complexity), significant challenges, IOT – Connected Devices, IOT- Command and Control, IOT- Data Ingestion, batch analytics, interactive analytics, real-time/streaming analytics.

Data Developer

A Data Developer is responsible for helping to select and implement the tools and processes required of a data processing pipeline in support of the customer requirements. The Data Developer may be a customer facing role, but the primary responsibilities include implementing ETL (extract, transform and load) pipelines, monitoring/maintaining data pipeline performance and implementing big data or advanced analytics solutions. The Data Developer is proficient in distributed computing principles and familiar with key architectures including Lambda and Kappa architectures, and has a broad experience across a set of data stores (e.g., HDFS, Azure Data Lake Store, Azure Blob Storage, Azure SQL Data Warehouse, Apache HBase, Azure DocumentDB), messaging systems (e.g., Apache Kafka, Azure Event Hubs, Azure IoT Hub) and data processing engines (e.g., Apache Hadoop, Apache Spark, Azure Data Lake Analytics, Apache Storm, Azure HDInsight). The ideal candidate has three or more years' experience working on solutions that collect, process, store and analyze huge volume of data, fast moving data or data that has significant schema variability.

Technical Skills

Advanced analytics, analysis services (tabular, multi-dimensional), application architecture, application development, application lifecycle management (ALM), big data, business intelligence, cloud data analytics, cloud storage, cloud systems management, cloud systems operations, cloud transformation, compliance (PCI, HIPPA, etc.), data architecting, data cleansing, data modeling (physical and logical), data movement, data potency, data transformation, data warehouse design, database architecture, database design, database lifecycle management, database management, database sharding, database tuning, diagnostics, distributed application design, distributed application development, distributed database design, event sourcing, identity and security, information architecture, information management, in-memory database architecture, IoT, Kappa architectures, Lambda architectures, MapReduce, master data management, mission critical database design and architecture, modern applications, monitoring, performance tuning, polyglot resiliency, predictive analytics, reporting services design and deployment, resiliency (clustering, etc.), scalability (up and out, high performance), security architecture, technical migration upgrades, technology architecture.

Non-Technical Skills

Technical recommendations, problem solving, risk management, proof of concept design, technical demonstration, consultative requirements clarification and issue troubleshooting.

Technologies

Azure Machine Learning, Analytics, AWS, AWS Athena, AWS Glue, AWS QuickSight, AWS RDS, AWS VM, AWS Redshift, AWS S3, Azure Active Directory, Azure Data Catalog, Azure Data Factory, Azure Data Lake, Azure Storage, Azure DocumentDB, Azure SQL Data Warehouse, Azure Import/Export, Azure HDInsight, Azure SQL Database, Azure Search, Azure Stream Analytics, Azure Event Hubs, Azure IoT Hub, Azure Web Jobs, Azure Cognitive Services, Birst, Business Objects, Cassandra, Cloudera, Cognos, Cortana Intelligence, CouchDB, Data warehouse, Database, DB2, Docker, Excel, Flume, IBM Bluemix, Google App Engine, Greenplum, Hadoop, HANA, HBase, Hive, Hortonworks, IBM, IBM Teradata, Informix, IoT Solutions, Java, Kafka, Machine Learning, MapR, MapReduce, ML, MongoDB, MariaDB, MySQL, Netezza, Networking, NoSQL, Oracle, Oracle Exadata, Oracle SOA, PBI, Pig, Power BI, PostgreSQL, Python, Qlik Tech, REST, Revolution R, Samza, SAP HANA, SAS, Security, Spark, SQL Server, SQL Server IaaS, SQL Server Integration Services, Storage, Storm, Streaming, Sybase, Tableau, Teradata, TSQL, webMethods, WebSphere, YARN.

Programming/Scripting Languages: C#, DMX, DAX, MDX, SQL, T-SQL, Java, Scala, Python, PowerShell, R, Ruby.

Platforms: Linux (Red Hat, Ubuntu, Debian, etc.), Windows.

Certifications

MCSE Business Intelligence, MCSA Cloud Platform Solutions Associate, MCSA Linux on Azure Solutions Associate, MCSE Cloud Platform and Infrastructure, MCSE Data Management and Analytics, AWS Certified Solution Architect, AWS Certified Developer, AWS Certified Developer, Big Data, Certified Analytics Professional, Certificate in Engineering Excellence Big Data Analytics and Optimization (CPEE), Cloudera Certified Associate Spark and Hadoop Developer, Data Warehousing, IBM Certified Data Architect/Engineer, Mining Massive Datasets Graduate Certificate (Stanford), Oracle, Salesforce.com, SAP, SAS Certified Big Data Professional.

Project Experience Types/Qualities

Advanced analytics (including machine learning), database modernization, further customer investment, hybrid solutions on premises or in the cloud, industry-visible, large project relative to size of customer, migrations and upgrades (SQL, etc.), on-premises to cloud, production environment, projects where data is born in the cloud, cross-platform SQL Server migration, size of project team (complexity), significant challenges, IOT – Connected Devices, IOT- Command and Control, IOT- Data Ingestion, batch analytics, interactive analytics, real-time/streaming analytics.



Cloud Architect

A Cloud Architect (CA) drives high-priority customer initiatives in collaboration with customers and your sales team. The CA is a technical, customer-facing role that is accountable for the end-to-end customer cloud deployment experience. CAs own the Azure technical customer engagement, including: architectural design sessions, specific implementation projects and/or proofs of concepts. The ideal candidate will have experience in customer-facing roles and success leading deep technical architecture and application design discussions with senior customer executives to drive cloud deployment. Bachelor's degree in computer science or related field preferred.

Technical Skills

Solid understanding of modern authentication protocols and a background in cyber security. Deep understanding of cloud computing tech, business drivers, and emerging computing trends. Deep technical experience in enterprise mobile, identity & access control, and security solutions. Working knowledge with AGILE development, SCRUM and Application Lifecycle Management (ALM) with one or more of the following programming languages: PowerShell, Bash, .NET, C++, Java, JSON, PHP, Perl, Python, Ruby on Rails, HTML, CSS, JavaScript, Responsive Web Design.

Non-Technical Skills

Proven track record of building deep technical relationships with senior executives and growing cloud consumption share in large or highly strategic accounts, driving decisions, resolving conflicts & ensuring follow through. Presentation skills with a high degree of comfort with both large and small audiences. Prior work experience in a consulting/architecture position within a software/services company. Problem-solving mentality leveraging internal and/or external resources. Exceptional verbal and written communication.

Certifications

MCSE Cloud Platform and Infrastructure, CompTIA Security+, CISSP, MCSA Cloud Platform Solutions Associate, MCSA Linux on Azure Solutions Associate, AWS Certified Solution Architect. Exam priorities: [Architecting Azure Solutions 70-534](#), [70-398: Planning for and Measuring Devices in the Enterprise](#).

Project Experience Types / Qualities

5+ years of architecture, design, implementation, and/or support of highly distributed applications (i.e. having an architectural sense for ensuring availability, reliability, etc.). 2+ years of experience in "migrating" on premise workloads to the cloud. 5+ years of success in consultative/complex technical sales and deployment projects (where necessary, managing various stakeholder relationships to get consensus on solution/projects). Oversight experience on major transformation projects and successful transitions to implementation support teams.

Technologies

Enterprise Mobility Suite, Intune, Azure Information Protection, Azure Active Directory, Octa, Auth0, LDAP, OAuth, SAML, Cloud App Security, Firewalls, Office 365, Windows Server Active Directory, Azure AD Connect, Active Directory Federation Services (ADFS), Mobile Iron, Airwatch, iOS, Android, Windows, Azure Virtual Machines, Virtual Networks, ExpressRoute, Operations Management Suite, Azure Site Recovery, Azure Backup, Azure App Services, Azure Storage, Azure Import/Export, Azure SQL Database, Azure Web Jobs, Azure ExpressRoute, MySQL, SQL Server, SQL Server IaaS, SharePoint on Azure, AWS EC2, S3, AWS DirectConnect.

Recruiting Resources

Top 10 Sources to Find Skilled Labor and What to Look For

Sourcing skilled labor can be a challenge. In our recent survey with MDC of 1,136 Azure partners, we found that referrals and LinkedIn rank among the top sources for finding candidates. For Data Science candidates, our interviews with partners emphasize recruiting from universities.

	TOTAL (n=1136)	SMB (n=886)	ENTERPRISE (n=250)
Referrals from employees or partnerships	70%	69%	73%
LinkedIn	59%	57%	66%
Posting on website	47%	45%	54%
Local universities	38%	36%	46%
Local technical communities	36%	35%	43%
Recruit from competitors	30%	29%	36%
Meet ups	29%	29%	30%
GitHub	8%	8%	7%
Stack Overflow	8%	8%	7%
Other job posting sites	6%	6%	4%

Source: Microsoft Cloud Practice Development Study, MDC Research, November 2016.

Now that you have an understanding of where to look, what are the most important factors to look for in a potential hire's skillset? In the Microsoft Cloud Practice Development Study, we asked the Azure partners this question. What they told us was that the top three most important factors were work history, cultural fit, and years of experience.

	TOTAL (n=1136)	SMB (n=886)	ENTERPRISE (n=250)
Work history	69%	68%	74%
Cultural fit	43%	40%	53%
Years of experience	42%	41%	47%
Professional certifications	32%	34%	22%
Referrals	28%	29%	26%
Professional training received	20%	21%	16%
Reputation through community	16%	15%	19%
Formal education	13%	14%	11%
Contract-to-hire or other means to test skills "hands-on"	13%	13%	13%
Publications	3%	3%	2%
Awards received	2%	2%	2%
Attitude	0%	0%	1%
Other	4%	4%	4%

Source: Microsoft Cloud Practice Development Study, MDC Research, November 2016.

Training & Readiness

Nurturing Data Science Capabilities

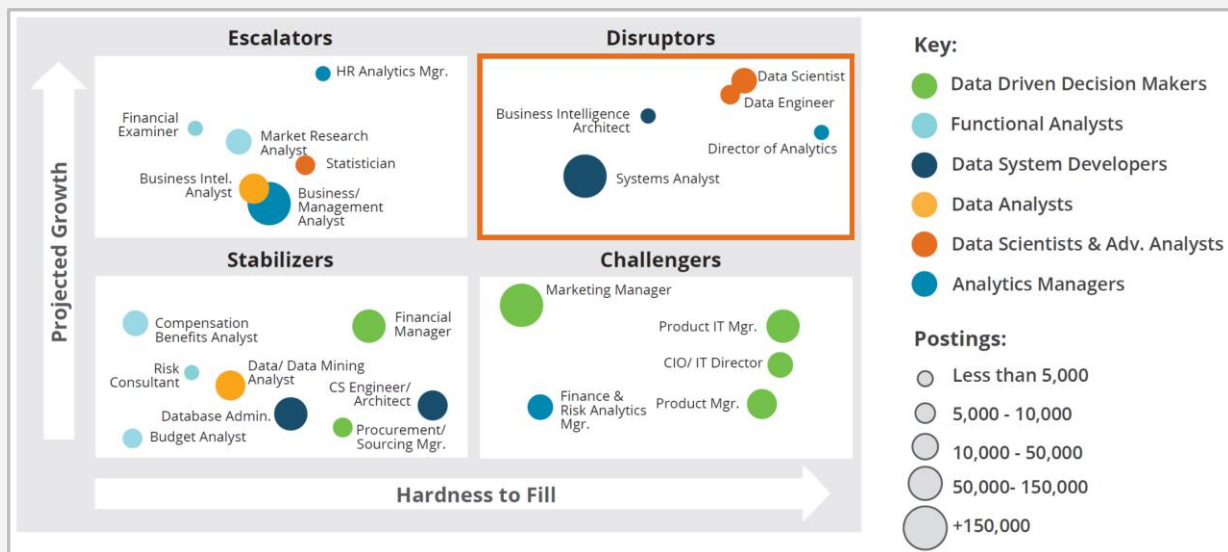
There is a shortage of data scientists, how can your practice grow in this climate?

In its [2016 report, the McKinsey Global Institute](#) states that the U.S. economy alone could be short by as many as 250,000 data scientists by 2024. According to the [2017 study by IBM, The Quant Crunch](#), it takes an average of 45 days to hire a data scientist, 5 days longer than the average position. With an expected shortage and elongated hiring lead times, the choice of what your practice can do now is a critical one. You are faced with a choice of either to wait to hire your next data scientist or nurture data science capabilities in existing resources.

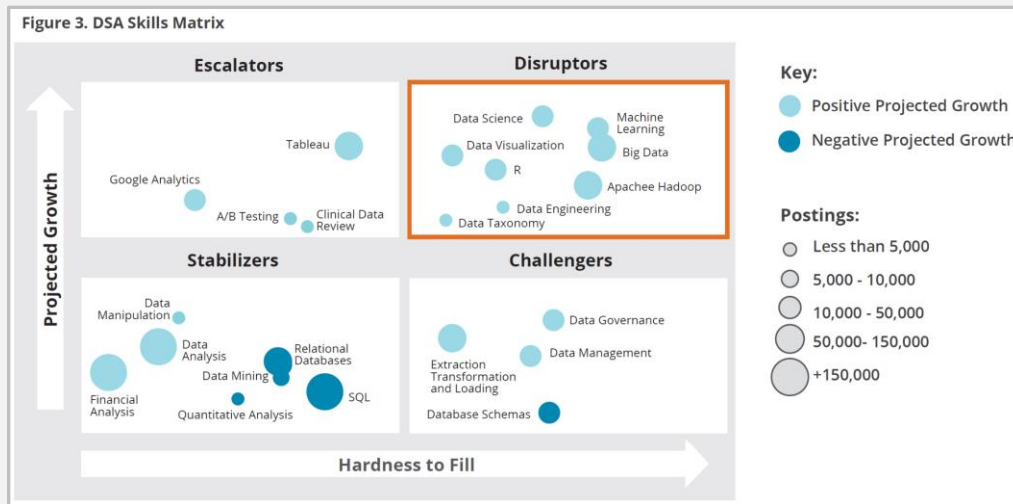
In The Quant Crunch report, those job roles placed in the disruptors category are deemed as:

- Having a high cost to hire
- Having strong needs for new training programs
- Creating the most risk to success in their absence

The report identified Data Scientist and Data Engineer as the two most significant such roles in terms of the difficulty to fill and the projected growth in demand for fulfilling that role.

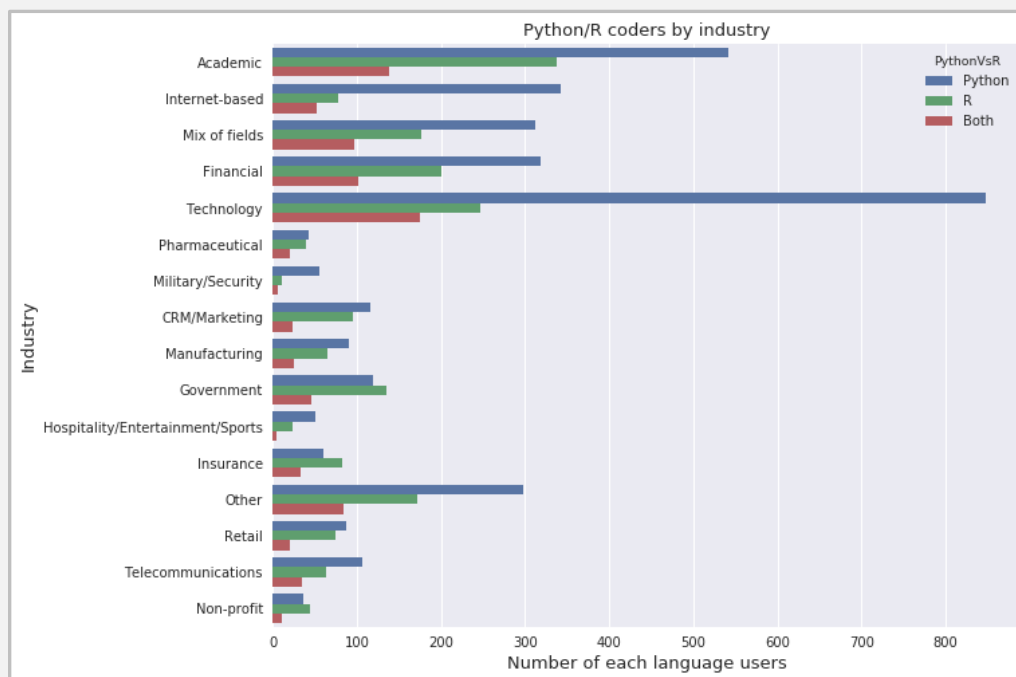


The skills that are most in demand for these disruptive job roles naturally veer towards big data, machine learning and data science acumen.



Source: <https://www.kaggle.com/mhjabri/what-do-kagglers-say-about-data-science>

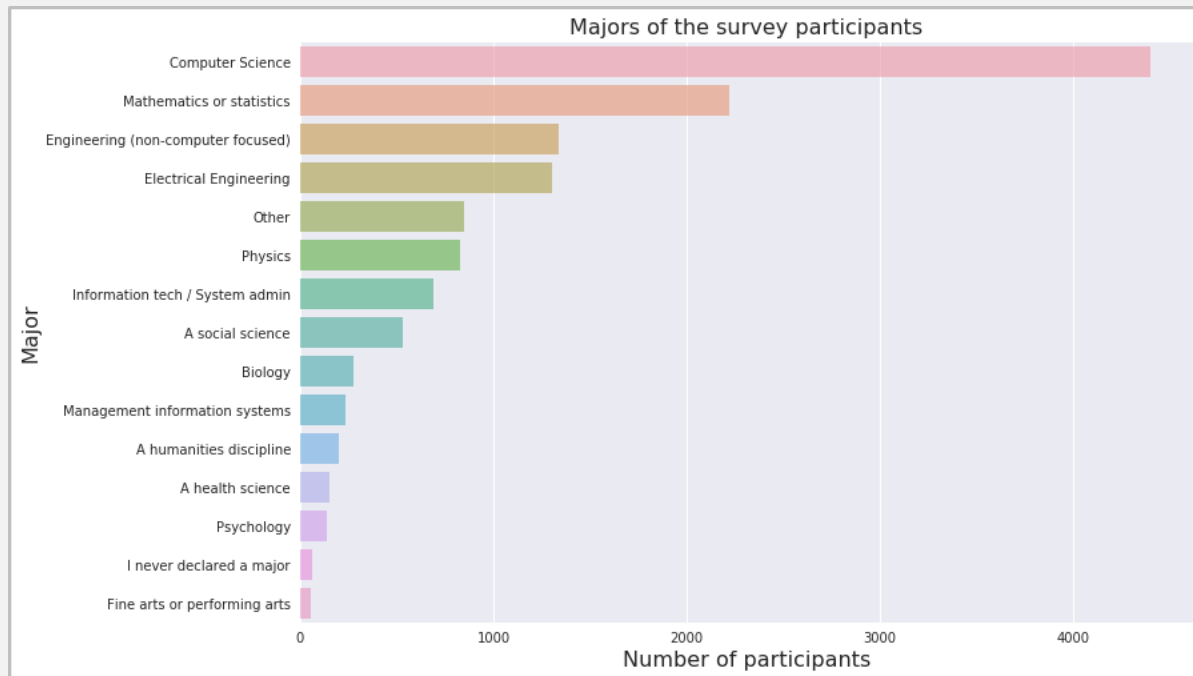
It is worth noting that the data for the above was collected in 2016. The year 2017 has seen significant growth in Python superseding R as the language of data science solutions. [The 2017 Kaggle ML and Data Science Survey](#) found that Python is indeed the dominant language of choice across industries.



Source: <https://www.kaggle.com/mhjabri/what-do-kagglers-say-about-data-science>

WHY NURTURE DEVELOPERS?

It is insightful to understand the education behind the aforementioned data science skills. Kaggle has created a community of data scientists and in their [2017 Kaggle ML and Data Science Survey](#) of 16,000 industry respondents, it is worthwhile to note that the most significant educational background: computer science. This may surprise those who would think the most common major is mathematics or statistics. In other words, most data scientists start as developers.



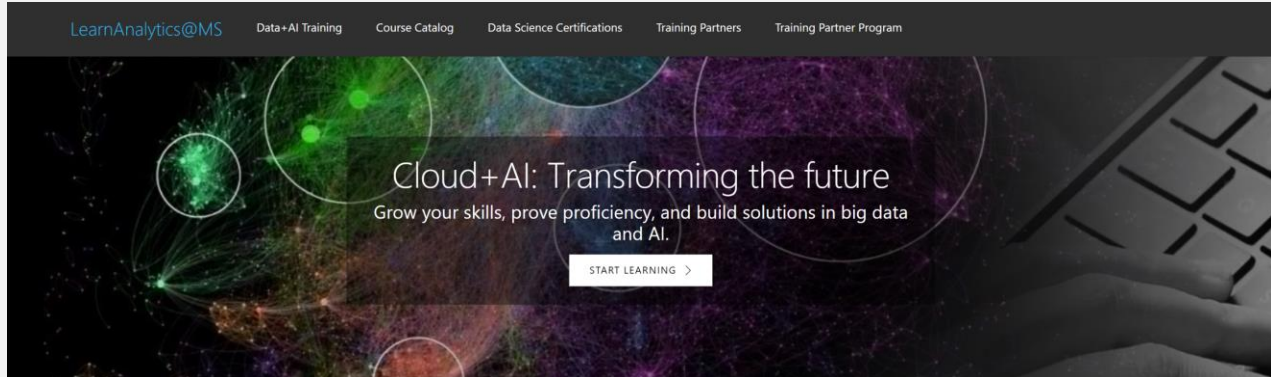
Source: <https://www.kaggle.com/mhjabri/what-do-kagglers-say-about-data-science>

While some developers will scoff at going back to the math they had to take in college, most are primed with enough of a mathematical aptitude and interest to learn the requisite statistics, probability etc. required of data science.

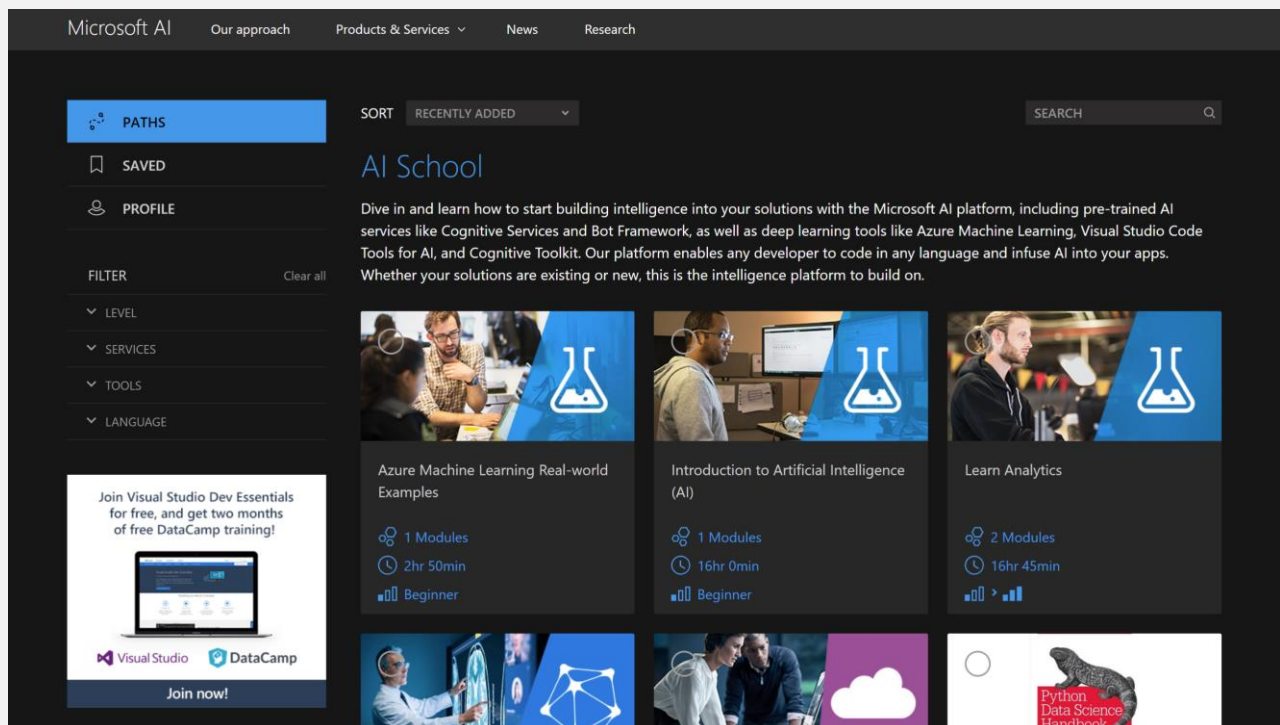
While Microsoft does provide numerous machine learning tools for the C# developer, such as Cognitive Services SDK's, Core ML for Xamarin, Acoord.NET for audio and image processing and the CNTK (Cognitive Toolkit) library in C#, an important challenge to the typical Microsoft development team is that .NET developers will need to learn Python because it is so pervasive in the tools of data science, and is likely to be the first language into which innovative contributions from the community are made available.

HOW TO NURTURE DATA SCIENCE CAPABILITIES?

When launching training initiatives, partners should begin by examining the resources at the [Microsoft Learn Analytics](#) site, which provides access to self-paced training and certification materials, as well as a curated list of training partners Microsoft has validated.



Additionally, partners should explore the wealth of resources Microsoft provides in the [AI School](#) website. The learning paths provide a curated collection of tutorials and guides for all skill levels from beginner, to intermediate, to advanced presented in an easily searchable fashion that enables the partner team to find the right content, at the right level and for the amount of time they have to commit to the effort.



Partners should also consider the following resources from the Microsoft Team Data Science Process for training existing team members and nurturing their data science capabilities:

- ➔ [Learning Path: TDSP for Data Scientists](#)
- ➔ [Learning Path: TDSP for DevOps](#)
- ➔ [Using the Team Data Science Process with the new Azure Machine Learning](#)

Besides technical training and certification (which is covered in the sections that follow), there are non-traditional ways to level up the team's data science skill. Some of the partners interviewed for this book suggested:

- **Innovation hours:** a percentage of the employee's weekly time (as much as 20%) is allotted to let them work on their own projects of interest, alone or with fellow employees. Partners suggest that the creativity unleashed during these free thinking projects have yielded insights and intellectual property.
- **Hackathons, competitions and community labs:** Partners also suggested encouraging employees to participate in community events that leverage open source data sets to try and solve specific problems. These range from 1-2 day hackathons where the team is fully engaged in the project at hand, typically over a weekend. The data science community is famous for its competitions, such as [Kaggle](#), where entrants compete for monetary awards in teams trying to solve difficult, real-world problems. Finally, community labs are opportunities for employees to work together with academia, domain experts and others learning the trade to help the local community solve worthwhile data challenges.
- **Partner with startups:** Some partners interviewed expressed a novel approach to skilling up the team. In the absence of enterprise customers, they have partnered with startups working on interesting problems requiring data science. The start-up receives support from the partner at no or low cost (or for some small amount of equity) and partner team gets real-world experience.

Additional Training Resources for the Data Science Team

[Microsoft Learning Partners](#) are available worldwide to help enable your team for Microsoft Azure via live instructor-led training. This can be scheduled as a dedicated delivery at your location or virtually using remote learning technologies. Many courses are scheduled as open-enrollment courses, which doesn't require you to schedule a dedicated class.

- [Pluralsight](#) is a key Microsoft partner that offers Azure training, as well as courses related to data science. Gain the know-how and confidence your job demands through these free online courses, delivered in partnership with Pluralsight.
- [O'Reilly Safari](#) provides subscription access to more than 40,000 books, videos, and interactive tutorials from over 200 of the world's best publishers, including O'Reilly, Pearson, Harvard Business Review, and Packt. It also offers live online training courses led by instructors from O'Reilly's network of tech innovators and expert practitioners. O'Reilly safari has a significant collection of materials relevant to the AI practice, including materials on the Data Science process, frameworks for deep Learning and machine learning, tools like Jupyter Notebooks, data platforms like Apache Spark, and the programming languages of data science like Python and R.



Preparing and Training Technical Staff for the Cloud

For technical staff to function as change agents supporting current and emerging cloud technologies, their buy-in for the use and integration of these technologies is needed. For this, staff need three things:

- An understanding of their roles and any changes to their current position.
- Time and resources to explore the technologies.
- An understanding of the business case for the technologies.



Use the following resources as part of your Azure onboarding for new and existing staff:

- [Azure Training and Certification](#) provides free online training options including online courses, learning paths, hands-on labs as well as resources to help you find learning partners who can help you achieve your skills development goals using Microsoft Azure services.
- [Microsoft Azure Hands-on Labs](#) provides free, self-paced labs to help you stay current with Azure. The live environments are fully self-contained. You do not need your own Azure subscription to complete the labs, just login with a remote desktop (RDP) client and get started.
- [Microsoft Virtual Academy](#) offers training from the people who helped to build Microsoft Azure. From the basic overview to deep technical training, IT staff will learn how to leverage Microsoft Azure for their business.
- [Microsoft IT Pro Cloud Essentials](#) is a free annual subscription that includes cloud services, education, and support benefits. IT Pro Cloud Essentials provides IT implementers with hands-on experience, targeted educational opportunities, and access to experts in areas that matter most to increase knowledge and create a path to career advancement.
- [The Microsoft IT Pro Career Center](#) is a free online resource to help map your cloud career path. Learn what industry experts suggest for your cloud role and the skills to get you there.
- [Microsoft Learning](#) offers a wide variety of official curriculum on-demand, as well as [edX courses that are taught by Microsoft experts](#), and help you learn through hands-on experiences with a broad reach of Azure technologies.
- [The Microsoft Partner Network \(MPN\) Learning Portal](#) provides a centralized interface with training opportunities and certification options organized by products, competencies, certifications, and job role.

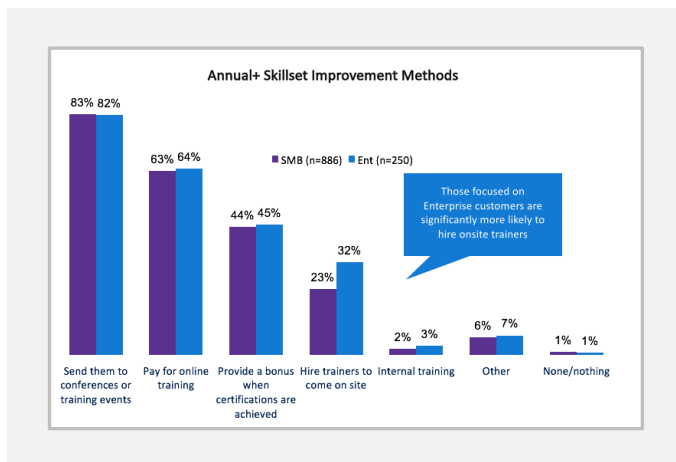
Follow a learning curriculum at your own pace to build the skills you need most to stay relevant. Suggested resources to help onboard your team for training success are available in this section.

General Technical Training

Whether you need to fill a skills gap or are looking to improve your team's skill surface area, technical training is critical to your success.



In our research, we found conferences and paid online training are the most common learning mechanisms.



Source: Microsoft Cloud Practice Development Study, MDC Research, November 2016

CLOUD AND ENTERPRISE PARTNER RESOURCES

The [Cloud and Enterprise Partner Resources Portal](#) provides a source of sales and technical training for partner practices and key areas of specialization. Resources include customer success stories, sales and technical training, tools, engines, and resources available to help build your skills around selling, deploying, and architecting cloud infrastructure and management, cloud application development, data platform and analytics, and security and compliance solutions.

MPN LEARNING PORTAL

The [Microsoft Partner Network \(MPN\) Learning Portal](#) provides a centralized interface with training opportunities and certification options organized by products, competencies, certifications, and job role.

CLOUD + ENTERPRISE UNIVERSITY ONLINE

Leverage the [Cloud + Enterprise University Online](#) to build knowledge, stay sharp, and prove your expertise on selling and supporting Microsoft cloud solutions through our live and on demand webcasts and virtual, instructor-led courses—giving you the flexibility to train at your own pace.

MICROSOFT INSPIRE CONFERENCE RECORDINGS

Even if you missed the annual live event, the [Microsoft Inspire Conference](#) provides many of its sessions as on-demand recordings — no conference pass required.

PARTNER COMMUNITY EVENTS, CALLS & WEBINARS

The [Microsoft Partner Enablement Blog](#) maintains a schedule of trainings available to partners. Visit often and plan your training calendar.

SMART PARTNER MARKETING

Leverage the [Microsoft Smart Partner Marketing](#) site as your starting point for training marketing resources.



“

You can get started in AI simply as a user of a library or AI service. Then go deeper increasing your breadth and depth of mastery. It's not all or nothing before you start creating value.

BOB SCHMIDT

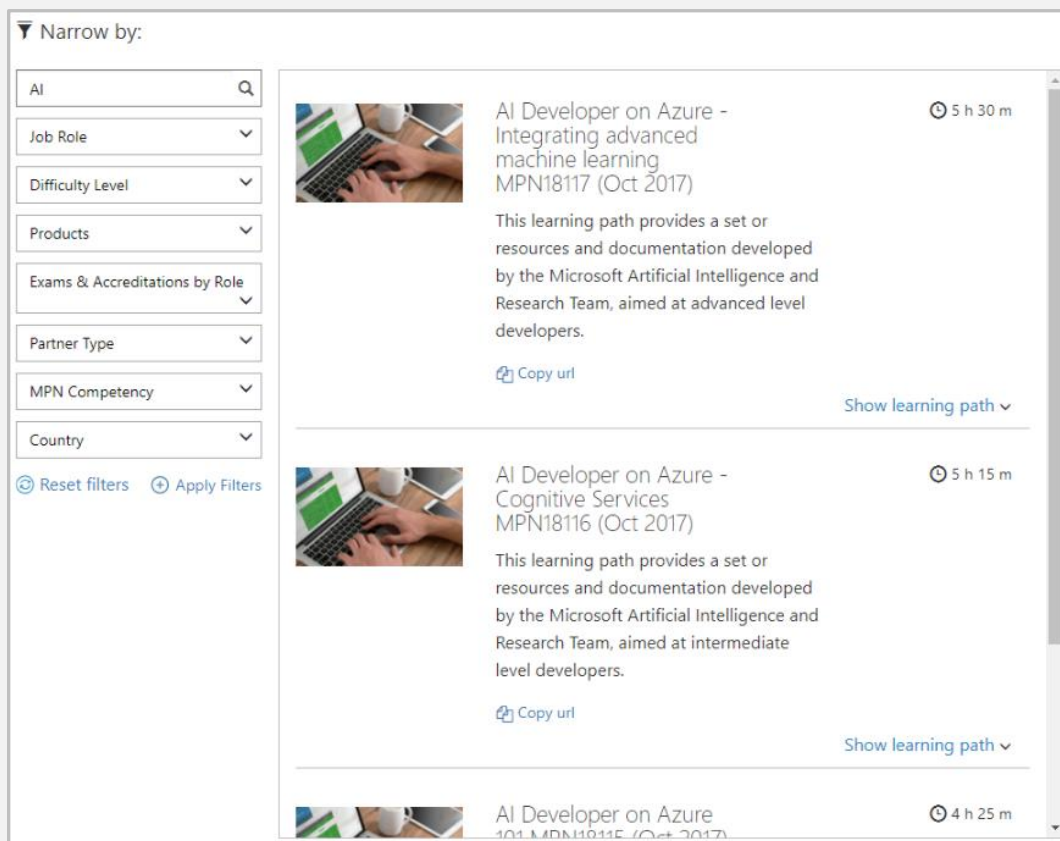
Freelance Data Scientist

Increase Readiness and Marketability with Learning Paths and Assessments, Competencies & Certifications

There are numerous assessments and certifications your team should consider as motivation for advancing their skills, creating proof points for your practice and enabling you to achieve Microsoft Partner Network Competencies.

LEARNING PATHS & ASSESSMENTS

Training content for the assessments can be found in the learning paths available from the [Learning Portal](#). Use the search field to identify new content as it becomes available.



The screenshot displays a search results page for learning paths. On the left, a 'Narrow by' sidebar contains several filter categories: AI (with a search input), Job Role, Difficulty Level, Products, Exams & Accreditations by Role, Partner Type, MPN Competency, and Country. Below these filters are 'Reset filters' and 'Apply Filters' buttons. The main content area shows three search results, each featuring a thumbnail image of hands on a laptop, the title of the learning path, its duration, and a brief description. The first result is 'AI Developer on Azure - Integrating advanced machine learning' (MPN18117, Oct 2017) with a duration of 5 h 30 m. The second is 'AI Developer on Azure - Cognitive Services' (MPN18116, Oct 2017) with a duration of 5 h 15 m. The third is 'AI Developer on Azure' (MPN18115, Oct 2017) with a duration of 4 h 25 m. Each result includes a 'Copy url' button and a 'Show learning path' dropdown menu.

Learning Path Title	MPN ID	Effective Date	Duration
AI Developer on Azure - Integrating advanced machine learning	MPN18117	Oct 2017	5 h 30 m
AI Developer on Azure - Cognitive Services	MPN18116	Oct 2017	5 h 15 m
AI Developer on Azure	MPN18115	Oct 2017	4 h 25 m

COMPETENCIES

One of the next steps is to ensure you align the technical team to the MPN competency for your practice.

The competencies applicable to the AI Practice are:

- [Data Analytics competency](#)
- [Data Platform competency](#)
- [Cloud Platform competency](#)

One of the next steps is to ensure you align the technical team to the MPN competency for your practice.

The following tables summarize the **skill requirements** needed by people in your organization to achieve either a Gold or Silver Competency for the competencies relevant to the AI practice. Some competencies have alternative options your organization can elect to meet to achieve the competency. You only need to meet the requirements of one option in any given competency.

DATA ANALYTICS COMPETENCY	SILVER REQUIREMENTS	GOLD REQUIREMENTS
Option 1: Data Analytics Specialist Option	<p>One individual must pass all of the exams in any focus area:</p> <p>Business Intelligence focus area:</p> <ul style="list-style-type: none"> • Exam 70-767: Implementing a SQL Data Warehouse • Exam 70-768: Developing SQL Data Models (Beta) <p>Advanced Analytics focus area:</p> <ul style="list-style-type: none"> • Exam 70-773: Analyzing Big Data with Microsoft R • Exam 70-774: Perform Cloud Data Science with Azure Machine Learning <p>Big Data focus Area:</p> <ul style="list-style-type: none"> • Exam 70-475: Designing and Implementing Big Data Analytics Solutions • Exam 70-775: Perform Data Engineering on Microsoft HDInsight <p>And, the same individual must pass the following assessment:</p> <ul style="list-style-type: none"> • Technical Assessment Data Analytics 	<p>Two individuals must pass all of the exams in any focus area:</p> <p>Business Intelligence focus area:</p> <ul style="list-style-type: none"> • Exam 70-767: Implementing a SQL Data Warehouse • Exam 70-768: Developing SQL Data Models (Beta) <p>Advanced Analytics focus area:</p> <ul style="list-style-type: none"> • Exam 70-773: Analyzing Big Data with Microsoft R • Exam 70-774: Perform Cloud Data Science with Azure Machine Learning <p>Big Data focus Area:</p> <ul style="list-style-type: none"> • Exam 70-475: Designing and Implementing Big Data Analytics Solutions • Exam 70-775: Perform Data Engineering on Microsoft HDInsight <p>And, the same individuals must pass the following assessment:</p> <ul style="list-style-type: none"> • Technical Assessment Data Analytics

Option 2: Data Analytics Beginners Option

Two individuals must pass the following assessment:

- [Technical Assessment Data Analytics Foundational](#)

And, they both must pass one of the following assessments:

- [Technical Assessment for Microsoft Power BI Data Analytics](#)
- [Technical Assessment Advanced Analytics for Data Analytics](#)
- [Technical Assessment Big Data for Data Analytics](#)

Four individuals must pass the following assessment:

- [Technical Assessment Data Analytics Foundational](#)

And, the same four must pass one of the following assessments:

- [Technical Assessment for Microsoft Power BI Data Analytics](#)
- [Technical Assessment Advanced Analytics for Data Analytics](#)
- [Technical Assessment Big Data for Data Analytics](#)

DATA PLATFORM COMPETENCY

SILVER REQUIREMENTS

GOLD REQUIREMENTS

Option 1: SQL Database Specialist Option

One individual must pass all of the following exams in any one of the two focus areas:

Database Administration Focus:

- [Exam 70-764: Administering a SQL Database Infrastructure](#)
- [Exam 70-765: Provisioning SQL Databases](#)

Database Development Focus:

- [Exam 70-473: Designing and Implementing Cloud Data Platform Solutions](#)
- [Exam 70-761: Querying Data with Transact-SQL](#)
- [Exam 70-762: Developing SQL Databases](#)

And the same individual must also pass:

- [Technical Assessment Data Platform Foundational](#)

Two individuals must pass all of the following exams in any one of the two focus areas:

Database Administration Focus:

- [Exam 70-764: Administering a SQL Database Infrastructure](#)
- [Exam 70-765: Provisioning SQL Databases](#)

Database Development Focus:

- [Exam 70-473: Designing and Implementing Cloud Data Platform Solutions](#)
- [Exam 70-761: Querying Data with Transact-SQL](#)
- [Exam 70-762: Developing SQL Databases](#)

And the same two individuals must also pass:

- [Technical Assessment Data Platform Foundational](#)

Option 2: Big Data Option

Two individuals must pass the following assessment:

- [Technical Assessment Data Platform Foundational](#)

And the same two individuals must pass one of the following assessments:

- [Technical Assessment for SQL Server 2016](#)
- [Technical Assessment for Azure Data Warehouse](#)
- [Technical Assessment for Azure Data Services](#)


Four individuals must pass the following assessment:

- [Technical Assessment Data Platform Foundational](#)

And the same four individuals must pass one of the following assessments:

- [Technical Assessment for SQL Server 2016](#)
- [Technical Assessment for Azure Data Warehouse](#)
- [Technical Assessment for Azure Data Services](#)

CLOUD PLATFORM COMPETENCY	SILVER REQUIREMENTS	GOLD REQUIREMENTS
Option 1: Azure Consumption Option	<p>One individual must pass one of the following assessments:</p> <ul style="list-style-type: none"> • Technical Assessment for Cloud Platform • Technical Assessment for Remote Desktop Services on Azure • Technical Assessment for Using Azure for Data Analytics and Data Platform Solutions • Technical Assessment for Using Microsoft Azure for Application Development • Technical Assessment for Using Azure for Internet of Things Solutions <p>Or, one individual must pass one of the following exams:</p> <ul style="list-style-type: none"> • Exam 70-532: Developing Microsoft Azure Solutions • Exam 70-533: Implementing Microsoft Azure Infrastructure Solutions • Exam 70-473: Designing and Implementing Cloud Data Platform Solutions • Exam 70-475: Designing and Implementing Big Data Analytics Solutions • MCSA: Linux on Azure 	<p>Two individuals each must complete one of the following assessments:</p> <ul style="list-style-type: none"> • Technical Assessment for Cloud Platform • Technical Assessment for Remote Desktop Services on Azure • Technical Assessment for Using Azure for Data Analytics and Data Platform Solutions • Technical Assessment for Using Microsoft Azure for Application Development • Technical Assessment for Using Azure for Internet of Things Solutions <p>Or, two individuals each must pass one of the following exams:</p> <ul style="list-style-type: none"> • Exam 70-532: Developing Microsoft Azure Solutions • Exam 70-533: Implementing Microsoft Azure Infrastructure Solutions • Exam 70-473: Designing and Implementing Cloud Data Platform Solutions • Exam 70-475: Designing and Implementing Big Data Analytics Solutions • MCSA: Linux on Azure



CERTIFICATIONS

Increase readiness and marketability with MCSA or MCSE certifications.

There are numerous assessments and certifications your team should consider as motivation for advancing their skills, creating proof points for your practice, earning [certification badges](#), and enabling you to achieve Microsoft Partner Network Competencies.

TITLE	DESCRIPTION	REQUIRED EXAMS
MCSA: MACHINE LEARNING	Demonstrate your expertise in operationalizing Microsoft Azure machine learning and Big Data with R Server and SQL R Services.	<ul style="list-style-type: none"> • 70-773: Analyzing Big Data with Microsoft R • 70-774: Perform Cloud Data Science with Azure Machine Learning
MCSA: CLOUD PLATFORM	Demonstrate your expertise in Microsoft cloud-related technologies to reduce IT costs and deliver more value for the modern business.	Pass two of the following exams: <ul style="list-style-type: none"> • 70-532: Developing Microsoft Azure Solutions • 70-533: Implementing Microsoft Azure Infrastructure Solutions • 70-535: Architecting Microsoft Azure Solutions • 70-537: Configuring and Operating a Hybrid Cloud with Microsoft Azure Stack
MCSA: LINUX ON AZURE	Demonstrate your ability to design, architect, implement, and maintain complex cloud-enabled Linux® solutions that leverage Microsoft Azure open source capabilities. This certification also validates your Linux system administration skills to show that you are fluent in today's cloud-native world.	<ul style="list-style-type: none"> • 70-533: Implementing Microsoft Azure Infrastructure Solutions • Linux Foundation Certified System Administrator

TITLE	DESCRIPTION	REQUIRED EXAMS
MCSE: CLOUD PLATFORM AND INFRASTRUCTURE	<p>The Microsoft Certified Solutions Expert (MCSE): Cloud Platform and Infrastructure certification validates that you have the skills needed to run a highly efficient and modern data center, with expertise in cloud technologies, identity management, systems management, virtualization, storage, and networking.</p>	<p>Pre-requisites:</p> <ul style="list-style-type: none"> • MCSA – Window Server 2016 • MCSA – Cloud Platform • MCSA – Linux on Azure • MCSA – Windows Server 2012 <p>Pass one of the following Azure exams:</p> <ul style="list-style-type: none"> • 70-532: Developing Microsoft Azure Solutions • 70-533: Implementing Microsoft Azure Infrastructure Solutions (recommended) • 70-535: Architecting Microsoft Azure Solutions • 70-537: Configuring and Operating a Hybrid Cloud with Microsoft Azure Stack • 70-473: Designing and Implementing Cloud Data Platform • 70-475: Designing and Implementing Big Data Analytics Solutions • 70-744: Securing Windows Server 2016 • 70-745: Implementing a Software-Defined Datacenter • 70-413: Designing and Implementing a Server Infrastructure • 70-414: Implementing an Advanced Server Infrastructure
MCSE: DATA MANAGEMENT AND ANALYTICS	<p>Demonstrate your broad skillset in SQL administration, building enterprise-scale data solutions and leveraging business intelligence data — both on-premises and in cloud environments.</p>	<p>Pre-requisites (one of the following):</p> <ul style="list-style-type: none"> • MCSA – SQL Server 2012/2014 • MCSA – SQL Server 2016 Database Administration • MCSA – SQL Server 2016 Database Development • MCSA – SQL Server 2016 Business Intelligence Development <p>Choose one of the following Azure exams:</p> <ul style="list-style-type: none"> • 70-473: Designing and Implementing Cloud Data Platform • 70-475: Designing and Implementing Big Data Analytics Solutions • 70-464: Developing Microsoft SQL Server Databases • 70-465: Designing Database Solutions for Microsoft SQL Server • 70-466: Implementing Data Models and Reports with Microsoft SQL Server • 70-467: Designing Business Intelligence Solutions with Microsoft SQL Server • 70-762: Developing SQL Databases • 70-767: Implementing a Data Warehouse using SQL • 70-768: Developing SQL Data Models • 70-773: Analyzing Big Data with Microsoft R • 70-774: Perform Cloud Data Science with Azure Machine Learning • 70-775: Perform Data Engineering on Microsoft HDInsight

Operationalize

Microsoft
Partner
Network

Artificial Intelligence

aka.ms/practiceplaybooks



Executive Summary



In the previous section, we reviewed how you should hire, train, and equip your staff. In this section, we will guide you through the steps to operationalize your business plan.

We walk you through the options for leveraging your internal use benefits that provide you complimentary software licenses and subscriptions for use within your organization, as well as you how can deepen relationships with your customer by re-selling it as an overall package along with your custom software, creating a new revenue stream for your business.

This section also provides guidance on how to operate your business, from how to build materials to support your sales and marketing efforts to the key contracts you will want to put in place.

Whether you're building products, providing managed services, or performing project work for customers, your success may be impacted by your ability to manage your customer records, your projects, and your support trouble tickets. We provide guidance on what tools and systems you should consider implementing.

We also cover how you can increase visibility for your practice by reviewing the Microsoft marketplaces and how to get listed on them as well as provide guidance on the social offerings your practice should setup.

We conclude this section with checklists and templates you can use to standardize your customer engagement process.

Top 5 things to do

Get your practice off ground by putting your plan into action. These are the top 5 things you should do to get the momentum going.

- ✓ Implement processes
- ✓ Claim your internal use benefits
- ✓ Set up key contracts and tools
- ✓ Set up customer support process
- ✓ Standardize your engagements using checklists

Implement a Process

The process you follow in delivering your AI solution to your customer is just as important as the technologies you use to deliver it. Fundamental to successful custom AI is data science and AI implementations are delivered following a data science process.

When implementing an AI solution, particularly one with custom AI that leverages custom machine learning or deep learning models, you should consider following a data science process. Data science itself requires a cross-disciplinary team for execution, and as such a well defined process that helps to orchestrate the collaboration is essential, and this holds true practices delivering AI solutions. Additionally, having a standardized, repeatable process will make it easier for team members to understand the work done by others, to add new members to the team, and enables better communication with all stakeholders.

From a high level, [the general data science process](#) can be summarized as consisting of the following phases, each of which may be revisited in iterations as needed:

- Ask an interesting question
- Get the data
- Explore the data
- Model the data
- Communicate and visualize the result

In the industry, partners are following [CRISP-DM](#), [KDD](#), the Microsoft [Team Data Science Process \(TDSP\)](#), or have created some variant as their own custom in house process.

When it comes to delivering solutions for an AI practice, consider leveraging the Microsoft [Team Data Science Process](#). Even if you use another data-science lifecycle, such as the CRISP-DM, KDD, or your organization's own custom process, you can still use the task-based TDSP.

Applying the Microsoft Team Data Science Process

The [Team Data Science Process \(TDSP\)](#) provides a systematic approach to building intelligent applications that enables teams of data scientists to collaborate effectively over the full lifecycle of activities needed to turn these applications into products. TDSP is an agile, iterative, data science methodology to improve collaboration and team learning. It includes a lifecycle definition, standard project structure, project planning and execution guidelines (including aspects of DevOps), and tools for productive data science.

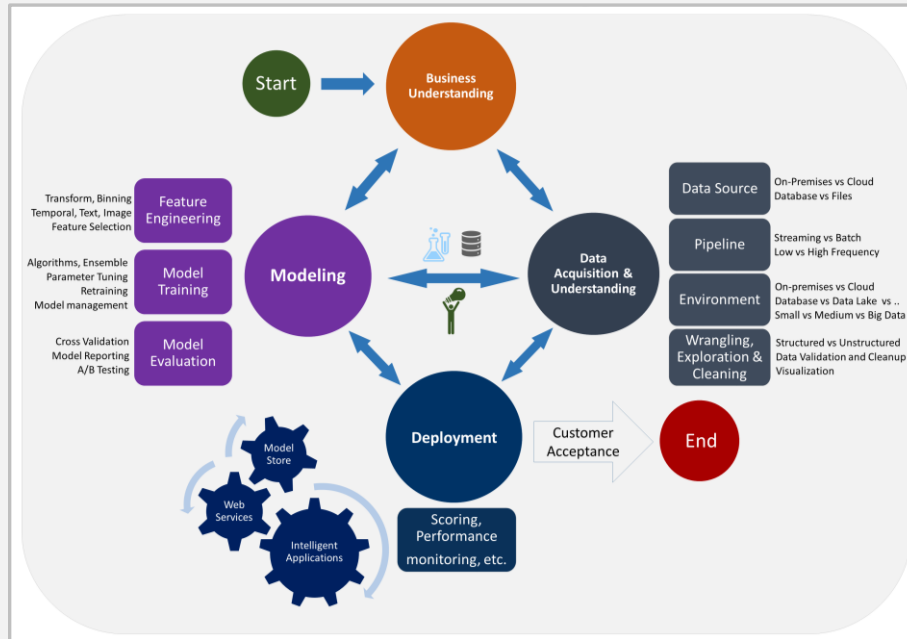
LIFECYCLE DEFINITION

The TDSP lifecycle is modeled as a sequence of iterative steps that provide guidance on the tasks needed to create, evaluate and use predictive models in production.

This lifecycle is designed for data-science projects that are intended to ship as part of intelligent applications. These applications deploy machine learning or artificial intelligence models for predictive analytics.

Exploratory data-science projects and ad hoc analytics projects can also benefit from the use of this process. But for those types of projects, some of the steps described by TDSP might not be needed.



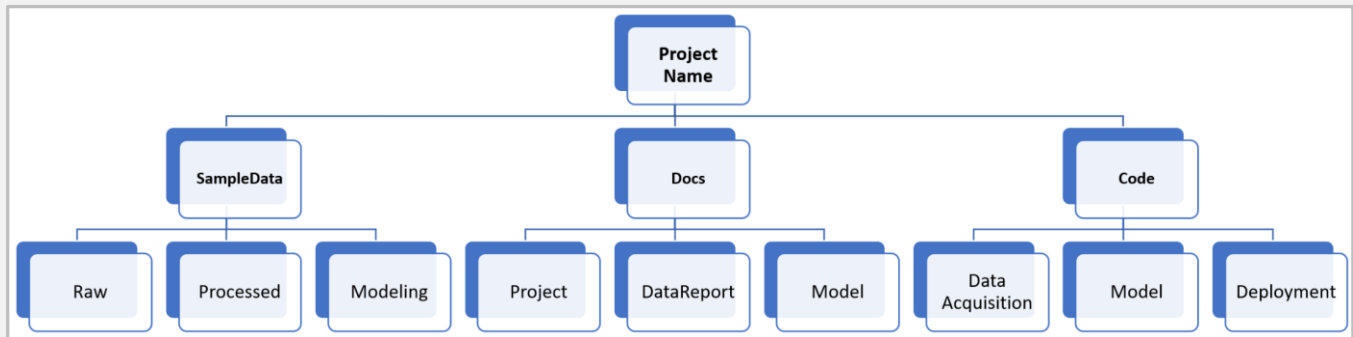


As the diagram illustrates, the data science lifecycle always begins with an understanding of the business – that is the problem domain, the requirements and the vision. From there the data science lifecycle begins with efforts around data acquisition and understanding (to get the data prepared for use in modeling). This is followed by the modeling effort, which examines the data to identify or engineer the key features that provide the greatest information or predictive strength, using this refined data to train the model and finally evaluate the performance of the model (e.g., in terms of measures like accuracy) and to validate that the predictions make sense in the context of the domain. Once ready, the trained model is deployed where it is ready for integration with intelligent applications. The final step is, naturally, customer acceptance of the delivered solution. It is important to note that the whole of the data science lifecycle is an iterative process, where each of the four major efforts may be revisited as the business understanding, data, and models are refined.

STANDARD PROJECT STRUCTURE

TDSP provides templates for the folder structure and required documents in standard locations. TDSP recommends creating a separate repository for each project in source control systems for versioning, information security, and collaboration.

As such it provides a standard project structure. This ultimately helps build institutional knowledge across the organization. All code and documents are stored in a version control system (VCS) like Git, GitHub, Microsoft Visual Studio Team Services, or Subversion to enable team collaboration. The project structure includes standardized folder structure, artifact templates you can use as a starting point, and tools for productive data science that integrate well with the process.



STANDARDIZED FOLDER STRUCTURE

The entire project structure can be cloned from the Azure TDSP ProjectTemplate repository on GitHub. At its root you provide a name for the project. Beneath the root is the SampleData folder that captures small sample data useful for basic execution of the developed solution and to provide example for the shape of the data. The Docs folder which contains architectures, data dictionaries, reports related to the data or the model, project management and planning docs, information obtained from stakeholders and documents that are used share information about the project. Beneath Code folder is stored all of the code of the project, organized in sub-folders according to the phase from the data science process (Data Acquisition & Understanding, Modeling and Deployment).

ARTIFACT TEMPLATES

The project repository includes key documents written in Markdown, such as the charter document which helps to define the project at the start of an engagement and the exit report which is used to provide a final report to the customer or client.

TOOLS FOR PRODUCTIVE DATA SCIENCE:


For performing data science steps, such as data exploration or modeling, and creating standardized reports for these stages of the projects, TDSP recommends creating team utilities or tools in Python or R, which are languages data scientists use most often. Such tools may be used by the team to improve efficiency, quality and standardization of data science workflow and reports.

As examples, TDSP provides two such tools, IDEAR (Interactive Data Exploration, Analysis and Reporting) and AMR (Automated Modeling and Reporting). IDEAR provides Jupyter notebooks (IDEAR in Python), and R scripts and markdowns (IDEAR in R Studio/Visual Studio or IDEAR in R Studio using Microsoft R Server) to standardize the initial data exploration, understanding and assessment process, along with reports you can export and include within the DataReport folder of the TDSP project structure. AMR (currently, available in R) provides an automated workflow for generating and comparing multiple modeling approaches on a given data-set. Currently, IDEAR and AMR are designed to work with structured (numerical and categorical) data only.

Data science teams may use these tools, or adopt or develop other such tools as needed.

STANDARDIZE INFRASTRUCTURE

TDSP encourages the standardization of infrastructure. Standardizing the infrastructure used during the data science process enables reproducible analysis. It also avoids duplication, which can lead to inconsistencies and unnecessary infrastructure costs. You should create and store in version control the tools that provision the shared resources, track them, and allow each team member to connect to those resources securely. It is a good practice to have project members create a consistent compute environment. Then, different team members can replicate and validate experiments easily and in a predictably consistent way.



APPLY TDSP WITH AZURE MACHINE LEARNING

There is a specific [GitHub repo that applies the TDSP](#) in the context of using Azure Machine Learning.

When you instantiate the TDSP from Azure Machine Learning, you get the TDSP-recommended standardized directory structure and document templates for project execution and delivery. Then you need to perform the following:

- Modify the documentation templates provided for your project
- Execute your project (fill in with your project's code, documents, and artifact outputs)
- Prepare the Data Science deliverables for your client or customer, including the ProjectReport.md report.

TDSP PLANNING AND PROJECT EXECUTION GUIDELINES

TDSP provides guidelines, tools and templates for project [planning](#) and [execution](#). For collaborative execution of data science projects, guidelines include using an Agile framework and Git version control. Aspects of DevOps can be integrated into such execution, as shown in this [training](#). This aspect of TDSP is being constantly expanded and refined. Further details and tools will be made available in the near future.

TDSP WORKED-OUT SAMPLES

To facilitate adoption of TDSP in data science projects using Azure Machine Learning and other Microsoft data platforms (SQL-server with Machine Learning Services, HDInsight , Azure Data Lake etc.), TDSP provides worked-out samples. These contain documentation and re-usable code. Users can look at these worked-out samples to see how finished projects could look like following TDSP lifecycle, structure and templates.

Claim Your Internal Use Benefits

A key benefit of being a Microsoft Partner is access to Internal Use Rights, providing your AI practice access to complimentary credits and licenses of Microsoft products and services, including Microsoft Azure, Office 365 and Visual Studio.

AZURE CREDITS

Providing access to Azure for your data science team is one of the key first steps to preparing for a successful AI practice. Microsoft provides several ways for your organization to gain access to Microsoft Azure for the development of new services, testing workloads, delivering services, or learning in general. For example, use your credits to enable your team to use the Data Science Virtual Machines with GPU's or perform scale out model training using HDInsight and Spark.

Members of the [Microsoft Action Pack](#) program receive \$100 monthly Azure credits.

Microsoft Partners with the Cloud Productivity Competency get Azure credits as a part of the Visual Studio subscription's core benefit (see the next section on Visual Studio).

Microsoft Partners with a Cloud Platform Competency at the Silver or Gold level get even more — \$6,000 per year and \$12,000 per year, respectively.

OFFICE 365 BENEFITS

As a Microsoft Partner, your core benefits include access to the Microsoft Office 365 Demo tenant that you can use to sell Microsoft Office 365, Power BI Pro, Microsoft Dynamics CRM Online and Project Online. You also get 25 seats of Office 365 E3 at the silver level or 100 seats of Office 365 E3 at the gold level from your core benefits.

As a Microsoft Partner with the Small & Midmarket Cloud Solutions competency, in addition to your core benefits, you get 10 seats of Office 365 E3 at the silver level and 25 seats of Office 365 at the gold level.

As a Microsoft Partner with the Cloud Productivity or Communications Competency, you get 25 seats to Office 365 E5 and 100 seats with the Gold competency.

VISUAL STUDIO

If your organization has Visual Studio subscriptions, you should know that each subscription has a set amount of Azure credits built in that the subscriber can use. The credit amount varies depending on the type of subscription purchased. You can also use MSDN software within your MSDN subscription on Azure Virtual Machines for development and test at no extra charge. The rate you will pay does not include any licensing costs — even virtual machines with SQL Server, SharePoint Server, or other software that is normally billed at a higher rate.

\$50 AZURE CREDIT	\$100 AZURE CREDIT	\$150 AZURE CREDIT
<ul style="list-style-type: none"> • Visual Studio Professional – annual • Visual Studio Professional with MSDN • Visual Studio Test Professional with MSDN 	<ul style="list-style-type: none"> • MSDN Platforms 	<ul style="list-style-type: none"> • Visual Studio Enterprise – annual • Visual Studio Enterprise with MSDN • Visual Studio Enterprise with MSDN (BizSpark) • Visual Studio Enterprise with MSDN (MPN)



Ways to Purchase Azure

There are a few ways you can purchase Azure and if you are a CSP, you have a built-in usage for testing your solution.

ENTERPRISE AGREEMENTS

Another option for getting access to your technical professionals is to purchase an Enterprise Agreement (EA). This arrangement is ideal for larger organizations that require the ability to create subscriptions for different departments, and even implement charge-back based on the department. Azure subscriptions within an EA agreement are managed through the Azure EA portal and allow for delegated administration and the ability to set quotas at the department or subscription level. For more information on how to get started with purchasing an enterprise agreement for Azure usage or adding Azure to an existing EA, visit: <https://azure.microsoft.com/en-us/pricing/enterprise-agreement/>.

PAY AS YOU GO AND TRIAL ACCOUNTS

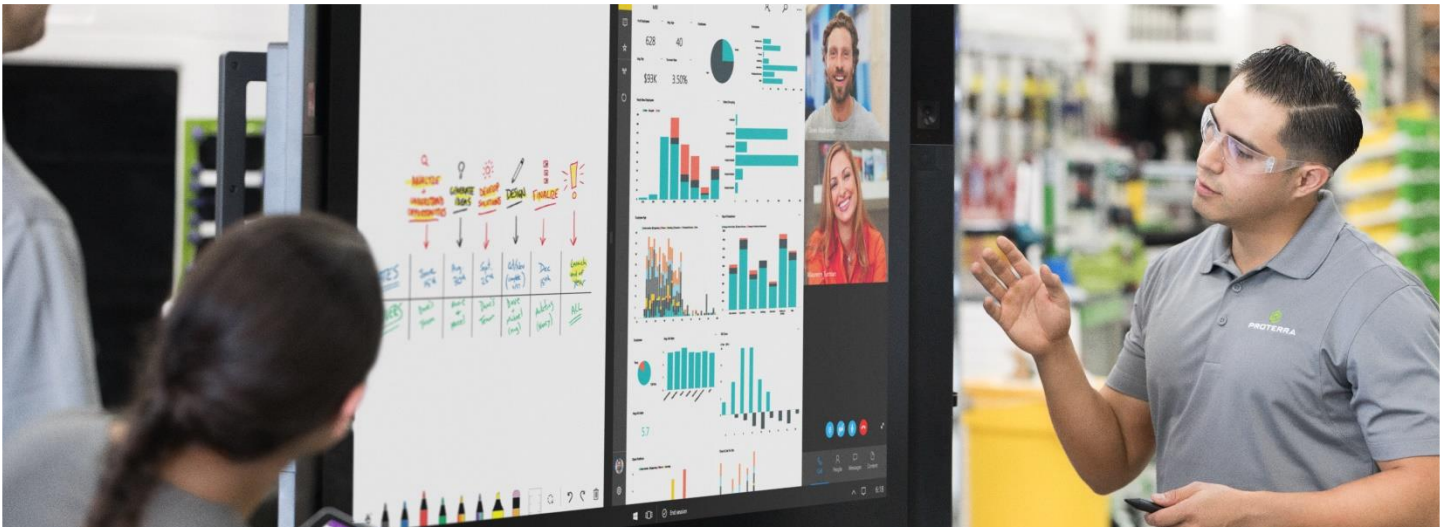
You can also create a free trial with Azure and allow it to convert to a pay-as-you-go subscription. An Azure free trial is valid for 30 days and allows up to \$200 in Azure credits. After the initial 30 days, any Azure usage is billed directly to you on your credit card. You can start a free trial by browsing <http://azure.microsoft.com> and clicking the free trial link.

OPEN LICENSE

You can also purchase Azure through a reseller using the Microsoft Open License Program. Open Value is the recommended program if you have a small to midsize organization with five or more desktop PCs and want to simplify license management, manage software costs, and get better control over your investment. It also includes Software Assurance, providing access to valuable benefits such as training, deployment planning, software upgrades, and product support to help you boost the productivity of your entire organization. For more information on the Microsoft Open Licensing program, visit: <https://www.microsoft.com/en-us/licensing/licensing-programs/open-license.aspx>.

CSP SANDBOX

Make sure you to take advantage of the CSP sandbox capability. Every Microsoft Partner onboarded in CSP has access to \$200 worth of test accounts for every subscription they provision.



Key Contracts and Tools for Your Practice

Practices need to use a set of legal documents to ensure compliance and deliverables, and an implementation process, to track the progress of a project both in terms of progress against a project plan and project budget, as well as protect your IP.

KEY CONTRACTS

Leverage the [Key Contracts for Your Practice guide](#), to learn more about developing service level agreements, master services agreements, a statement of work, and a mutual non-disclosure agreement.

MICROSOFT PROJECT ONLINE

[Microsoft Project Online](#) is a flexible online solution for project portfolio management (PPM) and everyday work. Delivered through Office 365, Project Online provides powerful project management capabilities for planning, prioritizing, and managing projects and project portfolio investments — from almost anywhere on almost any device. Project Online can be used by administrators, portfolio managers and viewers, project and resource managers, and team leads and members.

VISUAL STUDIO TEAM SERVICES

[Visual Studio Team Services](#) provides various tools for tasks like running agile teams, providing support for Kanban boards, handling work item backlogs, scrum boards, source control, continuous integration and release management. Source control functionality provides Git support, which enables integration with GitHub if such integration is desired.

While Visual Studio Team Services will help you manage the technical aspects of your project, cost-containment requires a different set of tools.

MICROSOFT DYNAMICS 365 FOR PROJECT SERVICE AUTOMATION

[Microsoft Dynamics 365 for Project Service Automation](#) provides users with the capabilities required for setting up a project organization, engaging with customers, project scheduling and costing, managing and approving time and expenses, and closing projects. It is specially targeted to address the needs of a Project Services based practices, as it is designed for professionals who manage projects and the associated customer engagement process end to end.

GITHUB

[GitHub](#) provides the hosted environment for the AI implementation team to version control and share their source code, notebooks and other artifacts both privately (e.g., internally to a team) and publicly (e.g., an open source project) and collaborate on development projects.



Collaboration Tools and File Sharing

Collaborating with customers through the lifecycle of a project or the duration of a managed services agreement is critical. There are several services that can help you share project plans or set up lists for shared data.

JUPYTER NOTEBOOKS

Jupyter Notebooks were introduced in the earlier section of the Team Data Science Process. These online, web-based notebooks enable the AI implementation team to share code and collaborate on data wrangling, data understanding, data preparation, model training and model evaluation. They also provide convenient mechanism to share results (in the form of notebooks that include rich text, tabular data and charts) with customers and other stakeholders in a read-only fashion.

MICROSOFT TEAMS

[Microsoft Teams](#) is the latest collaboration tool from Microsoft and is designed to make your content, tools, people, and conversations available in a single location.

YAMMER

[Yammer](#) is an enterprise social network collaboration offering to help teams collaborate and share files with each other.

ONEDRIVE FOR BUSINESS

[OneDrive for Business](#) is an enterprise file sharing service that is designed for automatic synchronization of files between your computer and the cloud. OneDrive makes it easy to share files with your customers or partners.

SKYPE FOR BUSINESS

[Skype for Business](#) is an enterprise online meeting and conference service designed for business communications.

SURFACE HUB

[Microsoft Surface Hub](#) is a Skype Online-integrated collaborations device, or “meeting room in a box.” In addition to the built-in team experiences like Skype for Business, Microsoft Office, and Whiteboard, Microsoft Surface Hub is customizable with a wide array of applications. Universal apps built for Windows 10 shine on Microsoft Surface Hub and scale to the large screen. You can also connect apps from your personal device and drive them from Microsoft Surface Hub.

Using CRM to grow your business

CRM solutions streamline processes and increase profitability in your sales, marketing, and service divisions.

A strong CRM solution is a multifaceted platform where everything crucial to developing, improving, and retaining your customer relationships is stored. Without the support of an integrated CRM solution, you may miss growth opportunities and lose revenue because you're not maximizing your business relationships. Imagine misplacing customer contact information, only to learn that your delay pushed your client into the arms of a competitor. Or, picture your top two salespeople pursuing the same prospect, resulting in an annoyed potential customer and some unfriendly, in-house competition. Without a centralized program where your people can log and track customer interactions, your business falls behind schedule and out of touch.

THE FUNDAMENTALS OF CUSTOMER RELATIONSHIP MANAGEMENT

CRM tools make the customer-facing functions of business easier. They help you:

- Centralize customer information
- Automate marketing interactions
- Provide business intelligence
- Facilitate communications
- Track sales opportunities
- Analyze data
- Enable responsive customer service

Running a successful business is no simple task. When marketing campaigns, data analysis, meetings, customer care, and more, all happen simultaneously, you need a powerful CRM solution to bring all these functions together in one place.

As a sales professional, you'll be working with the following types of records:

ACCOUNTS: Account records contain information about the companies you do business with.

CONTACTS: Contact records contain information about the people you know and work with. Usually, multiple contacts are associated with one account. Contacts could include people responsible for making purchasing decisions or paying invoices, support technicians, or anyone you work with at the company.

LEADS: Leads are potential sales, and you or your company can get leads from many different sources. For example, you can generate sales leads from marketing campaigns, inquiries from your website, mailing lists, social media posts, or in person at a trade convention.

OPPORTUNITIES: When you qualify a lead, it becomes an opportunity, or a deal that you're getting ready to close.

[Microsoft Dynamics 365](#) can be customized, so you can also work with records relevant to your field and the way your organization does business, including [sales](#), [customer service](#), [field service](#), [project service automation](#), and [marketing](#).



Define Customer Support Program and Process



Support overview

It has been said that an unhappy customer represents an opportunity to make a customer for life. Studies have found when a customer gets to the point of a complaint, they are very emotionally engaged. If you can turn that negative around to a positive, you may just have a customer for life.

When it comes to support, there are two perspectives you should consider. First, how will you support your customers when they have engaged you for project services, managed services, or are utilizing your intellectual property. Second, where do you go for [Azure support from Microsoft](#) for a solution you are building, or because you need assistance on behalf of your customer?

THE ITEMS YOU WILL NEED TO WORK THROUGH INCLUDE:

- Defining your support model
- Provisioning your support infrastructure
- Defining and implementing your escalation process
- Selecting and enabling your support options for Azure

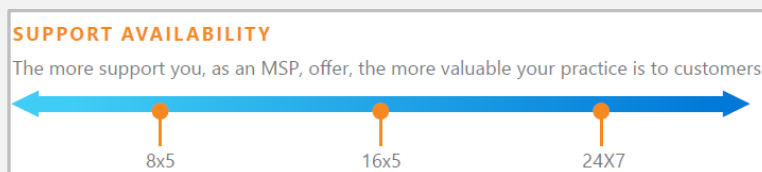
We cover each of these topics in the pages that follow.

Supporting Your Customers

Let's begin with the first scenario in which you support your customers directly. It should go without saying that one of the most important functions for your MSP practice will be supporting your customer once their applications and data are firmly in the cloud or sitting in a hybrid deployment. No matter how well a cloud or hybrid environment is planned, provisioned, operated, or monitored, problems will arise — and those problems will need to be remediated. It's your job as an MSP to offer support to your customers to deal with outages, breaches, inefficiencies, and disaster scenarios. MSPs need to consider the level of support that makes sense for their practice — in terms of resources and revenue — as well as what makes sense to the customers they serve.

SUPPORT MODEL

How do you package and sell your support? The typical options are to provide support either on a retainer basis (where the customer pays a monthly fee for up to a certain number of "use it or lose it" support hours) or per incident (where the customer pays a fee every time they utilize your support). You must also define your support availability so your customers have a realistic expectation of when they can access your service.



ESCALATION PROCESS

How does a customer get help at the right technical level? For your support process to make economic sense, avoid having your most skilled and most expensive resources (e.g., architects, senior developers, data scientists, etc.) answer every support call. For your particular solution offering, consider implementing a tiered support offering of junior-level resources that are equipped to handle common issues. These resources should be equipped to escalate a customer support case to a more senior-level resource once the common issues have been ruled out. You will need to decide how many levels of tiered support to offer, but two to three tiers are most common. When defining your escalation process, do not forget about the basics. For example, how do customers get in touch with you for support in the first place? This could be a dedicated support telephone number, forum or chat room, Twitter handle, email address, etc.



Support infrastructure: How will you manage customer support requests and track them to closure? Many MSPs offer premium support offerings such as a Technical Account Manager who is responsible for tracking, reporting, and escalating an issue.

Support Options from Microsoft

How do you receive support for your implementation efforts or on behalf of your customer?

SIGNATURE CLOUD SUPPORT

[Microsoft Signature Cloud Support](#) is provided as benefit to Silver and Gold Partners. It primarily provides support for issues occurring in Azure subscriptions you own or on which you are a co-admin. It is not intended for use in supporting issues in subscriptions owned by your customers.

MICROSOFT ADVANCED SUPPORT FOR PARTNERS

[Microsoft Advanced Support for Partners](#) is the ideal solution for partners who are growing their cloud business. Not quite ready for Premier Support, but need a higher level of service than the Microsoft Partner Network core benefits provide? The Advanced Support program delivers the right level of support to meet you in the middle while your business grows. With Advanced Support for Partners, you get cloud support at an accessible price point, which helps you be a great ally to your customers and grow your business faster. The program includes valuable proactive and reactive services delivered by experienced Services Account Managers and Partner Technical Consultants. Advanced Support for partners enables you to provide support on behalf of your end customers, in addition to providing support on subscriptions you own directly. Designed from the feedback of over 1,500 partners like you, Microsoft Advanced Support for Partners addresses the specific needs of Cloud Solutions Providers (CSPs), born-in-the-cloud partners, and all other partners selling Microsoft cloud services.

MICROSOFT PREMIER SUPPORT FOR PARTNERS

[Microsoft Premier Support for Partners](#) delivers a managed support offering for you and your customers — proactive support services for developing, deploying, and supporting Microsoft technology whether on-premises, hybrid, or in the cloud. As the only partner program with complete, end-to-end managed support across the full Microsoft platform, Premier Support for Partners also provides a powerful marketing tool to gain competitive advantage in the marketplace.

Microsoft offers a range of paid [Azure support plan](#) options for customers from developers starting their journey in the cloud to enterprises deploying business-critical, strategic applications on Microsoft Azure. These options are available in tiers — **Premier, Professional Direct, Standard and Developer Support Plans** — that are available for purchase directly for those who are not Microsoft Partners. In addition to these paid plans, Azure offers **core support**, which is free and provides support via forums, and help with account billing or management questions.

SUBMITTING AZURE SUPPORT REQUESTS

Support requests need to be submitted using the Azure Portal. First you must log in to the subscription for which you want to receive support. Next, submit a support request. Once submitted, you can manage the incident from the Azure Portal.

PARTNER ADVISORY HOURS

Partner advisory hours are used as currency for technical presales and advisory services offered by the Microsoft Partner Services team.

As part of your company's Microsoft Partner Network membership, your organization receives partner advisory hours for attaining a Microsoft competency, membership in Microsoft Cloud Accelerate, and subscribing to Microsoft Action Pack Develop and Design.

PARTNER LEVEL	NETWORK MEMBER	ACTION PACK	SILVER	GOLD	THESE HOURS CAN BE USED TO:
	0 hours	5 hours (after first cloud sale)	20 hours	50 hours	

SUPPORT OPTIONS

PARTNER-FACING OPTIONS	RESPONSE TIME
Signature Cloud Support	Less than 2 hours
Microsoft Advanced Support for Partners	Less than 1 hour
Microsoft Premier Support for Partners	Less than 1 hour
Partner Advisory Hours	N/A

CUSTOMER-FACING OPTIONS	RESPONSE TIME
Premier	Less than 1 hour
Professional-Direct	Less than 1 hour
Standard	Less than 2 hours
Developer	Less than 8 hours
Core	N/A

Manage and Support an AI solution deployed in Azure


Support Resources

Supporting an Azure deployment involves transitioning from deployment focus to ongoing health and occasional troubleshooting. Microsoft Azure offers several services to help manage and monitor workloads running in Azure, documentation for troubleshooting the services for your practice, and the [Azure Community](#) where you can ask questions, get answers, and connect with Microsoft engineers and Azure community experts.

KEEPING AN EYE ON COSTS

If your AI solution is deployed to Azure, there are a lot of ways you can both forecast spend and keep track of your actual costs.

- You can get estimated costs before adding Azure services by using the [Azure Pricing Calculator](#). This calculator includes all Azure services, including those relevant to the AI practice like Cognitive Services, Azure Machine Learning and HDInsight.
- Once you have resources deployed to Azure, you should regularly check the Subscription blade in the [Azure Portal](#) for cost breakdown and burn rate. From here, you use the Cost analysis feature to analyze the cost breakdown by resource.
- You can also report on your Azure costs programmatically by using the [Azure Billing APIs](#). There are two APIs available that when used together enable you estimate your spend by resource: the Azure Resource Usage API enables you to get your Azure consumption data, and with the Azure Resource RateCard API you can get the pricing information for each Azure resource.
- There are other situations, such as for an EA, a sponsored Azure subscription, or subscriptions acquired thru a CSP provider. These have their own portals for analyzing consumption and costs. For details on how to monitor these, view [this Azure billing and cost management article](#).

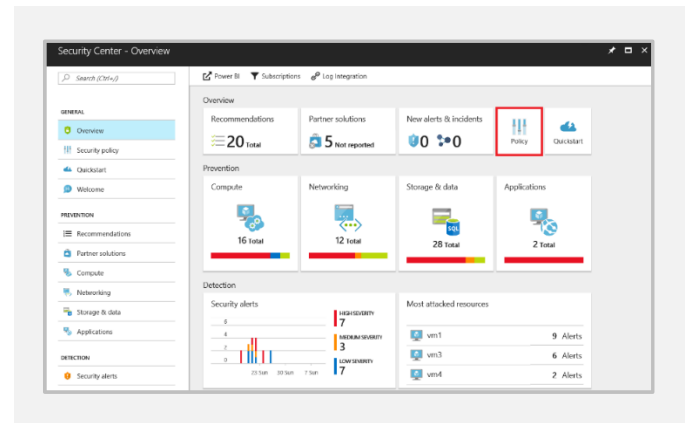


AZURE SECURITY CENTER

[Azure Security Center](#) helps you prevent, detect, and respond to threats with increased visibility into and control over the security of your Azure resources. It provides integrated security monitoring and policy management across your Azure subscriptions, helps detect threats that might otherwise go unnoticed, and works with a broad ecosystem of security solutions. Azure Security Center should be part of any managed service practice to assist with monitoring and support.

Security Center delivers easy-to-use and effective threat prevention, detection, and response capabilities that are built in to Azure. Some of its key capabilities are:

- Monitor the security state of your Azure resources
- Defines policies for your Azure subscriptions and resource groups based on your company's security requirements, the types of applications that you use, and the sensitivity of your data
- Uses policy-driven security recommendations to guide service owners through the process of implementing needed controls
- Rapidly deploy security services and appliances from Microsoft and partners
- Automatically collect and analyze security data from your Azure resources, the network, and partner solutions like antimalware programs and firewalls
- Leverages global threat intelligence from Microsoft products and services, the Microsoft Digital Crimes Unit (DCU), the Microsoft Security Response Center (MSRC), and external feeds
- Apply advanced analytics, including machine learning and behavioral analysis
- Provides prioritized security incidents/alerts
- Offers insights into the source of the attack and impacted resources
- Suggests ways to stop the current attack and help prevent future attacks



AZURE ADVISOR

[Azure Advisor](#) analyzes your resource configuration and usage telemetry to detect risks and potential issues. It then draws on Azure best practices to recommend solutions that will reduce your cost and improve the security, performance, and reliability of your applications.

OMS LOG ANALYTICS

Even if you are not offering OMS as part of your core offering, using [Log Analytics](#) for support and monitoring can be a huge time saver. Log Analytics can help you collect and analyze data generated by resources in your cloud and on-premises environments. It gives you real-time insights using integrated search and custom dashboards to readily analyze millions of records across all your workloads and servers regardless of their physical location.

ENGAGING MICROSOFT SUPPORT

If you are a CSP or have sold support as part of your managed services solution you are the front-line support for your customer. At some point, you may need to contact Microsoft to escalate an issue. Microsoft offers several options via forum support or via paid options as discussed in the preceding [Support Options from Microsoft](#) section.

Support Ticket Setup and Tracking

Customer Support

Setting up tickets, tracking issue resolution, and managing customer success are fundamentals of your practice.

Providing support to your customers from your practice is a non-trivial, omni-channel effort. Consider using Azure Machine Learning to monitor the performance of production deployed models. We suggest you implement [Microsoft Dynamics 365 for Customer Service](#) to help you quickly set up and start managing your overall customer support efforts.

MONITORING DEPLOYED MODELS WITH AZURE MACHINE LEARNING

You can use the model data collection feature in Azure Machine Learning to archive model inputs and predictions from a web service. For deeper insights into performance, you can capture model telemetry using the Azure Machine Learning SDK. The model telemetry can be used later for analyzing model performance, retraining, and gaining insights for your business.



MICROSOFT DYNAMICS 365 FOR CUSTOMER SERVICE

Microsoft Dynamics 365 for Customer Service is designed to manage the efforts of your customer support teams. It provides licensed users with access to core customer service capabilities for a significantly lower price than comparable offerings from other vendors, including enterprise case management, Interactive Service Hub, Unified Service Desk, SLAs and Entitlements, and other service group management functionality.

CREATE CONSISTENCY AND LOYALTY

Provide the seamless service your customers expect by meeting them where they are with the information they need, every time.

- Give customers great service on their channel of choice.
- Make help easy by providing relevant, personalized service.
- Proactively address issues by detecting customers' intent and social sentiment.

MAKE YOUR AGENTS' JOBS EASIER

Give your agents complete information — in a single customer service software app — to make smart decisions and provide great service.

- Reveal customers' case histories, preferences, and feedback.
- Provide guidance on entitlements and service-level agreements.
- Display it all in a single interface tailored to their job and skillset.

GET AN ADAPTIVE ENGINE

Respond quickly to customer and market changes within an agile, cloud-based environment that has digital intelligence built in.

- Adapt and customize easily using configuration, not code.
- Extend your functionality through a single interface.
- Rely on advanced analytics and a trusted cloud platform.

Implement Intellectual Property Offerings

Implementing IP in Your AI Practice Offerings

Consider these tips to start productizing your IP and go to market.

DEFINE YOUR SOLUTION

When we ask partners how they determined what IP they were going to build, we often get the same answer, which is that they realized most of their customers were asking for the same thing or something very similar. And rather than continuing to do high-cost custom work for every customer, they decided to productize what their customers were asking for. Bring your sales, marketing, technical, and delivery teams together to brainstorm and define what your solution will look like.



DETERMINE WHAT WILL DIFFERENTIATE YOUR SOLUTION FROM OTHERS IN THE MARKET

It is important that you think about your differentiation strategy. What is going to make your solution better than other similar solutions in the industry?

MAINTAIN RIGHTS TO THE IP

As partners make the transition from project or custom services to packaged IP, it is critical they revise their customer agreements so the partner can maintain the IP rights to the solutions.

PROTECT YOUR IP

As we mentioned in Understanding Intellectual Property, you should engage legal counsel to help you protect and maintain ownership of the IP you create. Key to partner success with IP is taking care with licenses, contracts and terms of use and the acquisition of patents.

ESTABLISH A RECURRING REVENUE MODEL

The beauty of deploying IP in the cloud space is that you can light up the recurring revenue model, which will have a positive impact on the valuation of your business and even help your cash flows in the future.

CONSIDER YOUR CHANNEL STRATEGY

One of the advantages of productizing your IP is that it opens up a lot of doors to sell your solution through channel partners.

Resources

- [Building IP to Drive Margins](#)
- [Create Stickiness with IP](#)

Setup Social Offerings

Blogging, Meetups, and More

Contributing to the technical community can help you increase credibility for your practice. It has the side benefit of strengthening the technical acumen of your delivery team by having them focus on a specific subject for a public-facing deliverable. Below are some suggested options to get started.

BLOGGING

Technical blogging is a great way to increase the skills of your technical team, as well as grow stature in the community at large with your organization. Blog posts should be well thought out and simple to digest. Visual aids such as diagrams or nicely formatted source code snippets go a long way towards readability.

MEETUPS, USER GROUPS & ASSOCIATIONS

Speaking at user groups and association events is another valuable tool to increase the skills of your team. Similar to blog posts, its great practice for honing vital communication skills with your team, as well as a great opportunity to dig deeper into a specific subject related to your practice. For an AI practice, consider the following:

- [Meetup.com Artificial Intelligence Meetups](#)
- [Association for the Advancement of Artificial Intelligence](#)
- [AI International – Societies around the world](#)

GLOBAL AZURE BOOTCAMP

Each user group will organize their own one-day deep dive class on Azure the way they see fit. The result is that thousands of people get to learn about Azure and join online under the social hashtag #GlobalAzure! This is a great opportunity to attend, participate as a speaker (reach out to your local organizer to see how you can help) or host your own. For more information, visit <http://global.azurebootcamp.net/>.

aka.ms/practiceplaybooks

WEBINARS

Webinars are another resource to extend your teams skills. Similar in scope to speaking at a meetup or user group, the webinar allows a much broader reach as attendees from all over the globe can attend.

MICROSOFT MVP COMMUNITY

For more than two decades, the Microsoft MVP Award is our way of saying thank you to outstanding community leaders. The contributions MVPs make to the community, ranging from speaking engagements and social media posts to writing books and helping others in online communities, have incredible impact. Among other benefits, MVPs get early access to Microsoft products and direct communication channels with product teams, and are invited to the Global MVP Summit, an exclusive annual event hosted in Microsoft's global HQ in Redmond. They also have a very close relationship with the local Microsoft teams in their area, who are there to support and empower MVPs to address needs and opportunities in the local ecosystem.

Contributing to the Azure community not only helps the reputation of your practice, but it can also hone much-needed skills for your delivery team.



Create Engagement Checklists & Templates

Standardize Customer Engagement

Repeatable processes make for profitable practices. Use the following example to kickstart your own checklist for executing a new engagement.

- Conduct envisioning session to capture vision for AI solution.
- Conduct initial requirements meeting.
- Identify product owner/manager(s).
- Conduct follow-up meeting to clarify and establish next steps.
- Discuss minimal viable product (MVP) criteria and decide on proceeding with a PoC or Pilot.
- Conduct an architecture design session to capture solution detailed solution architecture.
- Establish data science process (Microsoft Team Data Science Process, CRISP-DM, etc.).
- Identify milestones, tasks, evaluation criteria and exit criteria; share with customer.
- Provide cost estimates for research, data collection and preparation, modeling, development, cloud services, and ongoing maintenance/support.
- Address customer objections to proposed technology and services.
- Establish project repository for code, models, and project issues, tasks and documentation artifacts (e.g., GitHub or Visual Studio Team Services).
- Acquire data (or sample of data) for initial data assessment, prepare data for use in model training.
- Perform model development or integrate pre-built AI.
- Evaluate model or pre-built AI performance.
- Deploy solution.
- Follow up with customer and provide reports/status/demos on a regular basis (e.g., two-week sprint).
- Conduct a final handoff to customer.
- Conduct project debrief with customer.
- Customer conducts acceptance tests.
- Conduct internal project post-mortem.
- Monitor performance of model or AI in production.

Microsoft
Partner
Network

Go to Market & Close Deals

Artificial Intelligence

aka.ms/practiceplaybooks



Executive Summary

In previous sections in the playbook, we covered topics from how to build your practice by selecting products or services to specialize in, to building and training your team help turn your ideas into reality, to bringing your special offering to market and finding and keeping great customers. So, what's left to do? In this section, we'll discover strategies to compel potential customers that may be sitting on the fence to take action, from creating a good value proposition to building marketing and sales materials that tell your story.

It has been said that your current customers are your best customers. Do you know who your best customers are? What do they have in common? And how do you find more like them? We'll start by helping you build foundational marketing materials such as marketing personas, points of differentiation, value propositions, and customer business needs.

Once you've built the foundation, we'll look at how you can put these materials to work. We'll go through the different ways you can attract new customers and look at best practices. How do you put it all together? We'll discuss why integrated marketing campaigns work the best, and the tools you need to run them, such as a CRM system and marketing automation.

But marketing is only half of the story. Your sales team is the other half. Don't forget how the two work together and what marketing can do to support sales. The job of the marketing team is to build out not only customer facing materials, but also compelling materials that can be used to train and arm your sales team.

The sales end of the bargain is to close the sale. One way to do this is by writing a winning proposal. Another way is to build a proof of concept or prototype of your product or service offering, which could help a prospect understand what it is you're offering, or solidify their vision of what you can help make possible. Microsoft is committed to helping your business grow, and provides both co-selling and co-marketing opportunities.

Finally, don't miss the Microsoft resources available in the [Go-to-Market and Close Deals guide](#), which you can leverage to help build your marketing materials and campaigns, as well as resources to help your team close the deal.

Top 5 things to do

Add value to your practice and turn your prospective customers into lasting ones. These are the top 5 things you should do to go to market and get deals done.

- ✓ Identify your customer's business needs
- ✓ Write a compelling value proposition
- ✓ Leverage marketing to find customers
- ✓ Build marketing and sales materials
- ✓ Collaborate with partners

Marketing to the AI Buyer

Plan your customer’s journey to buying.

What is different about the AI buyer? According to the partners interviewed for this book, rarely is the buyer actively looking for AI. They are looking for partners to help them solve a domain specific problem. It is up to the partner to recognize the opportunity to apply and sell AI.

DO’s and DON’Ts for Marketing to the AI Buyer

DO target your existing customers with envisioning sessions and PoCs before marketing to win new customers

DO emphasize how the solution augments human ingenuity

DO help them envision the possibilities enabled with the use of intelligent technology

DO describe the benefits of in terms of the business needs (e.g., “Our technology helps you answer the questions: Who might drop-out of school? Why would

they drop-out of schools? How and when can we intervene?”)

DO explain how AI might benefit the customer’s digital transformation (see illustration that follows)

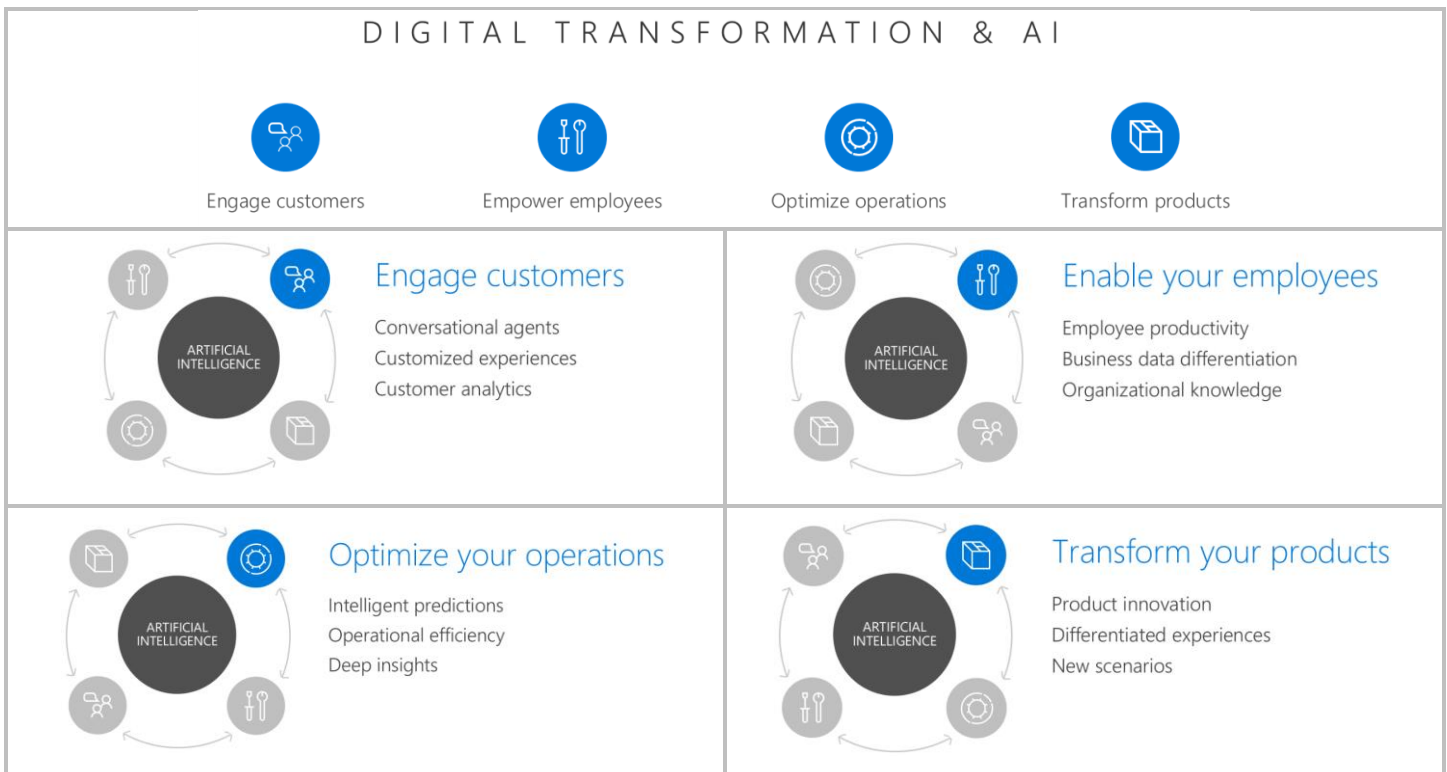
DO provide realistic benefits based on your experience with your solution

DON’T lead with AI as the value proposition

DON’T expect customers to ask for an AI solution

DON’T describe the benefits solely in terms of the “cool” technology (e.g., “Win with predictive analytics, deep learning, analytics and chat bots”)

DON’T overpromise the capabilities of AI





Buyer engagement

Buyers buy differently than in the past. With all the information on the internet, buyers tend to research and self-educate long before they engage with sales people. By the time they do engage with sales, they've already made some decisions.

To help illustrate this, just think about the way a buyer might go about buying a new car. Before going to the car dealership, the buyer will likely read about various car models on the internet, read reviews, and make some decisions. When the buyer is ready to visit a dealership, they already know what they want and how much they are willing to pay for that car. This poses a challenge for sellers. How can you get prospects to engage with you earlier in the process? Through marketing.

Another way partner businesses are changing is that when selling cloud-based solutions, you can gain recurring revenue streams. Recurring revenues provide business stability and confidence for business owners and managers, allowing them to make business decisions that may not be as easy when revenues are irregular and lumpy. While these recurring revenues are smaller on a per-transaction basis than buyers' large capital expenditures, you adjust for this. You'll need a higher volume of transactions.

To support that, you'll need a higher volume of high quality sales leads (through modern marketing techniques) coming in to create larger sales pipelines. Clients who are buying on a recurring basis represent great opportunities for you to upsell and cross-sell additional products and services.

Marketing is not an option anymore. Marketing helps you educate, identify, and engage with prospects earlier in the sales process. By identifying prospects who indicate interest in your products and services via their behavior (website visits, clicks, downloads, etc.), marketing can deliver high-quality leads.

Inbound marketing techniques such as search engine optimization and pay-per-click advertising make it easy for prospects to find you. Outbound marketing techniques, such as e-mail and telemarketing, enable you to tell prospects about your company's solutions.

Marketing is the toolset that addresses all these changes. Marketing today is digital and has the power to reach more people. Again, it's not to say that more traditional, non-digital marketing is ineffective. But to be found by prospective buyers that you don't have a relationship with, you need to employ digital marketing techniques. Modern marketing is focused on the prospects' and clients' views of the world.

Consultative Selling and Technical Pre-Sales

Discovering the art of the possible

From the very start of your engagement with a prospect, you need to be aware of the need for technical pre-sales assistance. Many times, you are dealing with business decision makers during the buying cycle. In that case, you are less likely to have a need for technical assistance. However more than ever before, technical staff are a part of decision making with AI practices as they help envision a solution to solve a customer need.

Your technical pre-sales staff should be very experienced users of your products and services. These employees need training or experience as a user of your products. Former support employees often make good technical pre-sales staff. The technical pre-sales staff is in place to explain technology, how it works, how it meets a business need and to answer any other questions. They should excel at the more complex issues that come from prospects, and be focused on pre-sales, working together with sales and marketing, who address the business benefits. One without the other cannot be effective. You need the sales staff to speak to business decision makers and envision the art of the possible, with AI solutions this often occurs jointly with technical expertise.

Examples of technical probing questions to ask during pre-sales conversations supporting an AI practice:

- What are the challenges you are looking to solve?
- Are you looking to improve communication, learn from your data (such as predicting future events)?
- Do you have the data to help you approach these challenges? In what formats?
- Is the data generated and captured with your system or is it external and provided by 3rd parties?
- What application development and technologies are within your existing team's comfort zone? Do you have any data scientists on the team?
- What application platforms would you like to target? Web, mobile, desktop, IoT, etc.
- Do you have any compliance or regulatory requirements that pertain to the handling of your data?

- Can you walk us thru the high level of where data enters your system and how it is ultimately consumed?

BEST PRACTICES – CONSULTATIVE SELLING:

Rather than just promoting an existing product, the salesperson focuses on the customer's problems and addresses the issue with appropriate offerings (products and services). The problem resolution is what constitutes a "solution".

The best reps combine solution selling with insights.

To gain credibility in the eyes of the buyer, the solutions sales rep must introduce content and data that adds value to the sales call.

Ask good questions. The successful solutions seller remains sensitive to the buyer's needs and asks important questions at the right moment.

Listen actively. Solution selling requires considerable understanding of the buyer's needs, which will only come from listening attentively. Solution sellers should actively listen as the buyer details their organizational needs, taking notes and asking considerate questions in the process.

Offer guidance. Solution sellers must guide the buyer towards the solution being offered. This guidance comes as the solution seller adopts something of a teaching role, helping the buyer to overcome business challenges by utilizing their deep knowledge of industry pain points and trends.

RESOURCES

- [Azure Briefing and Pitch Decks](#)
- [Azure Pre-Sales Resources](#)
- [Azure Training Resources](#)



Microsoft Technology Centers

The [Microsoft Technology Center \(MTC\)](#) can help you find the right solutions to transform your customer's business in a mobile-first, cloud-first world.

With over 40 locations around the globe, the MTCs bring together the right resources to help you accelerate your customer's digital transformation.

- **People:** The MTC staff is comprised of experts in Microsoft solutions. Their tenure in the industry ensures they will effectively guide your team to rapidly find solutions to your technology challenges.
- **Partners:** The MTCs have formed alliances with industry leaders who provide comprehensive resources, including hardware, software, and services to explore during your engagements.
- **Place:** The MTC environment provides rich interactive and immersive experiences for you to learn first-hand how Microsoft and partner technologies can help you take on your most difficult challenges.

Start your digital transformation by experiencing the Microsoft cloud at one of the 40+ centers around the globe.

The MTC can work with you to help you close sales with these engagement offerings:

- **Strategy Briefing:** This one-day briefing starts by examining your current IT environment and business objectives. Then it moves into the Envisioning Center, where you'll see Microsoft solutions in action through powerful demos and scenarios customized to meet your needs. The day includes mutual discovery, tailored product and technology drill-downs, and expert presentations. It culminates with the delivery of a clear and actionable picture of how Microsoft and partner technologies can help you reach your business goals.
- **Architecture Design Session:** This custom session focuses on your business objectives and aligns them with specific applications of Microsoft software to help you not only meet your goals, but also capitalize on them. We'll provide architectural guidance, consultation on preferred practices, and risk analysis to chief technology officers, architects, and senior members of your development team.
- **Proof-of-Concept:** In this multi week, in-depth workshop, our architects work closely with key members of your technical staff to transfer knowledge and prove out customized solutions. This workshop may also include detailed demos and training sessions. Your team will have a private, secure, and fully loaded development suite that's preconfigured prior to their arrival.
- **Workshops:** If seeing is believing, then imagine what a hands-on immersive experience can do! Attend a custom briefing that includes a facilitated, hands-on environment where you and your colleagues can experience the vision of Microsoft's platform and solutions firsthand.

Architecture Design Session (ADS)

An architecture design session is a working session between your experts and the customer. It should follow the envisioning session and build on the customer's vision already established.

This intensive, two-day session delivers in-depth technical information on integrating data from across your customer's entire organization and delivering it in an analysis-ready form. Presentations, demonstrations, and whiteboard discussions are customized to address your customer's needs. In many cases, the design session is used to identify candidate proofs of concept. The primary audience includes architects, developers & data scientists, test and quality assurance (QA) engineers, and technical staff, and here are some potential topics to cover:

- **Document solution architecture:** Document how all the services and components fit together to deliver the complete AI solution, end to end.
- **Identify pre-built AI & custom AI:** Identify what components of the AI solution will integrate pre-built AI, and which will require custom modeling efforts to produce custom AI.
- **Catalog data assets:** Collect the catalog of data assets that will be needed for any custom AI efforts. Be sure to document the pipeline with which this data is collected, the anticipated quality of the data, and any other factors that can complicate preparing the data for modeling.
- **Hone in on missing data:** Attempt to identify situations where the vision requires data that the customer does not currently have. Identify how this missing data will either be procured, created or its absence worked around.
- **Discuss risk:** During the ADS, it may become apparent that parts of the AI solution have a high risk of failure or the client is skeptical about achieving the vision. Consider marking these parts as candidates for a PoC.
- **PoC or Pilot:** Consider if the solution should be delivered as one or more PoCs, or if it is more appropriate to prepare for production directly by delivering a pilot.

- **Post-production monitoring:** Work with the customer to define how the AI will be measured after production delivery. How do you continue to ensure on-going performance?

Phases of a successful ADS

BEFORE THE ARCHITECTURE DESIGN SESSION

Prior to performing the ADS, it is important to conduct a simple session with the customer to establish the scenario. This session is oftentimes referred to as ideation or opportunity definition. The goal is to establish the five Ws (who, what, when, where, and why) of their needs, which can be used as a guide for the ADS, streamlining the brainstorming process, and informing the agenda and milestone goals you will bring to the session.

- **Perform an envisioning session:** Prior to the ADS, perform an envisioning session to identify preliminary opportunities that would benefit from AI.
- **Schedule a time for the design session:** This is normally 1–2 days.
- **Schedule a location:** Ensure you have whiteboards and a projector.
- **Schedule resources:** This could include experts from your team and a cross-cutting panel of technical and business stakeholders from the customer.
- **Build an agenda:** Establish milestone goals in advance so that the ADS doesn't get consumed discussing a single topic.
- **Prepare preliminary documentation and architectural diagrams:** Even if you only have the basic building blocks, it's good to come prepared with something you can modify during or after the session.

DURING THE ARCHITECTURE DESIGN SESSION

Begin by reviewing requirements with your customer. As you conduct the design session, whiteboard the requirements and proposed solutions, and arrive at a consensus for each major topic. During this time, be sure to capture photos of the whiteboard so as not to lose your notes.

There are typically the following phases during an architecture design session: discovery, envisioning, and planning.

DISCOVERY

- Customer background and business technology strategy
- Project background and its drivers/aims
- Functional and non-functional requirements
- Usage scenarios
- Technology landscape
- Data assets

ENVISIONING

- Key functions and capabilities
- Components of the solution
- External connections and integration points
- Security considerations
- Abilities considerations
- Map requirements and scenarios to components

PLANNING

- Establish proof points
- Exclusions, risks, and issues
- Pre-requisites
- Deliverables
- Resources
- Escalation, communication, and long-term plans

AFTER THE ARCHITECTURE DESIGN SESSION

During the ADS, you and your customer planned a high-level architecture framework and conceptual design for a solution that addresses their organization's business goals and technical requirements. In addition to a summary of the engagement, you'll deliver information about:

- Special areas of concern to your customer's organization, such as security, compliance, and compatibility.
- Deployment scenarios that map to established deployment and practices and that cite specific examples where applicable.
- Familiarity with the Microsoft technologies proposed for the solution, in addition to any trade-offs among the differing technology options.
- The capabilities of your solution to deliver business performance on premises or in the cloud.

The outcome should be polished architecture diagrams that can be reviewed and signed off on by the customer. If one or more proof of concepts are desired or a pilot is determined to be the path forward, provide a plan and a timeline to deliver.



Go-to-Market and Close Deals Guide

Leverage the Microsoft resources available in the [Go-to-Market and Close Deals guide](#), for details on marketing to the cloud buyer, aligning marketing goals with business goals, developing value propositions, and marketing and sales assets, resources, and best practices.



Microsoft
Partner
Network

Optimize & Grow

Artificial Intelligence

aka.ms/practiceplaybooks



Executive Summary

So far we've covered strategies for building your AI practice, finding and keeping customers providing them with ongoing support.

In this section, we'll focus on how to optimize your practice, strengthen your relationship with customers, and evaluate your performance to help you continue to delight prospects and customers.

Are your customers delighted by your services and products? Delighted and not just satisfied? In this section, you'll learn why customer lifetime value is so important, and how to create more customers for life. We'll share how to get to know your customers better by following their journey with secret shopping and analysis. We'll also explore the use of a "land and expand" strategy and see how getting to know your customers better can lead to incremental opportunities to provide additional services.

You will also discover ways you can keep your solutions top of mind for prospects and customers through nurture marketing, and how to grow and improve your lead generation practice through a well-planned referral marketing program. We'll show you how to make the most of your renewal process, and how to get ahead of deadlines.

We will help you learn how to grow your business by identifying the best customer personas and creating "look-alike" prospects, deepening your expertise in key verticals and marketing that expertise, and collaborating with other partners to offer your customers a more comprehensive level of service and support.

We will end by discussing how important it is to create advocates for your company. This includes turning a customer into a fan and collecting testimonials to create case studies that can be used in future marketing campaigns. Map your customer's experience and ask for feedback to ensure you are turning satisfied customers into delighted customers who can wait to tell your story!

Use the strategies we provide in this section and in the [Optimize and Grow guide](#) to optimize and grow your practice.

Top 5 things to do

Learn from your customers and experience to optimize your practice, and expand to new markets through strategic partnerships. These are the top 5 things you should do to optimize and grow your practice.

- Gather feedback from your customers
- Nurture existing customers
- Turn customers into advocates
- Generate referrals with marketing
- Nurture strategic partnerships

Understanding Customer Lifetime Value

Customer lifetime value (CLV) is the revenue from a customer over the lifetime of their relationship with you.

As most businesses have experienced within the tech industry, a lifelong customer is of far greater value than any one-off transaction. It's no longer enough for companies to invest their time and resources into the generation of single purchases. Especially in the cloud world, it is critical for businesses to develop relationships and solutions that engage a customer for life.

CLV allows businesses to step back and look at not just one sale, not just one customer, but the customer base as a whole. It's about defining the economic value of each customer within that base and using that metric to make data-based decisions. If you don't know what a client is worth, you don't know what you should spend to get or keep one.

Knowing the CLV helps you make critical business decisions about sales, marketing, product development, and customer support. For example:

- **Marketing:** What should my acquisition costs be?
- **Sales:** What types of customers should sales reps spend the most time on trying to acquire?
- **Product:** How can I tailor my products and services to my best customers?
- **Customer Support:** How much can I afford to spend to provide customer service to my customers?



CLV is also a good way to guide and reward your sales team. Pay them more for bringing in customers with high potential lifetime value. By measuring and monitoring your cloud customer CLV, you can:

- Gain insight into your customers' cloud consumption and usage.
- Qualify for MPN cloud competencies that will help you grow your business.
- Help your customers reach their desired business outcomes.
- Leverage insight for cross-sell/upsell and proactively engage customers for extension opportunities.

By increasing your customer adoption rates, you can increase your CLV, particularly with cloud customers. The more employees you can get to use your service or solution, the more likely you are to increase CLV. Here are some ways to improve adoption:

What do you need to do to be successful?

Key attributes of a successful adoption approach

Define a vision & identify business scenarios	Prioritize solution & create an adoption plan	Commit resources & execute an adoption plan	Measure, share success, & iterate
A deep understanding of the business goals, as well as people challenges and needs to achieve them.	A solution that people love and that helps them achieve business goals and get things done more effectively.	A strategy to drive adoption including communications, readiness, and community.	A benchmark, KPIs, and success stories to help demonstrate success internally, improve, & expand.



While it's important to define and track your metrics, there are some simple things you can do to increase your CLV.

Your customer strategy must be built on a culture of customer success and tracking customer satisfaction. Critical success factors, when it comes to CLV, are:

- Being sensitive to customer emotions.
- Maintaining good communication.
- Listening to customer pain points.
- Understanding that there are multiple layers to any one concern.
- Doing business with an understanding and empathy for where your customer is coming from.

Next Steps

- Identify customer lifetime value of your five best customers.
- Identify your average customer lifetime value along with your Microsoft contact.
- Identify actions to increase your average customer lifetime value.

Learn more about CLV in the Modern Microsoft Partner Series eBook, [Deliver Customer Lifetime Value](#). To model CLV as it relates to your business, explore the modeling tool available on the MPN portal.

Guide: Optimize and Grow

Leverage the Microsoft resources available in the [Optimize and Grow guide](#), for details on building customer lifetime value, executing nurture marketing efforts, optimizing and growing from feedback, refining your customer value proposition, growing partnerships, and measuring results.




AI Playbook Summary

Thank you for taking the time to review this playbook. We hope you have gained new insight on how to successfully grow your AI practice by taking advantage of unique offerings from Microsoft, engaging with your customers, & forming strategic partnerships.

Our goal, when creating this playbook, was to establish a baseline definition of Artificial Intelligence, organize resources and provide insight that you can use to quickly accelerate or optimize your AI practice. To this end, we laid out the practice's opportunity, emphasized that AI is a broad opportunity that presents itself across industry verticals, and provided relevant information on business strategies and technical topics to capitalize on the opportunity, within five sections that you can review in order, or individually at any time.

In the section, **Define Your Strategy**, we helped you define the strategy upon which to build your practice. Here we provided an introduction to Microsoft's approach to AI and the technologies you can leverage from the Microsoft AI platform, and examples of the various project services, managed services and intellectual property your practice could sell. The services critical to the AI practice that we detailed were the envisioning session, the proof of concept and support. The key actions we prompted you to take are: identify your unique value proposition, develop your solution offer leverage the Microsoft Partner Network, and plan your support options.

In **Hire & Train**, we focused on the importance of hiring the right team, including data/AI scientists, data developers and data engineers), and provided details around the skills, certifications, and experience you should look for in each role. In addition, we provided specific guidance for ongoing training and certifications, including those from Microsoft like the MCSA and machine learning.



In **Operationalize**, we suggested you put your plan into action. We provided detailed coverage of the Microsoft Team Data Science Process (TDSP) that you can use to create a repeatable process for delivering AI solutions. We recommended that you leverage your internal use benefits to get your Microsoft licenses and subscriptions to help reduce the costs of your AI solution in Azure, create your key contracts, set up your support process, set up your social offerings and organize your engagement process into checklists.

In **Go to Market & Close Deals**, we emphasized getting your practice off the ground by defining your sales process, building materials to support sales and marketing, finding new customers, and nurturing and investing in them to build lasting relationships. We examined how marketing to the AI buyer requires a different approach as most prospects are not asking for AI specifically. Similarly we identified how the sale is different – more akin to selling an on-going experiment than a one time project.

In **Optimize & Grow your Practice**, we stressed the importance of learning from your customers and experience with post-mortem analysis to optimize your practice and help expand to new vertical markets through strategic partnerships. The top five actions we recommended: gather feedback from your customers, learn from your project successes and failures, create case studies and a marketing plan to expand into new vertical markets, maximize your efficiency and profit to fuel growth, and establish and nurture strategic partnerships.

© 2018 Microsoft. All rights reserved.