

How AI Gets You from Insight to Action

Artificial Intelligence (AI) and machine learning (ML) can predict opportunities, automate your business processes, and drive unprecedented levels of agility. The primary challenges to adopting AI/ML are infrastructure and talent. Find out how the right partners can lower the barriers and unlock the potential in your data.

Moving to predictive

Someplace locked deep in your customer service support logs lies an insight that could improve customer satisfaction by over 20 points. Your marketing database may be teeming with ROI optimization strategies that are lying dormant, costing you millions. Your supply chain analytics system has billions of inventory data points; is there a pattern that can improve forecasting and automate order schedules?

This data is gold for leaders who are looking to disrupt their industries before they are disrupted. The question is, how to get at it? How to use that data in a way that predicts future trends, makes decisions automatically, increases your competitive advantage?

You've heard the buzzwords and read the hype. The promise of Artificial Intelligence (AI) and the revolution yet to come still seems like science fiction, particularly when compared to the real challenges of the here and now.

But smart and forward-looking organizations are already solving their here-and-now problems with AI, using it to move from big data analytics solutions that describe what's happened

in the past and prescribe recommendations to machine learning-based systems that predict what's going to happen and take action automatically.

What can I do for my business?

It seems like only yesterday when 'Big Data' was the buzzphrase du jour. But once your organization has gotten past the initial stages of a large-scale analytics deployment, AI is the logical next step, not to mention a critical protection for companies that are otherwise in danger of disruption.

Let's look at what you need to get there.

First, does your team have the skills, knowledge, and real-world experience to build, train, and deploy machine learning applications? Second, is your infrastructure prepared for the demands of machine learning applications, and are you prepared to upgrade or invest in new systems to support AI?

With the need for highly powerful compute platforms and systems that can handle massive amounts of data in real time, keeping up with ML may seem like a costly arms race. But choosing the right

AI IN ACTION 

Hotels.com entices visitors

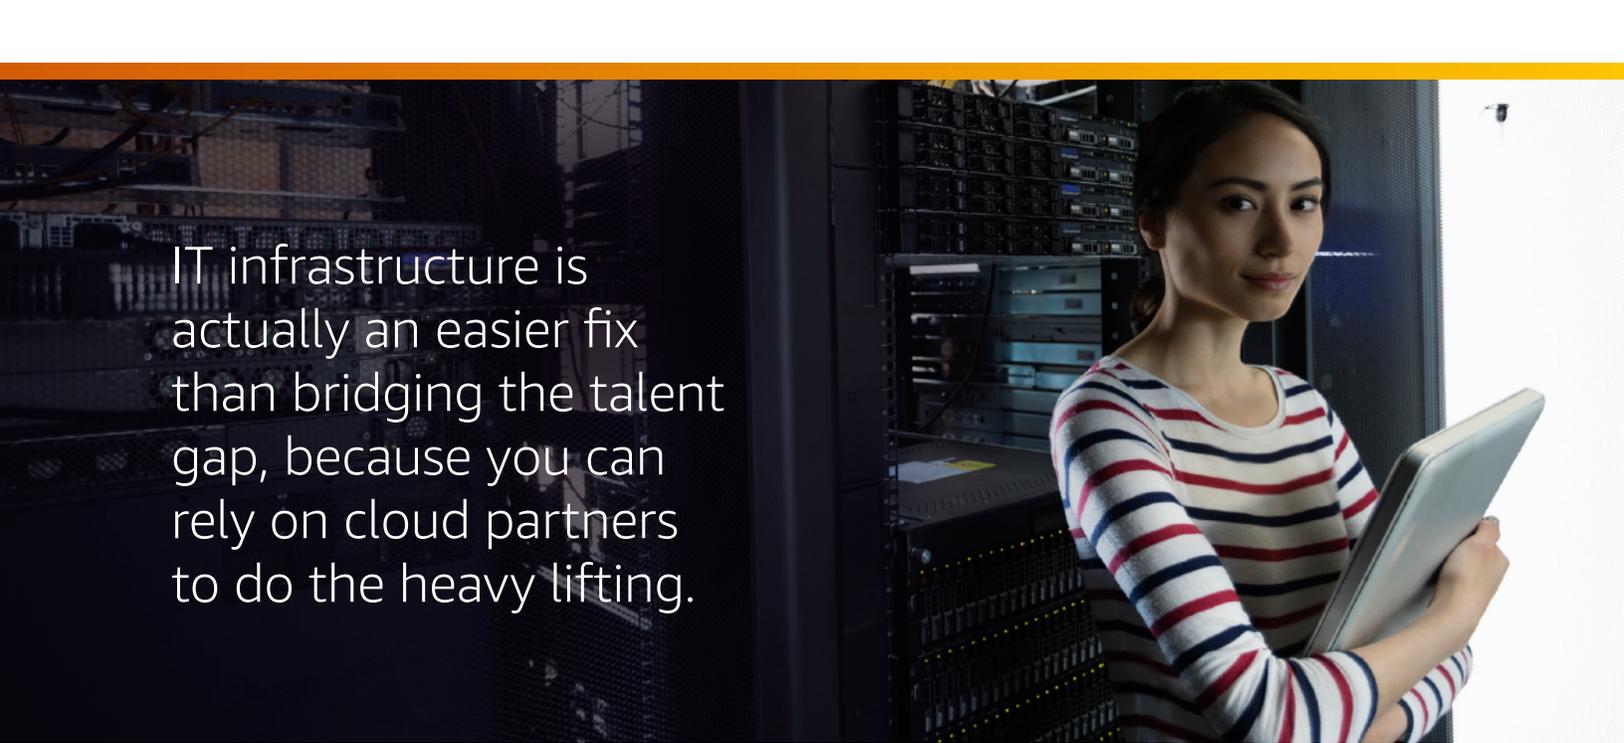
Target: Sales and marketing

AI used: Sentiment analysis

Hotels.com, a global lodging brand that operates 90 localized Web sites in 41 languages, is always innovating to move ever faster in a rapidly changing market. The company found that attractive photos incentivize site visitors to book rooms.

To identify the most appealing photos, Hotels.com enlisted Amazon Mechanical Turk, where human reviewers initially tagged thousands of images to 'teach' an ML system how to evaluate photos. The system now analyzes millions of photos across hundreds of thousands of properties. Running this project on AWS cloud resources reduces the time needed to rank images from one week to one day.

The group also uses AI to analyze sentiments, objects, and geographic locations across more than 25 million reviews to guide customers toward travel experiences that are uniquely suited to their preferences.



IT infrastructure is actually an easier fix than bridging the talent gap, because you can rely on cloud partners to do the heavy lifting.

partners and providers lets you take advantage of flexible, fast, and developer-friendly platforms while avoiding complexity and lock-in, as we'll see.

Solving for human and IT resources

For many companies aiming to develop ML applications, the harsh reality they face is a critical shortage of skills -- from data analysts and data scientists to software and machine learning engineers. There are far more job openings available for workers with high-value data and development skills than there are candidates to fill those jobs.

The solution for this challenge is twofold:

Turn developers into AI experts:

Empower existing developers who lack AI knowledge and experience with the tools they need to start developing, training, and deploying ML. Because the field is so new and standards are far from set, these tools must be as open and flexible as possible, capable of supporting virtually any current ML framework. At the same time, they need to be fast; there's no time for companies to take key developer talent out of production for six months to get trained up.

The ideal toolset will include application services that enable developers to plug-in pre-built AI functionality into

their apps without having to worry about the ML models behind the scenes, platform services that make it easy for any developer to get started building his or her ML skills, and frameworks and interfaces to support the established ML practitioner.

Secure the right infrastructure:

IT infrastructure is actually an easier fix than bridging the talent gap, because you can rely on cloud partners to do the heavy lifting. You're unlikely to have the specialized expertise needed to build and tune compute, network, and storage platforms to support ML in-house, but the cloud provides a relatively seamless and easy-to-consume resource stack for AI adoption that scales effortlessly and enables you to pay only for what you need.

AWS and Intel enable AI

Amazon Web Services (AWS) provides a smooth and simple on-ramp for companies looking to embrace AI. It supports a wide range of AI frameworks and tools, delivers a range of pre-built models to help developers get up to speed quickly, and it offers optimized resource stacks that are designed specifically for training and deploying AI models.

Here are the ways AWS and Intel address the challenges that companies face when considering AI adoption.

AI IN ACTION



Business Services Authority

NHS BSA improves call center performance

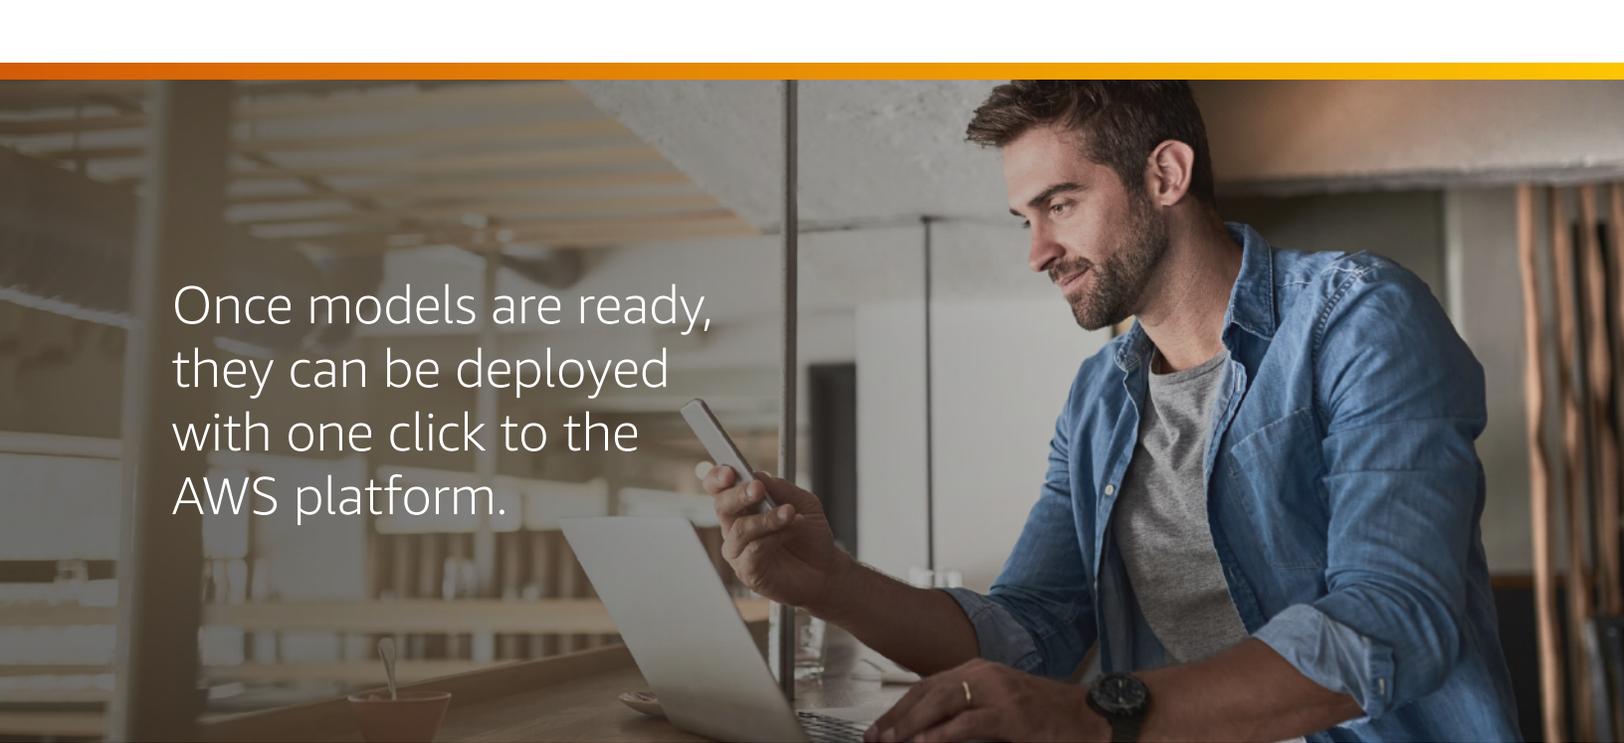
Target: Customer service

AI used: Natural language processing

The National Health Service (NHS) is the UK's largest employer and health provider. NHS Business Services Authority (NHS BSA) is a special health authority which provides a range of critical central services to NHS organizations, NHS contractors, patients, and the public. As such, their 450 call center staff handle around 5 million calls per year dealing with a range of routine requests for information related to the European Health Insurance Card (EHIC), and response times were lagging.

NHS BSA began routing incoming calls through an Amazon Lex-enabled chatbot, which understands natural language and intent without requiring callers to use specific phrases, with the goal to reduce call center load by providing simple answers to information, advice, and guidance on EHIC.

During a 3-week trial, Lex resolved 4,300 out of 11,000 calls, transferring the more complex queries to human operators. This 42% call deflection success rate led to happier customers and enabled call center workers, whose skills were put to better use, to manage more complex customer calls at a faster rate. Following the trial, the service is now live and will be expanding across other call center streams.



Once models are ready, they can be deployed with one click to the AWS platform.

Challenge: How do I get the right infrastructure?

Solution: You partner with AWS, and AWS gives you the right mix of compute, storage, and network resources to run your AI and ML workloads. Amazon Elastic Compute Cloud (EC2) delivers a variety of compute instances optimized for many AI scenarios. Amazon EC2 C5 instances, for example, run on customized Intel Xeon Scalable Processors and are ideal for training and deploying ML models. These instances deliver fast, reliable, and cost-effective performance for ML.

Challenge: How do I get the right talent?

Solution: You leverage AWS platforms and tools that help developers learn as they work. Amazon SageMaker is a fully managed machine learning platform that includes algorithm and model authoring tools, simplified connections to internal and cloud-based data sources, and a library of pre-built algorithms that are optimized to run on cloud resources. It's a one-stop shop for developers who want to learn and build ML applications. Once models are ready, they can be deployed with one click to the AWS platform.

In addition, AWS Greengrass, running optimized Intel AI technology, enables organizations to run inference models that were built and trained in the cloud

directly on IoT devices, delivering real-time insight locally and synchronizing data with cloud repositories when connections are stable.

To get your developers immersed in IoT and edge applications quickly, AWS DeepLens is the world's first ML-enabled video camera, able to run inference models on-board, using Intel Atom processors. AWS DeepLens enables developers to learn the basics of ML by exploring practical computer vision use cases, including pre-built models for face detection, object detection, activity detection, and more.

Challenge: How do I choose the right tools?

Solution: You find a platform that includes broad support. The ML space is still in its early days; you want to give your developers as many choices as you can, so they can pick the tools that work best for them. SageMaker supports model-building in all of the major frameworks available today, and it leverages Intel's optimized instruction sets to deliver unprecedented performance regardless of the developer's platform preference. With this kind of support, you can move forward confidently as the technology evolves.

Challenge: Where can I tap into experts who can help?

Solution: You look to experienced professionals who have assisted thousands of companies like yours with AI adoption. Intel's AI Academy provides your team members with essential learning materials, tools, technology, and a collaborative community that helps developers get from beginner to expert. Amazon's Machine Learning Solutions Lab pairs your team with Amazon ML experts to help prepare data, build and train models, and put models into production. ML Solutions Lab combines hands-on educational workshops with brainstorming sessions and advisory professional services to help you work backward from business challenges, and then go step-by-step through the process of developing ML solutions.

AWS and Intel offer an unmatched set of services and cloud infrastructure that are used by the broadest range of customers. In fact, more companies are doing machine learning on AWS than anywhere else.

Find out how AWS and Intel can help your organization embrace AI at ml.aws