
aws
intel
Unlock the potential of machine learning for executives in Australia & New Zealand
Lessons shared from leading local industry experts

Contents

aws machine learning

Machine learning is rapidly being developed, implemented and directed by executives to place their organisations as market leaders, drive growth, and futureproof their business.

SIMON JOHNSTON
Principal AI/ML Specialist ANZ
Amazon Web Services

Machine learning is enabling new features, products, and solutions that were not available just a few years ago. These solutions are rapidly bringing automation and augmentation capabilities to businesses and their workforce. This has led to an accelerated pace of innovation as machine learning also becomes a meaningful driver of value for many organisations. A direction that is allowing more businesses to find the balance between optimal business value and speed.

Machine learning has enabled businesses to transform value from their data to drive business growth from new customer experiences, reduce costs through process efficiencies from automation, or launch transformational initiatives aimed at innovation and competitive differentiation. The Fourth Industrial Revolution is the disruption of technology in every industry, transforming entire systems of production, management, and governance as defined by the World Economic Forum.[1] It is being witnessed by the pioneering business leaders of today through machine learning.

With the advent of the cloud, increased computing power, and the proliferation of data, the last decade has increasingly revealed the business value of machine learning within the enterprise. The democratisation of machine learning through the ease of access to managed services and platforms is reducing the barriers to entry for companies. However, the scale of adoption within Australia is still in the early stages, offering many exciting opportunities for businesses.

Australia's economic growth accelerated by machine learning innovation

The Australian Government recognises the importance of machine learning to the economy, as indicated within the AI Action Plan,[2] which outlines key societal and economic benefits: "The economic potential is enormous. Digital innovation, including new artificial intelligence technologies, is predicted to be worth $315 billion to the Australian economy by 2028".

The rapid growth that Australia is experiencing through digital technologies and innovation[3] shows that the need for enterprise leaders to embrace machine learning capabilities is immediate. Gartner forecasts that "by the end of 2024, 75% of organisations will shift from piloting to operationalising artificial intelligence, driving a five times increase in streaming data and analytics infrastructures."[4] The use cases, customers, and partner stories highlighted within this eBook demonstrate how businesses are capitalising on their data through implementing machine learning in their organisation. The Australian enterprise is ready to evolve.

Pathway to success

In this e-book, each customer and partner story highlights a journey. Companies that have reached the stage of delivering value with machine learning have been on an insightful journey, one which involves investment both financial and time into people, processes and technology. To support customers to accelerate this journey, Amazon Web Services (AWS) has developed a workflow based on experience from hundreds of thousands of customers. It is an incremental approach focused on delivering end-to-end machine learning solutions into production to deliver early value, with both tangible and intangible results, while building up the skills, operational change management, and technical capabilities required. The business potential of machine learning goes far beyond these selected use cases and is not a simple start stop – it is one that evolves over time for the business. This journey is depicted on page 34 of the eBook.

1 https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/
2 https://consult.industry.gov.au/digital-economy/ai-action-plan/supporting_documents/AIDiscussionPaper.pdf
3 https://www.weforum.org/agenda/2016/01/the-fourth-industrial-revolution-what-it-means-and-how-to-respond/
4 https://www.gartner.com/en/newsroom/press-releases/2020-06-22-gartner-identifies-top-10-data-and-analytics-technolo

Discover, deliver, scale: Inside nib Group's machine learning journey

"We pride ourselves on leading the way in machine learning. We want to make it quicker and easier for members to claim with us, and equally to be able to ensure our front-line teams are able to focus on important enquiries, rather than answering basic policy questions that can be answered via our chatbot."

Mathew Finch
Head of Emerging Technology and Data Platforms at nib

A leader in automation

With more than 1.4 million members in Australia and New Zealand, nib Group (nib) is one of the leading health insurers in the region. By leveraging the latest machine learning innovations, it responds to members faster, more accurately, and at lower cost than their previous solution.

nib's team of developers have been experimenting with automation and machine learning for several years. Eager to enhance the company's machine learning capabilities further, they turned to data science consultancy Eliiza, which is a premier partner of Amazon Web Services (AWS).

Eliiza pulled together a team of engineers and data scientists to help design a machine learning engine called Melvin, which is built using Amazon Textract. They also enhanced nibby, nib's chatbot, using Amazon Kendra – two innovative AWS solutions to two very different customer service challenges.

The first challenge: processing more claims, more quickly, at lower costs

In 2015, nib launched an innovative tool that allows members to submit health insurance claims via a mobile app. After photographing and uploading healthcare receipts directly to the app, members are reimbursed for valid expenses very quickly – usually within 24 hours.

Despite being a huge step forward for members, nib's claims team were spending too much time extracting data from receipts – such as the customer number, medication, dosage, dates, and provider number – and typing this information into a database.

"We began experimenting with ways to automate the rather routine tasks of transcription and data entry. Our end goal is to process more claims, more easily, in less time, freeing up the company's claims team to focus on more critical steps in the verification process," says James Wilson, CEO at Eliiza.

"We set out to build a machine learning engine that 'reads' data from receipts and pre-fills corresponding fields in a database. Our solution needed to be as fast and accurate as nib's claims team. It also needed to slot neatly into their existing claims processing architecture without compromising sensitive medical data."

A seamless approach to integrating machine learning into nib's existing architecture

nib's machine learning journey kicked off with a 'Discovery' workshop where stakeholders came together to identify the challenges that needed to be solved. This led to the creation of a Machine Learning Blueprint to guide the next steps and iterations.

In the 'Delivery' stage, Eliiza and nib began testing a small number of mock documents using different technologies such as Optical Character Recognition (OCR), image recognition technology used to convert image-based text such as handwriting or printed material into machine-encoded, electronic text. Ultimately, one solution emerged ahead of the rest: Amazon Textract, a machine learning service that automatically extracts text, handwriting and data from virtually any document.

"It quickly became apparent that Amazon Textract sets the global benchmark. We couldn't find anything that could match its performance, both in terms of its accuracy, and its ability to read lower-quality images. It spoke to a lot of the challenges we faced," says Finch.

There was a slight hitch: Amazon Textract was yet to launch in Australia. Undeterred, Eliiza and nib began using Amazon Textract to build their machine learning engine, dubbed Melvin, in early 2019, putting the pipeline on standby until Amazon Textract's Australian launch later that year.

"This actually had several benefits. It gave us time to implement error detection and correction mechanisms and deepen the level of automation. By integrating with other databases like the MIMS drug database, our solution can also verify the validity of claims," explains Wilson.

Melvin's components were developed and tested locally using existing nib data structures. Upon completion, Melvin slotted directly into the company's processing framework using Amazon Simple Queue Service (Amazon SQS), which facilitates asynchronous processing, and Amazon Lambda. Meanwhile, to minimise the risk of exposing sensitive customer data, Eliiza created a quarantine cage to ensure this data never leaves the AWS environment.

Scaling up to iterate faster

Since May 2020, nib has been using Melvin to extract and automatically enter data into its database with astonishing accuracy. Amazon Textract's accuracy exceeds 87% for pre-filled fields overall, while the accuracy of almost half pre-filled fields is 95% or higher. Around half of all claims are processed without any manual re-keying or data entry adjustments. All that's left for nib's team is to quickly review and process each claim.

As a result, Melvin shaves around 20 seconds off every claim processed, which means the company can process significantly more claims every day.

Melvin already processes more than 150 claims per day, and nib is exploring whether a percentage of claims can soon be processed with zero human intervention. "Ultimately, that's the end goal. Ideally, we would like to see a large percentage of claims processed automatically in real time, so that members receive a response within minutes. That's what we're building towards and data extraction is key to that goal," says Finch.

The next hurdle: meet nibby

nib launched its chatbot, nibby, in 2017. Created in partnership with AWS partner DiUS, it was built using Amazon Lex – the same deep-learning technology that powers Amazon Alexa. The chatbot deflects basic policy coverage questions away from nib's call centres, freeing up agents to focus on more complex enquiries.

"We chose Lex because it's a complete package, which we could seamlessly integrate with other aspects of our AWS environment. Lex took care of the heavy lifting without needing to employ data scientists to create machine learning models. We developed a proof of concept within four to six weeks," explains Finch.

The chatbot has been a huge success. Today, around 65% of chat-based enquiries are handled by nibby while only 35% are diverted to contact centres for resolution. The insurer now handles around 15,000 chats per month up from 4,000 chats when nibby first launched.

To build on this success, in 2019 nib teamed up with Eliiza to further improve nibby's speed and accuracy and better analyse nibby's data. But how?

Deep automation and intelligent search are a click away

By entrusting labour-intensive, repetitive tasks to AWS, they began experimenting with Amazon Kendra, an intelligent search service that is powered by machine learning.

When a member asks, "Am I covered for x-rays?", nibby can now respond in precise detail. Previously, nibby would have shared a link to a policy document. Unfortunately, most people avoided reading these documents and will most likely reach out via phone or online chat instead. Today, nibby has the content it needs to provide immediate, accurate responses to member questions about their specific insurance policy. To achieve this, Eliiza used Amazon Kendra to build nibby's Kendra Index. The Index was created by ingesting over 40 product disclosure statements, each several pages long, along with lengthy policy PDFs and FAQs.

"The setup was incredibly quick. With just a few clicks, Eliiza easily configured our Kendra Index and connected it to relevant data sources," says Finch.

James Dunwoody, Machine Learning Engineer at Eliiza, adds: "Unlike conventional search technology, Kendra's natural language search capabilities help nibby to answer questions quickly and accurately – no matter how deep the information lives within the Index."

This solution is on track to roll out in 2021. Eliiza is also working with nib to leverage nibby's chat data, with over 300,000 conversations collated so far.

Ultimately, nib and Eliiza hope to expand nibby to handle telephone enquiries too.

"We're now using Amazon Connect to make nibby not just chat-oriented but voice-based. This means our members will be able to converse with a very human-sounding bot, deflecting even more enquires away from our contact centres," says Finch.

"We receive 150,000 calls per month – that's 10 times more than our chat-based enquiries. If we can deflect just 10% of these calls that will represent huge savings and efficiency gains."

Going above and beyond: how nib has surpassed its self-service targets

By entrusting labour-intensive, repetitive tasks to AWS technologies, nib is on track to continually improve its self-service targets – a term that describes the percentage of member enquiries that require zero human intervention.

"One of our key metrics for success is self-service, which is how we judge all the work we do in the chat/voice and member enquiry space," says Finch.

Eight months ago, the insurer's self-service rate was between 35% and 40%. Today, it is 65%, which means just 35% of chats are deflected to contact centres – the rest are expertly handled by nibby, freeing employees up to handle the more complex cases.

"We never thought we'd get our self-service rate past 50% so we're really happy with 65%. This results from a combination of a few different technologies, including Amazon Kendra and Amazon Lex, which now sit behind nibby. Our goal is to keep pushing our self-service rate up as high as we can."

From Amazon Textract to Amazon Lex, and Amazon Kendra, nib has embraced a suite of AWS solutions to challenge the status quo, and deliver an innovative first for the industry.

It's an exciting era for data-driven companies that embrace artificial intelligence and machine learning in innovative ways. By automating time-intensive tasks like data entry or fielding member enquiries, nib is one of the disruptors leading the way.

Benefits

- Saves an average of 20 seconds handling time per claim by extracting and pre-filling data with accuracy of 87% to 95%
- 50% of claims processed require no further human intervention or data entry corrections, allowing employees to focus on more complex cases
- In just 18 months, nib's self-service rate has increased from 35–40% to 65%
- nib chatbot, nibby, now fields 15,000 chats per month, up from 4,000 in 2017

Find out more here >

"It quickly became apparent that AWS Textract sets the global benchmark – we couldn't find anything that could match its performance."

Mathew Finch
Head of Emerging Technology and Data Platforms at nib

Domino's Pizza Enterprises delivers in record time using AWS for predictive ordering

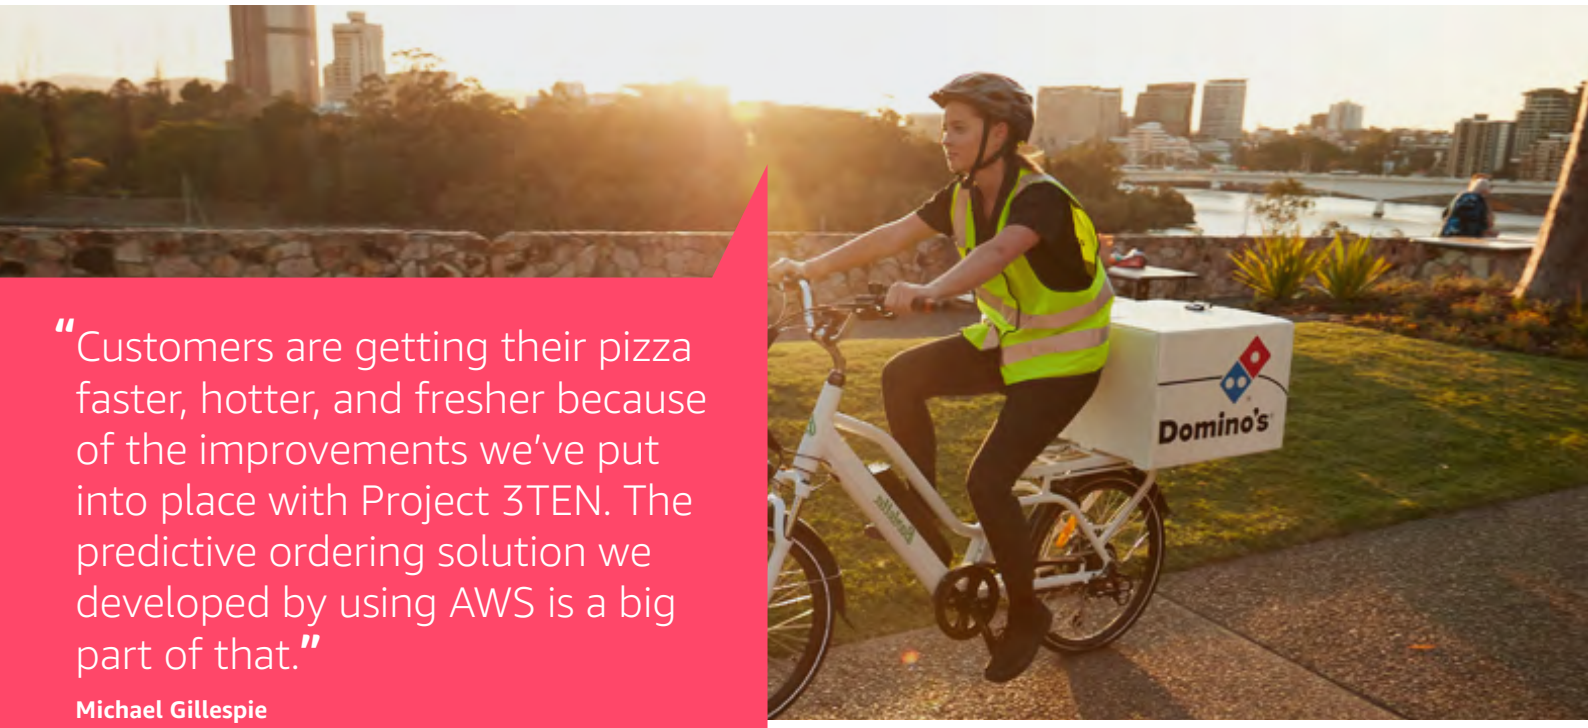

"Customers are getting their pizza faster, hotter, and fresher because of the improvements we've put into place with Project 3TEN. The predictive ordering solution we developed by using AWS is a big part of that."

Michael Gillespie
Chief Digital and Technology Officer
Domino's Pizza Enterprises Limited

When it comes to the global pizza business, Domino's Pizza Enterprises Limited (Domino's) has a large slice of the pie. The company, which is the largest Domino's franchise holder, represents the Domino's brand in Australia, New Zealand, Belgium, France, the Netherlands, Japan, Germany, Luxembourg, and Denmark. Domino's maintains a network of more than 2,600 stores globally and is based in Brisbane, Australia.

Domino's is an increasingly digital business with more than 70% of sales coming from online orders. "Our investment in technology is a key ingredient to our growth as a business," says Michael Gillespie, Chief Digital and Technology Officer for Domino's. "We strive to use it to reduce pick-up and delivery times because we've identified that the sooner we can get a pizza to our customers, the more satisfied they are with their meal."

To enable faster pick-up and delivery, the company recently launched Project 3TEN, an initiative that aims to have a pizza ready for pick-up within three minutes or safely delivered within 10. This is all possible through efficiencies in cooking methods and transportation and by opening more stores closer to customers. To support this initiative, the company wanted to use predictive technologies to help reduce pizza-making and delivery times. "We can use technology to increase the speed of our ovens and we can give our drivers more efficient technology like e-bikes or scooters but we also needed to look at how we can help stores anticipate what customers order," says Gillespie.

Developing accurate, predictive ordering on AWS

Domino's turned to Amazon Web Services (AWS) and Max Kelsen, an Advanced Consulting Partner in the AWS Partner Network (APN), to help create a predictive ordering solution. "We knew we needed a smart, accurate system, and we liked what we saw in AWS machine learning technologies," says Gillespie. "We knew that by using AWS services we could develop a solution that would give our stores a glimpse into the future by predicting what pizzas would be ordered next."

The company created a data lake consisting of key order information by taking advantage of Amazon Simple Storage Service (Amazon S3) for data storage and AWS Glue for data querying. It also uses Amazon SageMaker to build and train machine learning models to predict the likelihood that an order will be placed. A store can then begin making that order even before it is actually placed.

As a trial, Domino's initially deployed its predictive ordering solution in some of its stores in Australia. Store employees could view an ordering screen displaying specific pizzas with various colour indicators corresponding to the likelihood of those pizzas being ordered. "This isn't making pizzas and leaving them in a hot box for half an hour – this is getting the pizzas lined up, coming out of the oven, and ready to go as an order is placed," Gillespie says.

The company enhanced the solution during the trial so it would be ready for a wider rollout. Domino's then began deploying the solution to stores in New Zealand, France, the Netherlands, Japan, and Germany.

From order to delivery in 10 minutes or less

Using the AWS-based solution, Domino's has given its stores a tool to help drive down pick-up and delivery times for customers. For example, in 2019, a Domino's store in Australia averaged delivery times of under five minutes, from order to doorstep, across an entire week. "It's exciting that nothing changes from the customer's perspective, except that the post-order experience can be much quicker," says Gillespie. "Customers are getting their pizza faster, hotter, and fresher because of the improvements we've put into place with Project 3TEN. The predictive ordering solution we developed by using AWS is a big part of that."

Domino's is gaining a competitive advantage in the marketplace. "When our customers are hungry, they're hungry now, and we want to deliver their meal as quickly and safely as possible," says Allan Collins, Domino's Chief Marketing Officer for Australia and New Zealand.

Enabling fast, easy deployment for franchisees

Another advantage is that franchisees don't need to spend much time being trained. "One thing we always try to keep in mind when developing technology is ease of use – the technology has to be accessible and easy to implement," says Gillespie. "We knew if we had barriers to deployment, it would be challenging to roll this out." Domino's Technology Team worked hand-in-hand with the Operations Team to ensure a smooth rollout, with all new processes seamlessly dovetailing into existing operations.

Making customers happier

The solution is contributing to a boost in customer satisfaction. "Our data shows that the stores with faster pick-up and delivery times have higher customer satisfaction scores," says Collins. "And those customers are more likely to come back and recommend our brand to others."

Domino's has deployed the solution to stores in a number of other countries. "AWS has been a great company for us to work with because it's helping us provide a new, unparalleled level of service to our customers," says Gillespie. "In addition, we hope to take what we've achieved with AWS so far and push it to new heights in the future."

To learn more, visit aws.amazon.com/sagemaker and aws.amazon.com/travel-and-hospitality

About Domino's Pizza Enterprises Limited

Based in Brisbane, Queensland, Domino's Pizza Enterprises Limited is the largest pizza chain in Australia and the world's largest franchise holder for the Domino's Pizza brand. The company has more than 2,600 stores in Australia, New Zealand, Belgium, France, the Netherlands, Japan, Germany, Luxembourg, and Denmark.

Benefits of AWS

- Assists Domino's stores in achieving goal of pizza delivery in 10 minutes or less
- Deploys accurate, predictive ordering solution quickly and easily
- Enables fast, easy deployment for franchisees

Find out more here >

> "That's our market differentiator. Some of our competitors still take 45 minutes to an hour to deliver a meal, when we can have a pizza delivered in 10 minutes or less following an order. That really impresses our customers."
>
> **Allan Collins**
> Chief Marketing Officer ANZ
> Domino's Pizza Enterprises Limited

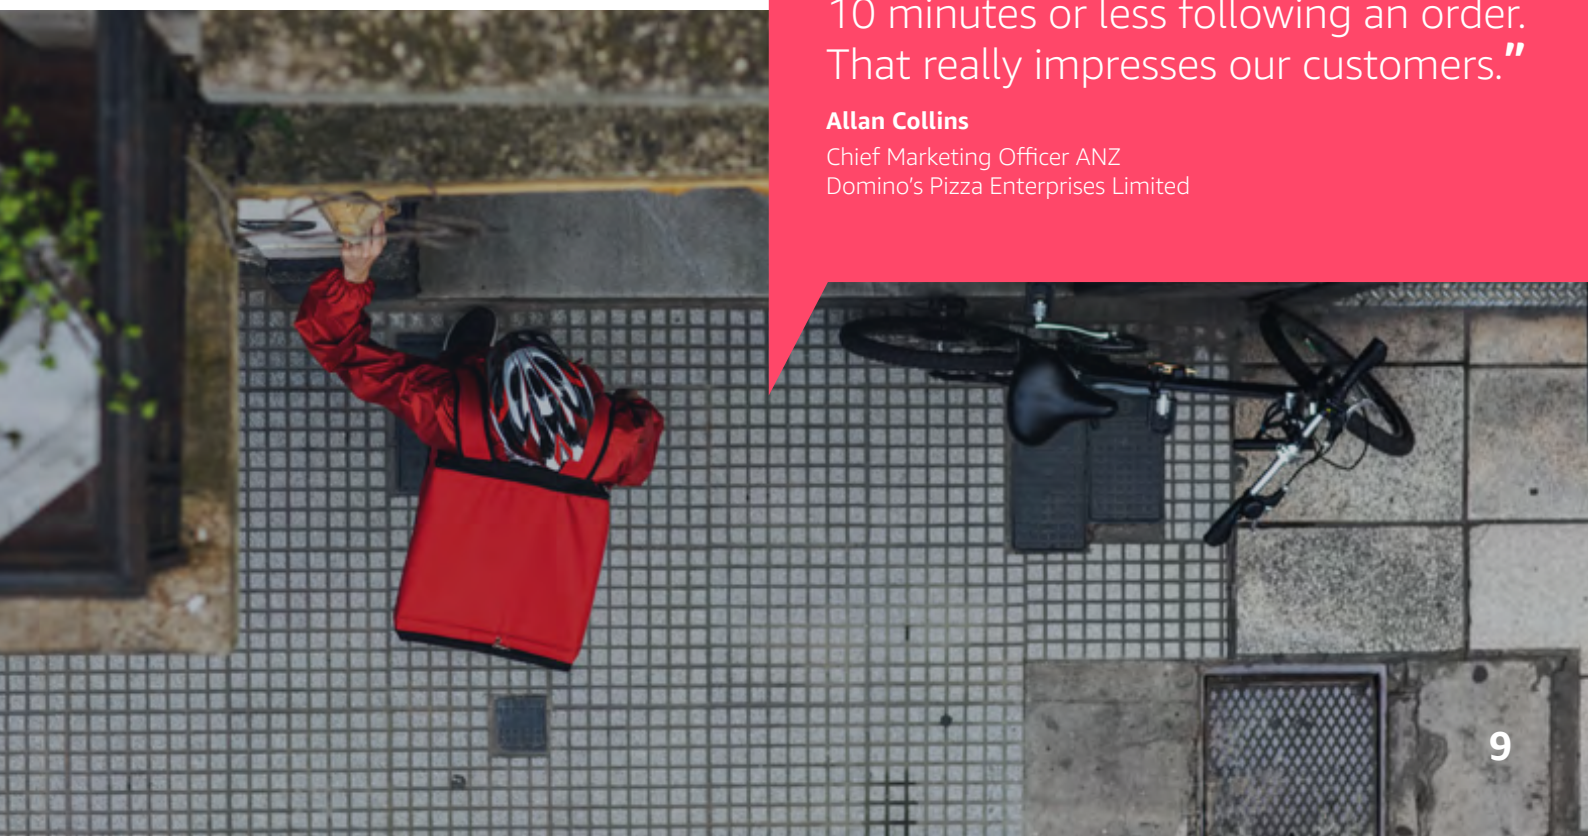

Tape Ark

Tape Ark and AWS develop an out-of-the-box archiving system

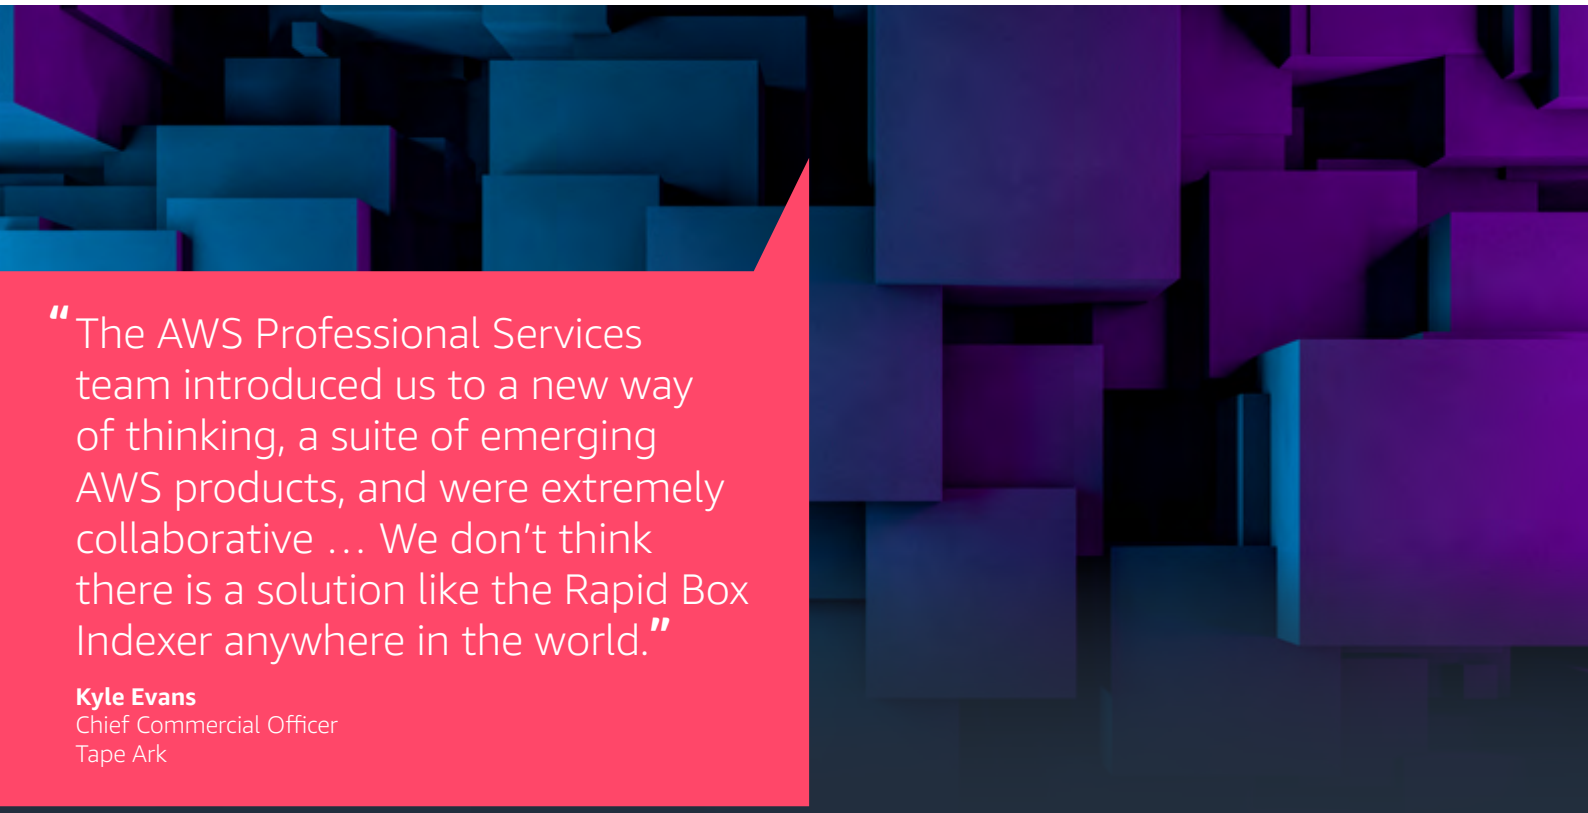

> "The AWS Professional Services team introduced us to a new way of thinking, a suite of emerging AWS products, and were extremely collaborative … We don't think there is a solution like the Rapid Box Indexer anywhere in the world."

Kyle Evans
Chief Commercial Officer
Tape Ark

Tape Ark is on a mission to help organisations to manage their backup data and ageing corporate documents by leveraging data analytics, artificial intelligence and machine learning. As the world's leading specialist in tape-to-cloud migration, having processed over five million data tapes to date, Tape Ark is now helping customers to take control of their paper-based data.

For many industries, storing documents is a legal requirement of doing business, yet the document management sector – worth $529 million in Australia alone, according to IBISWorld – is ripe for disruption.

"More often than not, companies tend to lose track of what's inside each box. This makes it difficult to make decisions about box disposal or document scanning. As time passes and their archive box collection grows, the problem becomes even more unwieldy," says Kyle Evans, Chief Commercial Officer at Tape Ark.

One of Tape Ark's customers, for example, has more than one million boxes in storage – its collection dates back to the 1930s.

"We realised that storing paper documents in archive boxes is a pain point for many companies. In order to decide which documents to keep, digitise or destroy, they traditionally look to scan all of the documents, or bring the archive boxes back to their office for a subject matter expert to manually review every box. The default option is to do nothing but this means forking out monthly warehousing fees for the foreseeable future. We were already using artificial intelligence and machine learning in other areas of our business when we had the idea of using these tools to invent a smarter solution: one where they can learn what is inside each box for a fraction of the cost," says Evans.

The key steps in Tape Ark's machine learning journey: Discovery, Delivery, Scale

Tape Ark's machine learning journey began with a discovery workshop in July 2020 led by the AWS Professional Services (ProServe) team. A dedicated team of data scientists, engineers and business experts came together to develop a Machine Learning Blueprint, outlining key steps and iterations. By December, Tape Ark was ready to deploy a proof of concept (POC) called the Rapid Box Indexer.

The Rapid Box Indexer allows organisations to view the contents of archived boxes from afar via Tape Ark's customer portal. It is the first service of its kind, according to Evans:

"Until now, companies had to physically retrieve boxes from offsite warehouses in order to audit their contents. Each box is opened by a member of staff, who manually examines, documents or scans items before adding them into a spreadsheet or database. A single box, depending on its contents, can take more than an hour to manually index and document. Multiply that by tens of thousands – if not hundreds of thousands – of boxes, which is typical for large organisations, and you're looking at endless hours of work," says Evans.

To automate this process, the Rapid Box Indexer uses machine learning to index and record box contents.

Intelligent image and video analysis software is also used to categorise information and add metadata, providing much deeper insights than traditional indexing methods.

Unpacking Tape Ark's suite of AWS machine learning innovations

"The ProServe team were amazing," says Evans. "They introduced us to a new way of thinking, a suite of emerging AWS products, and were extremely collaborative. It is great to work with a likeminded team who think big and really want to address the customer's problem."

Together, Tape Ark and ProServe built the Rapid Box Indexer using Amazon Textract, which uses machine learning to extract text, handwriting and data from virtually any document just like a person would. Amazon Rekognition is used to automate image and video analysis using machine learning and speech recognition software, while Amazon Comprehend enables entity detection and sentiment analysis. Amazon S3 is used to store and protect data, while Amazon Lambda enables very fast information processing.

Users can now 'see' inside each box using three layers of data – videos, images and text – via the Tape Ark portal, along with searchable tags. Armed with this information, it is much easier to decide which documents to keep, digitise, or destroy. They can also audit boxes long after video, images and text are captured – without recalling a single box.

"In our view, it's like having all of your offsite boxes with you onsite so you can flip through their contents as needed. This allows our customers to home in on whatever they are hoping to find," says Evans.

Scaling up: taking Tape Ark's Rapid Box Indexer to the world

After developing a POC in Australia, Tape Ark is now piloting the Rapid Box Indexer from its Houston facilities in the United States.

"We are currently using the Rapid Box Indexer to process 7,500 boxes, which is a reasonable scale for our first pilot. Customer feedback will help us to refine the Indexer and make it even more valuable. In 2021, we will start rolling it out en masse," says Evans.

"Our aim is to expand the Indexer to meet the needs of any industry and answer some of the hardest document management questions: Can we prioritise documents for digitisation? Will this help to reduce the size of storing physical box archives? How will this reduce monthly warehousing costs?"

According to AWS comparisons, the Rapid Box Indexer can be up to 20 times cheaper than traditional archiving systems. This estimate is based on the costs of storing boxes in warehouses, which are "significantly higher than the costs of digitising and storing data in the cloud," explains Evans.

About Tape Ark

Tape Ark is bringing the management of offsite, archive tape data into the 21st century by securely migrating ageing corporate data from tape media directly to the public cloud. By embracing digital and virtual data storage technologies, Tape Ark is re-imagining the way physical data is stored off-site, bringing physical tape storage into the new millennium.

Benefits of AWS

- Reduces physical warehousing costs by up to 20x
- Provides three layers of data – text, image and video – users can see what's inside each box, textually and visually, from afar
- Simplifies and automates low-value tasks like data entry and indexing to drive business efficiency

MECCA
BRANDS

How MECCA is mirroring in-store customer experiences online—with dramatic results

"At MECCA, we want to curate a seamless experience for our customers online. The challenge is providing the equivalent bespoke, customised service to our customers regardless of whether they are in a boutique or not."

Sam Bain
Chief Digital Officer
MECCA

Founded in 1997, MECCA brings together some of the world's most coveted beauty and skincare products, selling over 120 brands to two million customers in Australia and New Zealand. It represents 10% of Australia's $4.2 billion beauty market and around 25% of the prestige beauty sector, according to IBISWorld.

One of the secrets to MECCA's success is its high-touch, high-service beauty boutiques. According to Vogue, MECCA spends around 3% of its revenue training staff to offer tailored, personalised advice to every customer.

MECCA is now looking to replicate this unique face-to-face experience online.

"When you go into a MECCA store, the personalised service and recommendations you get from our hosts is something that we want to translate into the online world," explains Lauren Shepherd, Head of Customer Relationship Management and Loyalty at MECCA.

MECCA is an early adopter of e-commerce, launching online shopping in 2001. It turned to its strategic partners, Amazon Web Services (AWS) and Servian, for help leveraging data and analytics to improve customer engagement online.

Delivering a hyper-personalised product and a bespoke journey are at the core of improving customer experience (CX) to drive business growth. Yet according to Gartner, over 70% of CX leaders "struggle to design projects that increase customer loyalty and achieve results."

Using machine learning to deliver a highly tailored customer experience

MECCA's data transformation journey began in 2018, when it worked with AWS to build a data platform and employed Tableau as the business intelligence reporting tool.

After seeing a positive uptake across the business, MECCA's technology team wanted to push the envelope by using data to add more value for customers. Their end goal is to enable different parts of the business to self-service – a form of business intelligence that enables anyone to drill into data that is relevant to their role.

MECCA's CRM and Loyalty teams were especially eager to use data to predict which products are most likely to appeal to different customers.

"We had lots of information about our customers, but we weren't really using it to optimise their experience – each customer would receive the same content and promotions," says Shepherd.

The MECCA data team in partnership with Servian commenced a trial of Amazon Personalize, a machine learning service that provides a simple framework to build and run personalisation models. It is used by innovative retailers to provide specific product recommendations, personalised product re-ranking and customised direct marketing.

Kickstarting MECCA's machine learning journey: from Discovery to Delivery

Before it could transition to artificial intelligence-powered personalised marketing, MECCA needed to cleanse and restructure 23 years' worth of data contained within its email platform to create an organised, clean and easily accessible data system.

This began the 'Discovery' stage of MECCA's ML journey. An interdisciplinary technology team of CRM, data and Servian experts came together to develop a Proof of Concept (POC). It was their job to anticipate and address any roadblocks that would derail the POC, which uses Amazon Personalize to provide product recommendations as part of a targeted mascara marketing campaign. By utilising existing metadata, the campaign targets customers 90 days after their last purchase with mascara promotion emails.

Moving into the 'Delivery' stage, the technology team made multiple changes such as migrating to an Amazon S3 data environment and building new integrations to the marketing system to allow for automated end-to-end data processing. They also spent time helping executive stakeholders and managers understand how personalisation and the shift to machine learning will positively benefit their respective areas of the business.

This was key to the POC's success, according to Paul Erskine, IT Delivery Manager for Digital and Data at MECCA: "Many executives had reservations about the complexity of data science in general. They had questions like: 'Who will manage the model if someone leaves the business? What is the cost versus value? Who will support it over time?"

To address these concerns, the technology team shared its vision and plans at a data governance forum. They revealed the product recommendations generated by Amazon Personalize, and explained how data science can be used to optimise customer conversion and engagement rates. Their evidence was so compelling, MECCA's executive team gave the POC a green light.

The final stage: Scaling up to deliver self-service and dramatic conversion results

When Amazon Personalize was released in Australia at the end of 2019, MECCA and Servian immediately deployed it – one of the first Australian companies to do so. Within hours, it began to produce tailored product recommendations, and today, it produces product recommendations across MECCA's entire catalogue.

"One of the benefits of using Amazon Personalize is the ease of training custom models using existing data on AWS managed services. This allows developers, not just data scientists, to build recommendation algorithms," explains Erskine.

MECCA also deployed a long short-term memory (LSTM) propensity model to identify the best timing for product replenishment. In AB testing, 50% of emails contained personalised product recommendations while 50% did not. Those with personalised recommendations achieved significantly higher conversion rates.

From zero machine learning to 10 million automated recommendations every week

MECCA is now running its personalisation model weekly for all active customers, generating more than 10 million product recommendations every week across all marketing campaigns.

Amazon Personalize is also outperforming MECCA's previous system for implementing product recommendations using native capabilities from its email management system.

According to Shepherd: "We tested the Amazon Personalize recommendations against system recommendations from our email provider. In theory, the recommendations coming through the email platform are also based on purchase history, but they don't take into account as many measures as the Personalize model, making it less effective."

"We have really proven that by showing our customers products that are relevant to their life stage, their journey, and their purchase history, they're far more likely to convert."

MECCA continues to work with AWS and Servian to reimagine the digital experience and delight loyal customers. Its goal is to leverage the power of machine learning to forecast what customers will like, and optimise MECCA's ability to meet demand – all while improving its underlying data set to build ever-more predictive models.

About MECCA

Since 1997, MECCA has helped its customers to look and feel their best by offering the world's best line-up of beauty and skincare brands, coupled with exceptional service and a rapidly growing online business. It employs 4,000 MECCA team members across 100+ retail stores in Australia and New Zealand. Growth is fuelled by opening new stores and harnessing technology to constantly innovate and evolve its concepts, experiences, and service offerings.

Benefits of AWS

- Increased email click-through rates by 65% and email revenue by 76.4%
- Generates more than 10 million product recommendations every week across all marketing campaigns
- Empowers MECCA developers to build product recommendation algorithms using existing customer data – no machine learning expertise required.

Find out more here >

"Since integrating Amazon Personalize, we are seeing a 65% increase in email click-through rates and a 76.4% increase in email revenue relating to the products recommended."

Sam Bain
Chief Digital Officer
MECCA

Solve Geosolutions and DiUS partner to transform the mining industry using artificial intelligence on AWS

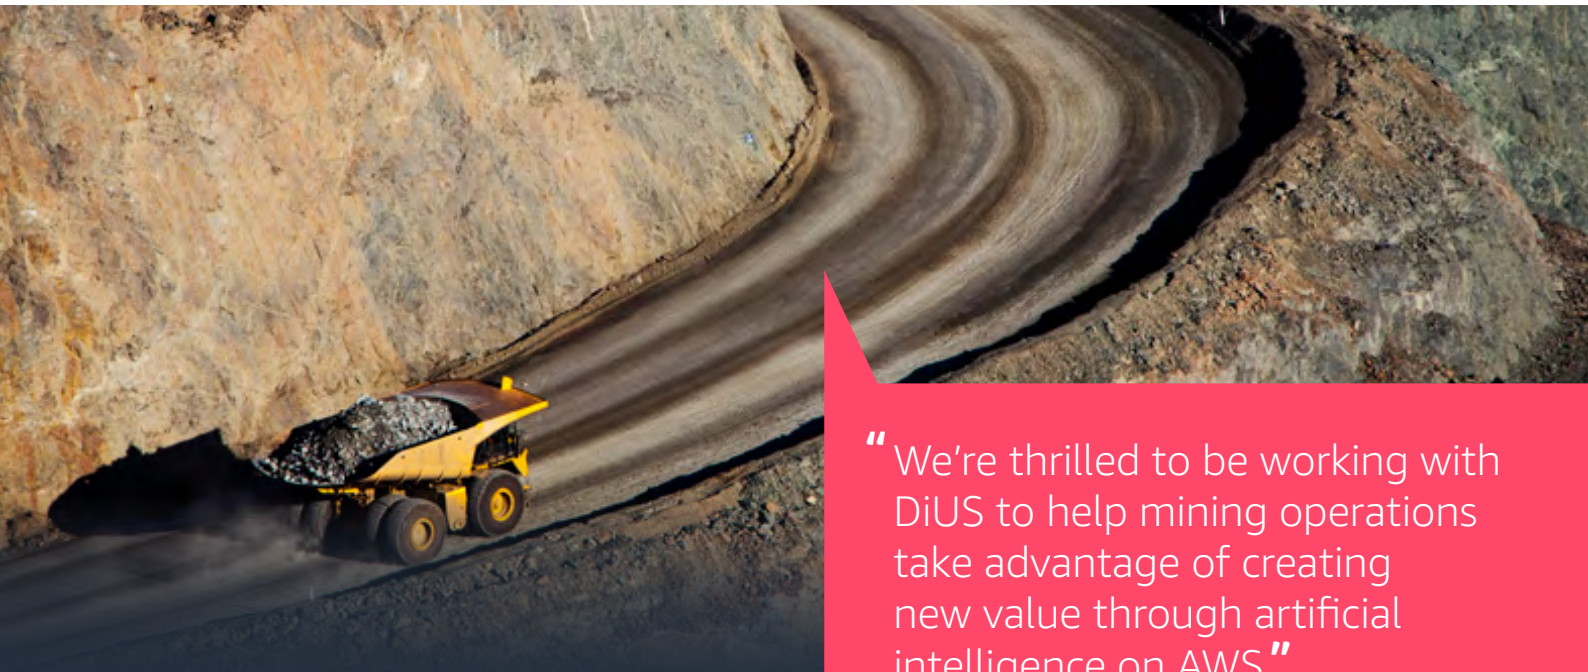

> "We're thrilled to be working with DiUS to help mining operations take advantage of creating new value through artificial intelligence on AWS."
>
> **Liam Webb**
> Director at Solve Geosolutions and
> Chief Executive Officer at DataRock

Helping mining companies limit manual tasks and become data-driven

Improvements in imaging technology enable mining companies to capture, store, and analyse drill-core imagery in innovative ways. Technological advances also help them improve operational efficiencies by using better imagery, collecting many more data points, and revisiting historical data. However, these new capabilities can also require more time and manual resources, particularly from geologists, who must sift through a more significant and complex scope of data. Many mining companies therefore must decide if reinventing themselves in a data-driven world is a worthwhile effort. For the team at Solve Geosolutions, the decision is yes.

Through education, collaboration, and advanced tooling, the company's mission is to help mining companies and geologists take advantage of new data analytics capabilities to drive insights, operational improvements, and, ultimately, pursue new business outcomes Solve Geosolutions seeks to change the way mining companies use technology, such as machine learning, to limit geologists' manual inspection of drill-cores and help them focus on the heart of what they do: using their domain knowledge to understand and interpret geology.

"We have a deep background in machine learning across our team. Our mission is to be at the cutting edge of machine learning to solve geological problems," says Brenton Crawford, director and geoscientist at Solve Geosolutions. "We want to change the way geologists work. Our goal is to help them make the most of the geological data."

Using machine learning and working with DiUS to change the way mining companies approach data

Already an expert in the application of machine learning to solve geological problems, Solve wanted to build an automated and scalable solution to help mining customers tap into increasingly complex sources of data and objectively analyse large repositories of drill-core imagery. The company sought to explore the application of image processing algorithms to high-resolution images and automate the extraction of geological information to deliver more consistent results.

The company sought to use deep learning to help companies generate new insights at scale. However, the team felt they were missing some core skills, like deep learning expertise in cutting-edge technologies, to develop and launch such a model. The team decided to build its solution on Amazon Web Services (AWS) and reached out to AWS for consulting partner recommendations. AWS quickly brought DiUS into the conversation.

DiUS is an APN Advanced Consulting Partner and Machine Learning Competency Partner founded by two individuals passionate about innovation and taking advantage of emerging technologies. "Almost everything DiUS does is cloud-native and on AWS," says Frank Losinno, Client Engagement Principal at DiUS and Chief Technology Officer at DataRock. "We always have the AWS Cloud underpinning how we develop products and solutions." With a strong focus on machine learning, specifically deep-learning analysis, and proven expertise in AWS technology, DiUS complemented Solve's expertise in analysing and understanding exploration and mining data.

The Solve team began working with DiUS to build a smaller-scale proof of concept for its solution. "We turned to DiUS to help us build a model using deep learning to help with complex image processing to extract geological information from core imagery. Anything we asked DiUS to do, they did. It felt anything was possible from the beginning," says Liam Webb, Director at Solve Geosolutions and Chief Executive Officer at DataRock. "Working with a highly professional, well-run, agile, and flexible company like DiUS makes a big difference in our ability to accomplish our desired outcomes."

Creating a solution to streamline image processing and drive faster results for geologists and mining companies

Together, DiUS and Solve developed a solution to apply a range of deep-learning based image analysis techniques to extract geological information from drill-core imagery.

Within weeks, DiUS and Solve built a deep-learning model on AWS to extract a specific geological element requested by a client. The solution used Amazon Simple Storage Service (Amazon S3), a deep neural network architecture called Mask R-CNN, and Amazon Elastic Compute Cloud (Amazon EC2) P3 Instances to accelerate machine learning and high-performance computing applications. DiUS also used AWS Deep Learning AMIs with Open Source Frameworks and implemented its model in TensorFlow, hosted as a docker container on AWS Fargate. The application to prepare the dataset and apply batch predictions was a static react front end served from an Amazon S3 bucket using Amazon CloudFront and Amazon Cognito for authentication.

The results of the segmentation models were really impressive. "We very quickly found we were able to detect important geological features in images at a scale, resolution, and consistency impossible to collect manually by a geologist," Says Webb.

Taking innovation to the next level with DataRock, a SaaS solution on AWS

Following the successful deployment of its models and several new project leads, both companies came together to discuss the possibility of developing a joint venture and business plan to launch a new end-to-end solution, unlike anything currently available, to market. Through their joint venture, the companies could focus on building a commercialised product that could help mining companies take advantage of image segmentation technology to automatically analyse drill-core imagery and deploy machine learning models to help drive more significant insights and better decision-making in the field.

The result of the venture is DataRock, a machine learning-powered, cloud-based image analysis platform specifically designed for exploration and mining. The startup's goal is to help mining companies be more flexible and agile in where they drill, saving time and money to put towards other business-critical projects.

"In the face of increasing industry challenges such as low commodity prices and declining grades, DataRock provides accurate, fast, and consistent high-resolution information about a mineral deposit's geology that aids decision-making throughout the entire mining cycle, delivering important productivity and throughput savings to a mine's bottom line," says Webb. For example, it may have previously taken a company up to six weeks to obtain geo chemical assay results from a mining site. Using DataRock, companies can get proxies within hours. This near real-time power enables more intelligent decision-making regarding the viability of a particular mining site.

DataRock is using many AWS services to power its platform, including Amazon SageMaker to build and deploy its machine learning models, Amazon Simple Queue Service (SQS) to orchestrate the process, AWS Lambda for processing and serving the UI, Amazon API Gateway for the REST API, and Amazon S3 for storage of core image sets. "We are also currently working with the AWS GroundTruth team to collaborate in building out feature sets for our workflow and product needs," says Losinno.

Having recently launched DataRock, the Solve and DiUS teams are just beginning to go to market and educate mining companies around the globe about the benefits of the DataRock solution. "We're thrilled to be working with DiUS to help mining operations take advantage of creating new value through artificial intelligence on AWS," says Webb.

About DiUS

DiUS is an Australian technology services company that helps organisations both large and small to solve hard problems, get new ideas to market, and disrupt and transform business models. The company strives to deliver sustained success by delivering a platform for growth. Its services span the entire product life cycle, from ideation and service design and POCs to creating a minimum viable produce and building a product to scale. DiUS' end-to-to end software and hardware capabilities, coupled with its ability to explore emerging technologies for commercial application, make it unique in the market. DiUS is an APN Advanced Consulting Partner and holds the AWS Machine Learning Competency.

About Solve Geosolutions

Solve Geosolutions is a data science consultancy based in Melbourne. Founded in 2015, the company specialises in statistical analysis and machine learning-based solutions for exploration and mining. Solve's aim is to develop and implement data science-based solutions that improve the way companies mine and explore large and complex multivariate datasets. In addition to designing end-to-end solutions, Solve strives to empower its clients to think about their data in new and innovative ways through collaborative engagement and training.

Find out more here >

Elenium enables safer and smoother airport travel using AWS

"We wouldn't have the intellectual freedom to even think about these ideas if we didn't have AWS to help us take them to market."

Aaron Hornlimann
CEO
Elenium Automation

Innovative technology has extraordinary potential for streamlining passenger flow through airports, enhancing safety, and providing a more enjoyable experience to all. Australian-based startup Elenium Automation (Elenium) is leading the way, using automation to create advanced self-service technology that delivers a connected and seamless traveller experience.

Elenium employed a suite of services from Amazon Web Services (AWS) to create a touchless system that eliminates outdated, confusing, and time-consuming processes and that increases passenger flow. When the COVID-19 pandemic began, Elenium again used AWS to quickly deploy an early-detection health tool to make airport travel even safer.

Rapidly deploying new technology using AWS

Founded in 2016, Elenium aims to help airports and airlines securely modernise their operational technology. With its focus on building sophisticated hardware solutions for airports, Elenium and its team of 80 employees were not interested in building advanced software solutions from scratch and then also having to manage them. The company determined that outsourcing to a managed cloud service provider was the most efficient way to deploy its products. The company turned to AWS because of its breadth of quickly deployable managed services.

In late 2018, Elenium began developing Voyager, a program that uses biometrics, voice recognition, artificial intelligence, and cloud technologies to replace outdated processes – such as scanning boarding passes and manually entering data – in place at many airports. "We developed the first working prototype in four months," says Aaron Hornlimann, Elenium CEO, "and that includes building a whole new set of hardware, bag-drop design, new kiosk design, boarding gate design, and a new virtual duty-free shopping experience, as well as all the AWS-backed software behind that." In March 2019, Elenium introduced Voyager at an exposition in London, essentially setting up a virtual airport in which users could download Voyager onto their mobile phone and interact with it to check in, drop off luggage, shop, and more. The process was simple, but behind it is an integrated array of AWS services working seamlessly with Elenium's hardware.

Turning a selfie into an all-access pass at the airport on AWS

When passengers check in for their flight, they use the Voyager mobile app to take a selfie that is submitted through Amazon Rekognition, which uses highly accurate facial analysis for identity verification. Amazon Rekognition generates an ID from the photo that passengers use throughout their check-in process. The data is housed in Amazon DynamoDB, a key-value and document database that delivers single-digit millisecond performance at any scale, which then creates a map of passengers' progress as they pass from touchpoint to touchpoint – from security to an airport retail shop, for example. From that information, Voyager predicts what data will be needed at the next touchpoint, caching that data to speed up the overall airport experience.

Voyager also uses Amazon Polly, which turns text into lifelike speech, and Amazon Lex, which turns speech into text to create a conversational interface, to provide support and instructions to passengers during the check-in process. Passengers can talk to the app on their mobile device or at Voyager kiosks in airports to check in bags, figure out seat changes, perform automatic visa checks, scan passports, manage bookings, and more.

Another key feature is Voyager Bag Drop, which eliminates the need to use a bag tag by using computer vision to identify passengers' luggage. Once travellers place their bags on the conveyor belt, Voyager sends a confirmation message that the luggage was received along with a photo of it so travellers have a record of its condition before the journey. Passengers can also purchase duty-free products at the Voyager Virtual Shopping Wall – powered by AWS – and have them delivered to their homes or hotels rather than carry them on board.

At the time of boarding, Voyager uses Amazon Rekognition to automatically validate each passenger using biometrics (i.e. facial structure), which is much more secure than the traditional process of printing and scanning boarding passes. "The chance of someone getting a copy of your boarding pass and using it to access your frequent flyer number or anything else that's normally stored in those barcodes is greatly minimised," explains Hornlimann.

Since March 2019, Elenium has implemented various Voyager features to create a connected and seamless traveler experience in thirteen airports worldwide and is currently deploying in another seven. Airports with Voyager technology report that the system is able to reduce processing time by up to 60%. New Zealand's Auckland Airport has installed Voyager kiosks and reduced its average check-in time from twenty minutes to eight minutes. Etihad Airways is now implementing Voyager Bag Drop units and plans to incorporate the entire Voyager process in the next six months, becoming the first airline to deploy the complete contactless self-service technology. This adoption process, says Hornlimann, "is much faster than it has ever been. A check-in kiosk was effectively the same for a decade before we saw any kind of innovation there. But now we're seeing that adoption much faster. And that can only happen because of cloud service providers like AWS."

Enabling safer airport travel during the COVID-19 pandemic

Soon after Voyager's launch, Elenium began to investigate touchless technology as a solution to make self-service more accessible for passengers with reduced mobility. When the COVID-19 pandemic brought new safety concerns to air travel, Elenium quickly realised that Voyager's contactless, voice-controlled technology was especially relevant and useful because it could help people pass more safely through airports without having to touch a single surface, reducing the risk of viral or bacterial transmission. The company also realised the potential for this technology across a wide array of sectors in which public screening may be vital in the future, including health care, aged care, sporting arenas, and more.

In less than eight weeks, Elenium was able to evolve its customer-facing Voyager kiosks into health-screening points that use contactless detection to measure three vital signs: temperature, heart rate, and respiratory rate. If the tool detects an abnormality, it connects the person to a medical professional through a virtual conference or diverts the person to manual screening for a second round of screening. The Elenium tool has been deployed in Abu Dhabi International Airport, and a US trial at airports and healthcare facilities will begin soon. Throughout this process, AWS has helped Elenium innovate to bring confidence to the public, the air travel industry, and other sectors that will be required to transform their practices after the COVID-19 pandemic. "We wouldn't have the intellectual freedom to even think about these ideas if we didn't have AWS to help us take them to market." says Aaron.

Changing the relationship between aviation and technology

Building on AWS, Elenium deployed a touchless self-service system that facilitates a faster, more secure passenger experience in airports. When the COVID-19 pandemic began, Elenium used AWS services to rapidly adjust its focus, adapting self-service kiosks to create health-screening points that make airport travel safer. In the end, what Elenium really accomplished on AWS was changing the aviation industry's relationship with technology.

About Elenium Automation

Australian company Elenium Automation develops innovative technologies that improve the passenger experience and passenger-flow management in airports worldwide. Founded in 2016, it has supplied more than 700 devices to 13 airports globally.

Elenium Automation is recognised as an AWS Travel and Hospitality Competency Partner. Achieving this competency differentiates Elenium as an AWS technology partner that leverages deep industry expertise to help customers transform their business from operational efficiencies to guest experiences. Elenium leverages the wide range of AWS services to build advanced self-service technology to help make the world safer by making various customer-facing devices touchless and evolving them into health-screening points using contactless technology.

Benefits of AWS

- Processes passenger traffic 60% more efficiently in an airport
- Reduced average check-in time by twelve minutes
- Enabled a production cycle of just four months
- Developed and deployed a new tool in less than eight weeks

Find out more here >

> "The advent of cloud computing marks a fundamental shift for the aviation industry"... "It changes how quickly airports can adapt and consume technology."
>
> **Aaron Hornlimann**
> CEO
> Elenium Automation

Transport for NSW uses AWS and machine learning to make real-time decisions

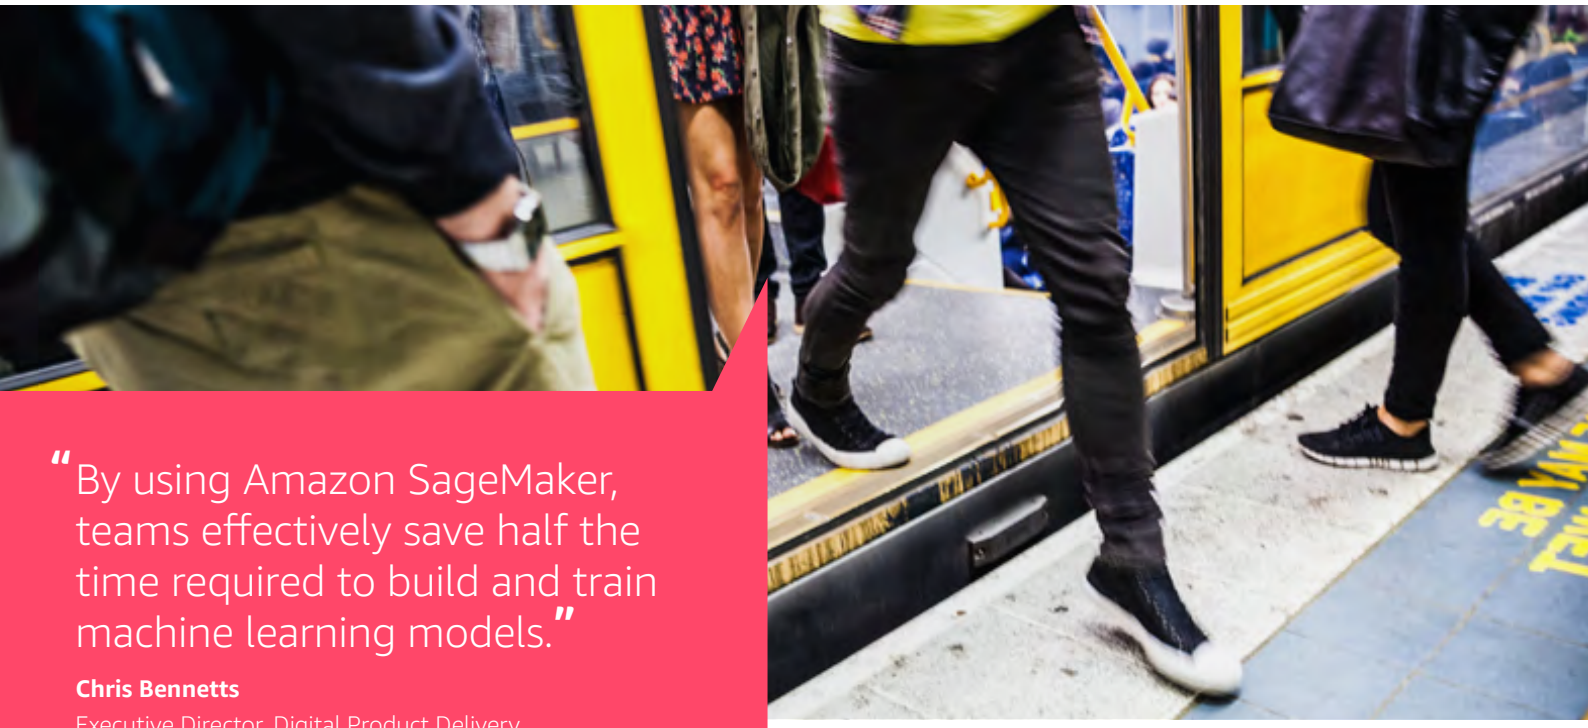

> "By using Amazon SageMaker, teams effectively save half the time required to build and train machine learning models."

Chris Bennetts
Executive Director, Digital Product Delivery
Transport for NSW

Public-private partnerships yield digital innovation

Travellers to Sydney, Australia, have many choices when it comes to public transportation. From ferries to metro lines to buses and light rail, the state of New South Wales (NSW) maintains one of the world's largest public transportation systems in terms of geography, in addition to regional roads. Transport for NSW, the government agency responsible for managing this system, issues Opal RFID transport cards that are used by most passengers for public transport services. Passengers can also tap onto the Opal network using their credit card, debit card, or linked device. From 2018 to 2019, Transport for NSW's train patronage alone comprised 424 million passengers, followed by bus patronage at 349 million.

To improve the customer experience and better connect communities, Transport for NSW continually invests in technology, especially big data processing. Amazon Web Services (AWS) is a partner of the Future Transport Digital Accelerator in Sydney. The Accelerator facilitates direct collaboration between the public and private sectors by connecting teams from the NSW transport cluster with industry partners and startups in the digital space. Several third-party mobile apps have launched from accelerator projects, allowing passengers to check the status of trains, capacity of buses, and more.

Checking the pulse of the network

Transport for NSW has been using AWS for its customer-facing assets since 2014, choosing AWS for the breadth and depth of managed services available. "A lot of our improvements are done by small, four to five people teams who are skilled across different technical disciplines. That's only possible because AWS has invested in opening up these services as part of their infrastructure offering," says Chris Bennetts, executive director of Digital Products Delivery at Transport for NSW.

In October 2018, Transport for NSW began experimenting with machine learning to transition from historically based analytics to a forward-looking model with predictive capability. One of its small teams had been using the Tableau data visualization tool and homegrown algorithms to deliver analytics reports, but results were always two or three days behind. "Executives didn't have the pulse of the network, and the analytics team lacked the latest technology tools to help our operators make better decisions today for tomorrow," explains Mr. Bennetts.

Multiple data sources, including Opal data, feeds from a third-party vendor and takes many hours to ingest—making it difficult to access, process, and analyze datasets at speed. Through the accelerator, the agency worked with Contino, a Premium Consulting Partner in the AWS Partner Network (APN), and Vizalytics Technology, an APN Select Technology Partner, to develop dashboards that deliver near real-time insights and data-driven recommendations. Transport for NSW is now able to predict patronage numbers across the entire transport network, enabling the agency to better plan workforce and asset utilization, and improve customer satisfaction.

Training machine learning models using AWS services

Transport for NSW developed its analytics platform with Contino as a highly scalable, serverless solution built using DevOps practices to maximize business agility. To quickly and easily train a machine model that could predict customer patronage across the network, Transport for NSW ran one year of Opal and weather data through Amazon SageMaker to get predictions for the coming year. Machine learning model training now takes just 35 minutes using the Amazon SageMaker DeepAR forecasting algorithm and Amazon Elastic Compute Cloud (Amazon EC2) C5 Instances, with each transport mode consisting of 168,000 data points. Transport for NSW stores weather data on Amazon DynamoDB, which works in conjunction with Amazon Simple Storage Service (Amazon S3) and Amazon Kinesis Data Firehose to enable the ingestion and transformation of high-velocity data feeds. The agency takes advantage of AWS Lambda to minimize operational requirements while reducing cost.

"The advantage of applications like Amazon SageMaker and Lambda is the ability to quickly spin up a dashboard and build it in a couple of weeks to show our executives," Mr. Bennetts says. "By using Amazon SageMaker, teams effectively save half the time required to build and train machine learning models, by not having to code neural networks." Mr. Bennetts has even built a simple machine learning model on Amazon SageMaker himself, demonstrating the ease of using the tool.

Supporting timely decisions with near real-time dashboards

Thanks to additional dashboards developed with Vizalytics, train controllers benefit from data-driven recommendations on how to mitigate and efficiently recover from service interruptions. "The most important impact of this technology is the confidence it gives our operational teams to make better and more timely decisions that affect customers," Mr. Bennetts says. For example, operators can react immediately to delays on a train line and deploy a supplementary bus service or add a train when usage spikes in a particular location. Such actions have led to improved customer satisfaction.

Because they can predict patronage figures up to 48 hours in advance, and now have access to real-time contextual insight on impending weather events, management can provide more transport services to customers when they need them. The dashboard solution for NSW TrainLink, for example, has eliminated the previous 24 to 48-hour lag in receiving KPI reports. Management can predict how weather fluctuations will impact utilization across each mode of transport and enable better long and short-term planning. In addition, data is easily shared across the enterprise, which facilitates agility and wider involvement in decision making.

Putting analytics at the core

Overall, this machine learning initiative has made analytics a central part of the agency's efforts. "This whole process has given us a lot of confidence around how we need to procure and build tools going forward," Mr. Bennetts says. With machine learning tools now at their fingertips, he adds, teams may choose to build in-house or together with the startup community through the accelerator, instead of issuing big requests for proposals to consortiums of companies.

"We want to shift our way of working to partnering and incrementally delivering capability, using a test-and-learn mindset," Mr. Bennetts says. "In transport, this change has been well received. We have really turbocharged the transport sector globally by using real-time data to improve customer satisfaction."

Taking an incremental approach to machine learning

Reflecting on Transport for NSW's machine learning journey, Mr. Bennetts explains how this step-by-step approach has come to define the agency's path to innovation. "When you're using these emerging technologies, sometimes it's hard to envision what's possible," he says. "You have to iteratively build by capability, and from there comes more questions and you begin to build a backlog of what comes next. This was very true for the evolution of our dashboards."

The agency has since built the Vizalytics Interactive Passenger Report to give a real-time replay of passenger flow through the system, similar to a weather radar loop. Executives are using the report as a tool to help improve customer service. Data feeds from this line will be particularly important because this is the first service of its kind in Australia. Mr. Bennetts concludes, "Our machine learning use cases are fairly straightforward, but even the most basic use cases can be completely transformative to operational and customer teams."

About Transport for NSW

Transport for NSW is the government agency responsible for the Australian state's extensive public transport network, including rail lines, buses, ferries, and more. Its role is to lead the development of a safe, efficient, and integrated transport system to make New South Wales a better place to live, work, and visit.

Benefits of AWS

- Improves customer satisfaction with more efficient transport services
- Facilitates development in small, cross-skilled teams with managed services
- Performs machine learning model training in just 35 minutes with no special skill set required
- Enables operators to make quick decisions using real-time dashboards
- Predicts patronage figures 48 hours in advance with machine learning-driven analytics

Find out more here >

"Machine learning model training now takes just 35 minutes using the Amazon SageMaker DeepAR forecasting algorithm and Amazon Elastic Compute Cloud (Amazon EC2) C5 Instances, with each transport mode consisting of 168,000 data points."

Chris Bennetts

Executive Director, Digital Product Delivery
Transport for NSW

Voice of the Partners

AWS Partner Network

Helping companies build successful Amazon Web Services (AWS) based business practices and solutions

The AWS Partner Network (APN) is the global partner program for technology and consulting businesses that leverages AWS to build solutions and services for customers. The APN helps companies build, market, and sell their AWS offerings by providing valuable business, technical, and marketing support. There are tens of thousands of APN Partners across the globe. More than 90% of Fortune 100 companies and the majority of Fortune 500 companies utilise APN Partner solutions and services.

The Voice of the Partner segment features four ANZ AWS Machine Learning Competency Partners below that share their industry experience and expertise by working with hundreds of customers to deliver machine learning workloads to production. These include, Max Kelsen, Intellify, DiUS and Deloitte. These AWS Machine Learning Competency Partners have demonstrated expertise delivering machine learning solutions on the AWS Cloud. These partners offer a range of services and technologies to help you create intelligent solutions for your business, from enabling data science workflows to enhancing applications with machine intelligence.

Max Kelsen

Machine learning is the 'new normal'

Forging ahead with machine learning

Nick Therkelsen-Terry,
CEO, Max Kelsen

Since its establishment in Brisbane in 2015, Max Kelsen has worked closely with a wide range of organisations in Australia and around the world to improve their competitiveness and make better decisions by applying machine learning.

At Max Kelsen, we believe machine learning is directly linked to the generation of business value, and ultimately, competitive advantage for our customers. If it is deployed with a business value generation problem to solve, machine learning is one of the most powerful economic and commercial tools created in the past century – and it's only the beginning.

The power of machine learning is only now being understood in business and human life sciences environments, particularly in Australia. It can provide insights into behaviour, economically or even at the DNA level, in weeks, versus what would have been months, years, or not at all. For example, Max Kelsen has worked with Domino's to help 800 stores around Australia significantly improve their ability to predict how many staff they will need in two-week blocks – a concept called Cognitive Rostering. Applied to the motor vehicle insurance industry, Max Kelsen's machine learning models have been able to accurately predict at-fault causes of accidents in claims with approximately 99% level of accuracy – a significant improvement from the 60% accuracy generally achieved by human analysis.

In human life sciences research, our teams are collaborating with Professor John Fraser's COVID-19 Critical Care Consortium, using cutting-edge techniques to learn causative relationships in how the disease manifests at the acute stage – potentially saving lives.

> "We must, as a nation, grasp the opportunity to take a global lead in the development of machine learning skills and their application to the major industries which generate economic wealth."
>
> **Nick Therkelsen-Terry**
> CEO
> Max Kelsen

aws machine learning

What's next for machine learning?

The implications of machine learning becoming the new normal are profound, from an economic and societal perspective. Organisations that adopt machine learning and use it in their daily business activities will gain sustainable competitive advantage over those that don't. This applies to most industries as machine learning unlocks both meaning and insight into your customers, into your supply chains, and into your historically complex and messy data, which traditionally you have not been able to make sense of. Prior to the emergence of machine learning, attempting to derive insights from big data would have demanded hundreds of hours of painstaking manual data classification, data entry and of course, analysis, across a wide range of variables. Even worse, that data would often be located in disparate places and be unstructured, making the gaining of insight almost impossible.

Machine learning has changed that. As economies and industries recover from COVID-19, they will look even more closely at how they can create intimate, long-lasting and mutually beneficial relationships with customers; find new ways of efficiency in supply chains to meet new market demands, and in the case of medical research, gain insights from crunching massive amounts of data, faster machine learning will play a key role in all these fields. For example, it can help companies, such as mining companies, improve the efficiency and safety of their operations; it can enable environmental scientists to gain clarity on the movement of marsupial populations after bushfires. The uses of machine learning are unlimited, and the competitive advantages it unlocks are unparalleled.

Taking a global lead in machine learning

We must, as a nation, grasp the opportunity to take a global lead in the development of machine learning skills and their application to the major industries that generate economic wealth.

This is the challenge for business and policy decision makers in the next decade. The emergence of government-backed initiatives such as the Queensland AI Hub, which was formed this year to help create more jobs in machine learning application, is a great start. But we need to do more. Being a consumer of machine learning is one thing. Being a player in its future is what Australia should and can become.

Partners are strategically important

At Max Kelsen, we recently expanded our Amazon Web Services (AWS) relationship, becoming one of only a few Australian organisations to achieve AWS ML Competency status as part of the AWS Partner Network (APN). The company was also named in September as AWS' APN Data, Analytics and Machine Learning Partner of the Year for ANZ. Having a deep, strong relationship with AWS not only provides Max Kelsen with access to AWS' industry leading services, but also to the deep pool of talented people who are developing them. Our relationship with AWS underpins our ability to innovate, differentiate and help solve real world customer problems in record time.

Find out more about Max Kelsen and AWS >

Intellify

Humans of machine learning

Lessons learned on large-scale machine learning transformation programs

Koorosh Lohrasbi
Chief Technology Officer, Intellify

The prediction machine

The 'Prediction Machine' is here to stay, and it will be one of the most potent change agents that businesses have encountered. Thankfully, we are entering the exit ramp of the artificial intelligence hype cycle, and the reality has started to kick in.

The post-hype reality

The reality is that the technical components of machine learning are the easy bits. Thanks to services such as Amazon SageMaker, the barriers of entry to world-class machine learning-tools are lowering every month. However, despite all the advances, it is still tremendously challenging for organisations to adopt prediction machines at scale. A recent Big Data and AI survey by Gartner estimated that around 85% of machine learning projects fail to reach production. Of those that do get successfully implemented, only a minority deliver the value initially envisioned.

We the people

Let's assume that we wish to fundamentally transform a business process, say, insurance claims fraud detection. From a purely technical point of view, there exists significant complexity in bringing together data from multiple systems, understanding that data and building out prediction machines that can accurately flag and triage potential fraudulent claims. Now layer on the complexity of understanding the existing claims process flow, integrating the machine learning system into the organisation and getting these 'domain expert' operators to trust the system.

> "Without a solid executive sponsor who fully recognises the business opportunity and challenges involved in a solution, the chance of success is cut by more than **half.**"
>
> **Koorosh Lohrasbi**
> Chief Technology Officer
> Intellify

All while at the same time potentially sending signals to these operators that their old job could be on the line. Do you see the problem? It is people and culture are what need the majority of investment in the transformation journey. Sadly, we still see many companies just throwing money at a problem, putting all of their bets on technology, algorithms, data scientists and engineers, assuming everything else will fall into place.

So, what causes data science projects to fail so frequently? In a study at Intellify, we reviewed the 100+ projects we have delivered for over 40 customers around the world to find three key drivers play the most significant role in successful adoption.

1. Executive sponsorship and enablement

Without a solid executive sponsor who fully recognises the business opportunity and challenges involved in a solution, the chance of success is cut by more than half.

Investing in machine learning initiatives cannot be compared with a typical transformation project such as changing an enterprise resource planning (ERP) system or cloud migration. The outcomes of these initiatives typically have more initial uncertainty, and that is why machine learning programs need an executive sponsor who has set clear expectations with the business and is willing to experiment.

2. Machine learning is truly a team sport

No one likes silos, yet we keep creating them in our businesses. Why does this happen so often? Because it's easy to put a boundary around a complex domain. Cross-functional teams are a way to cut through these silos. Using small cross-functional, multidisciplinary teams with a clearly articulated goal has a much higher rate of success.

3. Enable people through investing in the right platforms

The final reason we see a lot of great machine learning ideas fail is the lack of a reliable and scalable machine learning platform. Just like other core business applications in the enterprise, lines of business need to rely on machine learning applications built around security, scalability, reliability and performance.

At Intellify, we recognise the vital need for these platforms and we invested significantly in building enterprise-grade machine learning and data analytics platforms, which we provide as IP to shorten our customers path to production in each of our projects. Additionally, we offer the platforms as a managed service to minimise the undifferentiated heavy lifting for our customers.

Intellify continues to grow rapidly as a data science and engineering consulting firm with deep technical expertise. Still, we are under no illusion that real change only occurs when people, culture and process change alongside technology.

Find out more about intellify and AWS >

DiUS

Machine learning at the edge

Insights in adopting machine learning through AWS cloud and edge technology

Marc Phoa, Client Engagement Principal Machine Learning, DIUS

At DiUS we work at the intersection of cutting-edge technology development and practical business applications. As machine learning rapidly accelerates, we have delivered significant savings to businesses through algorithms that reinvent traditional processes and augment human decision-making processes. More recently, our focus has expanded to include edge machine learning, the game-changing combinatorial effect of machine learning and edge computing. Delivering analysis and action via edge machine learning right at the point of interaction has far more transformative potential for how businesses and consumers operate – enabling brand new models of interaction and impact.

Best practice for cloud operations is moving fast towards Hybrid Cloud – one connected footprint to expand reach and differentiate through seamless and consistent customer experiences across all operations. Basically, it is cloud operations everywhere using on-premise, public cloud and edge nodes. In turn, this will power a demand for data analysis and action at the edge.

The advantages of adoption edge machine learning

In edge machine learning, the algorithms are processed locally on a hardware device, using data that is created or collected on the device. It can create actionable insights without a network connection. Delivering analysis and action via edge machine learning enables rich, individualised, experience-based instant interactions that consumers now expect. As edge machine learning advances, it will set the bar for how future interactions work.

> "According to IDC's Worldwide Edge Spending Guide, edge products and services are powering the next wave of digital transformation to create a global edge computing market worth **$250.6 billion** by 2024.[1]"
>
> **Marc Phoa**
> CEP, Machine Learning
> DiUS

Edge machine learning enables real-time responses in situations where milliseconds matter. Consider self-driving vehicles – a well-known example of edge machine learning – and the importance of reaction time. Processing and analysing a fast flow of video and images in the cloud can be impacted by location and internet connectivity, resulting in a reaction lag with potentially devastating results. Processing the same data using the machine learning models on the car's computer, at the edge, the reduction in latency improves response times so that the self-driving car can identify that pedestrian that's stepped into its path quickly enough to take action and avoid hitting them.

1 https://www.idc.com/getdoc.jsp?containerId=IDC_P39947

aws machine learning

Improving the speed and reliability of operational productivity

Fast reaction times are important in many other industrial operations to support worker safety and productivity. For example, creating an algorithm that processes a local video camera stream in real time to alert or automatically shut down a machine when a worker enters an exclusion zone or hazardous area, ensuring compliance with OH&S practices to prevent accidents on site.

Reliability is a significant factor in the suitability of edge machine learning solutions in mining or other industries operating in remote areas. An edge device can continue to operate autonomously with intermittent connectivity or remain effective in a remote location. DiUS is extending its machine learning-powered solution, DataRock, to operate on the edge in order to support remote mining and engineering operations. By enabling the digital photos of geophysical measurements from drill-core trays to be collected and processed using DataRock's machine learning models on the same device, important geological information about the site is delivered swiftly to help efficiently direct mining operations in areas without internet connectivity. For example, it may have previously taken a company up to six weeks to obtain geochemical assay results from a mining site. Using Datarock, companies can get proxies within hours. This near real-time power enables more intelligent decision-making regarding the viability of a particular mining site.

The potential of edge machine learning is almost infinite. In fact, it's so new that businesses and technologists are only just scratching the surface of what it can do and where it could be best applied. At DiUS, we're excited by the possibilities it presents in helping us unlock previously impossible technology solutions for our client's problems.

Find out more about DiUS and AWS >

Deloitte.

Scaling machine learning

Overcoming people, process, and technology challenges for adopting machine learning

William Dodd, Partner,
Analytics and Artificial Intelligence,
Deloitte Consulting

Cloud technologies are redefining how organisations innovate and helping them introduce greater agility across business operations. Moving legacy infrastructure and applications to the cloud as part of these transformations is also accelerating the adoption of artificial intelligence, with a growing number of services like AWS SageMaker now democratising the delivery of intelligent solutions.

While the generalised artificial intelligence often seen in popular culture is still some distance away, organisations are increasingly deploying pragmatic applications across finance, operations and human relations, using machine intelligence to act through automation of repetitive business processes, interact with humans using natural language, and provide understanding from the data they analyse.

"It is clear the widespread use of artificial intelligence by organisations is now becoming a matter of **how** and **when,** not **if.**"

William Dodd, Partner
Analytics and Artificial Intelligence,
Deloitte Consulting

aws machine learning

Addressing the people, process and technology challenges

Organisations are investing heavily in data science and artificial intelligence capabilities to scale these applications, and although this represents a significant asset, often the value realised from these endeavours is underwhelming. As a global firm, Deloitte recognises the huge impact that artificial intelligence can have on businesses and understands that successful adoption is often dependent on overcoming people, process and technology problems that invariably constrain technology adoption.

While model development has traditionally been the domain of data scientists, maintaining, deploying, and monitoring model performance is not – it sits at the juncture of business, data science and traditional technology delivery. Scaling and uplifting data science and machine learning capability requires a robust and repeatable delivery methodology, which drives consistent communication of data science outcomes across the business.

Machine learning operations to scale

ModelOps is a rapidly emerging field at the intersection of artificial intelligence, DevOps, and Data Engineering, seeking to manage the complexity inherent within maintaining both the underlying software and code base of a model and also the data pipelines relied upon to produce results. Building and embedding ModelOps within the delivery of artificial intelligence technologies helps to accelerate delivery of accurate models, and facilitate continuous development, ultimately generating greater business value and relevance.

An Australian telecommunications provider was struggling with a predominantly manual process for classifying and addressing customer complaints. Prior attempts to produce an artificial intelligence-based modelling solution had been unsuccessful because model deployment was often challenging and slow, tool selection wasn't aligned to business processes and performance was not continuously monitored, creating challenges in assessing if the model had degraded and required retraining.

Through implementing a fit-for-purpose ModelOps Workbench powered by AWS, including SageMaker for Automated Feature Engineering and AutoML capabilities, Deloitte and the telecommunications provider were able to establish an artificial intelligence-powered model that auto-classified over two-thirds of customer complaints at greater than 85% accuracy, all delivered from concept to production in under four weeks.

Combining people, process and technology capabilities together with the cloud services that enable repeatable and automated delivery, provides organisations with the ability to realise lasting business value from artificial intelligence as we move into an ever-intelligent future.

Find out more about Deloitte and AWS >

Demystify AI with the power of AWS & Intel

Amazon Web Services (AWS) and Intel® have a 14+ year relationship dedicated to developing, building, and supporting cloud services that are designed to manage cost and complexity, accelerate business outcomes, and scale to meet current and future computing requirements. Intel and AWS aim to drive artificial intelligence based impact for its customer and partners to help solve their business challenges and drive stronger performance and optimisations through groundbreaking innovations built by the power of our combined technologies.

Artificial intelligence and machine learning are helping businesses around the world improve decision-making, enhance customer engagement, and increase their bottom lines. Businesses today have more data than ever before, and that data houses invaluable insights. This treasure trove of data has the potential to accelerate decision-making, improve forecasting, reduce operational overheads, predict equipment failures before they happen, and overall, gain a competitive advantage in an increasingly global marketplace. But how? Artificial intelligence and machine learning give you the tools you need to harness your businesses data and turn it into powerful insights. And whether you realise it or not, you have been using these tools almost every day. Artificial intelligence is enabled by machine learning. Machine learning represents the tools, techniques, and technologies that enable developers to build products or automate processes that include human-like capabilities. A machine learning system can accurately predict or understand something based on data it is given.

Simplify AI with Intel® Artificial Intelligence Technologies and AWS.

Today, implementing artificial intelligence is essential to better serving customers and unlocking new business value. But deploying the infrastructure needed to support artificial intelligence at scale remains a challenge for many enterprises, given the complexity of achieving high performance and accuracy - all while managing costs and utilisation. With Intel® architecture, you can deploy applications into production both quickly and cost-effectively. And with access to Intel's global partner ecosystem, you'll find expertise and solutions to suit your every artificial intelligence need.

From hardware that excels at training massive, unstructured datasets, to extreme low-power silicon for on-device inference, Intel® AI supports cloud, enterprises and research teams with a portfolio of multipurpose, purpose-built, customisable and application-specific hardware that turn model into reality. Intel® has also worked closely with the artificial intelligence ecosystem for a number of years, both optimising and developing a broad range of software tools, frameworks, and libraries that will satisfy the most demanding data science needs – from development to deployment and scaling.

Take advantage of the incredible performance of Intel® technology on AWS. Learn more about Amazon EC2 compute-intensive C5 instances tuned for artificial intelligence workloads as a result of the longstanding collaboration between Intel® and AWS. C5 instances for Amazon EC2 are:

- Ideal for compute-intensive scientific modelling, financial operations, machine learning inference, high performance computing (HPC) and distributed analytics that require high-performance for floating point calculations
- A custom cloud solution based on Intel® Xeon® scalable processors, featuring Intel® Deep Learning boost with Intel® AVX-512 Vector Neural Network Instructions (VNNI)
- Offer up to 96 vCPUs, 192 GBs of memory
- Get started on one of the Available instances: c5.12xlarge, c5.24xlarge, c5.metal (bare-metal) by taking advantage of the AWS free tier

---

Build your organisation's artificial intelligence with AWS and Intel and experience unleashed performance and reduced total cost of ownership for your artificial intelligence systems. Your teams will get access to:

- Intel-based optimisations across all major AWS artificial intelligence services like Amazon SageMaker, DL AMI, Personalize, Comprehend, Rekognition etc. to further accelerate your artificial intelligence journey and reduce costs significantly
- Partner-built models in AWS machine learning Marketplace that can quickly get your organisation started with the most common use cases
- Optimised infrastructure – such as the Amazon EC2 C5 and C5n instances – running on custom Intel® Xeon® Scalable Processors for top performance on compute-intensive machine learning training and inference workloads

---

AWS and Intel® offer a full portfolio of business development and technical training resources, giving your organisation the skills it needs to push artificial intelligence forward with confidence.

Unleash the power of artificial intelligence with the latest Gaudi AI Processor – powering better price performance in upcoming EC2 Instances on AWS.

As demand for artificial intelligence training workloads increases worldwide, so do requirements for greater training efficiencies – for both time- and cost-to-train. The Habana Gaudi accelerator is built from the ground up to specifically address these artificial intelligence training needs. In the 2020 AWS re:Invent keynote, Andy Jassy announced new EC2 instances that will include up to eight Gaudi accelerators, and deliver up to 40% better price performance than current GPU-based EC2 instances for machine learning workloads. Explore top-line view of Gaudi hardware and software architecture and how business leaders will be able to take advantage of it on AWS.

Launch your machine learning journey with AWS

Organisations that have reached the stage of delivering value with machine learning have been on a brave journey, one which involves investment both financial and time into people, processes and technology. To support customers to accelerate this journey, AWS has developed a workflow based on experience from hundreds of customers. It is an incremental approach focused on delivering end-to-end machine learning (ML) solutions into production to deliver early value, with both tangible and intangible results, while building up the skills, operational change management and technical capabilities required.

Click here to start your journey. **Lets get started**