



Empowering employees with AI: A look into how Microsoft uses AI to do business



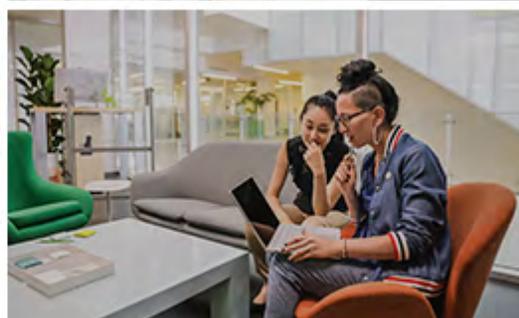
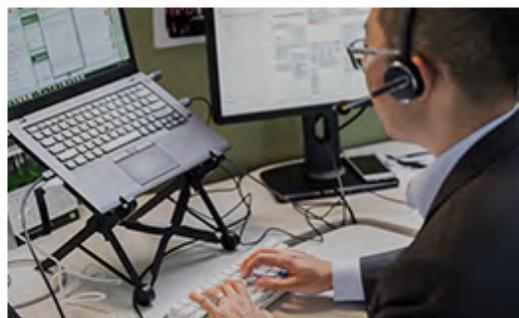
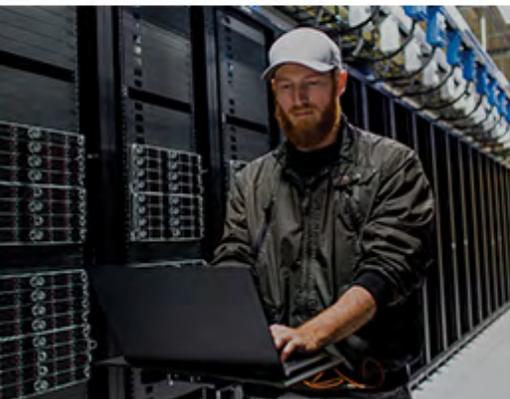


Table of Contents

- 01** Introduction

- 02** Marketing: Intelligent lead scoring

- 03** Sales: Increasing conversion rates

- 04** Customer service: Feedback analysis

- 05** Finance: Late payment prediction

- 06** Finance: Chatbots

- 07** Additional Microsoft AI use cases

- 08** Turning potential into reality

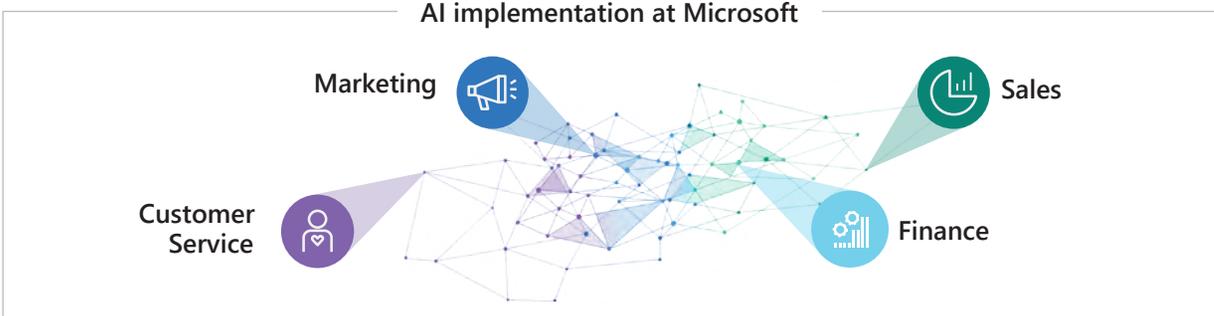
Sharing our experiences and succeeding together

AI presents a tremendous opportunity to drive unprecedented value from data across industries. Like the advent of the internet and cloud computing, AI has the potential to transform internal business processes and products. In fact, Gartner says AI augmentation will create 2.9 trillion USD of business value in 2021.¹

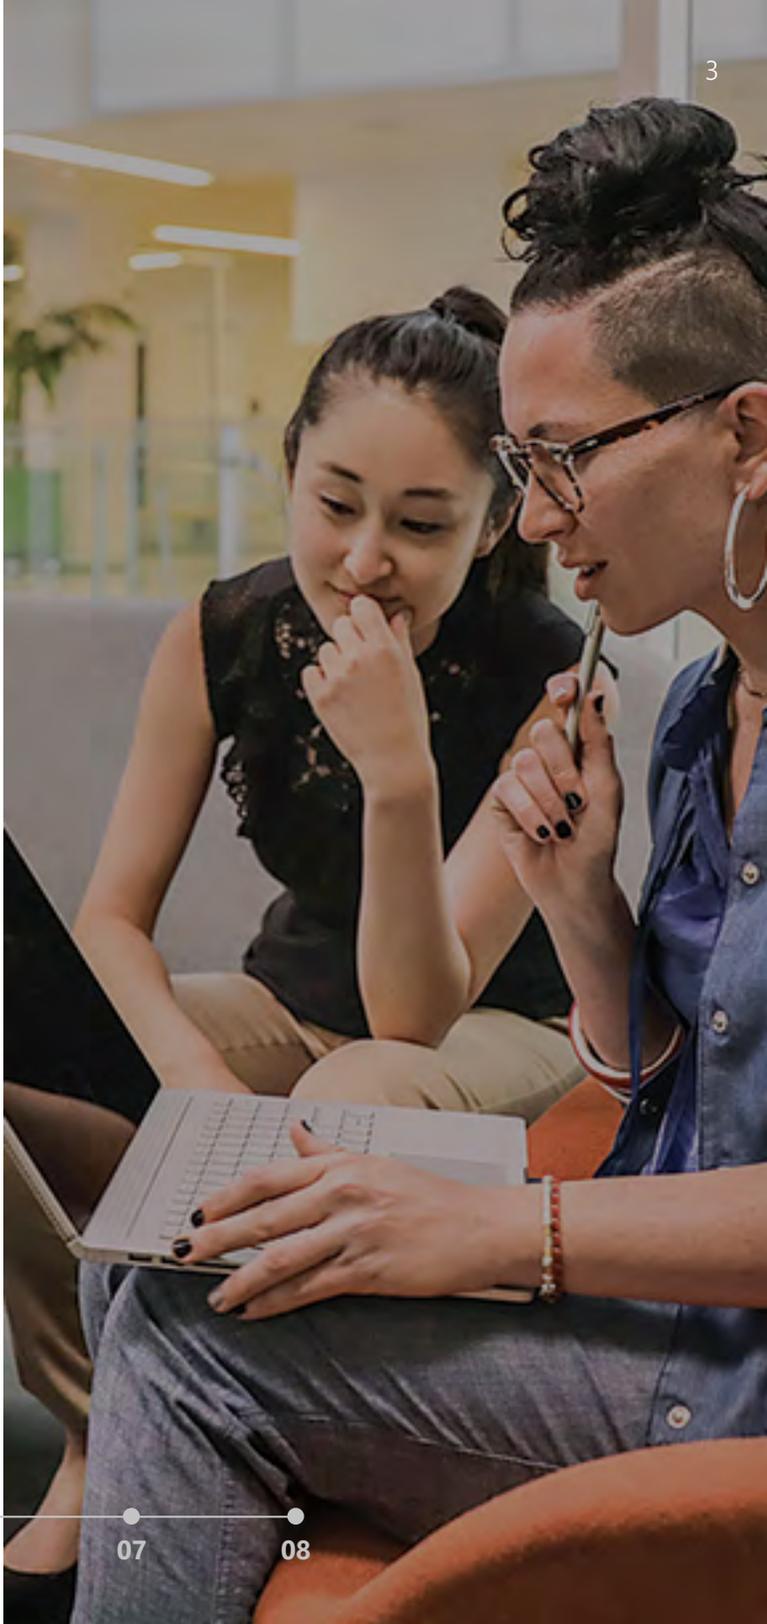
But it's not always easy to translate the potential of AI into tangible benefits. The data that drives AI is the most critical asset modern organisations possess, but its value is directly proportional to the number of people who can access, understand, trust and derive insights from it.

Customers looking to capitalise on AI often ask us to share the ways we are utilising it at Microsoft. In this eBook, we are excited to share our learnings by exploring a few key examples from core business processes familiar to most organisations: sales, marketing, customer services and finance.

Each example is part of a greater holistic effort to foster an AI-ready culture throughout Microsoft. We are working to create a modern data foundation by aggregating clean, connected and authoritative data that any team can use for AI. We're also democratising AI for all of our employees, not just developers or data scientists. We hope that our experiences with AI transformation can help other companies envision their own unique path to success with AI.



1. Gartner Press Release, Gartner Says AI Augmentation Will Create 2.9 Trillion USD of Business Value in 2021, August 5, 2019

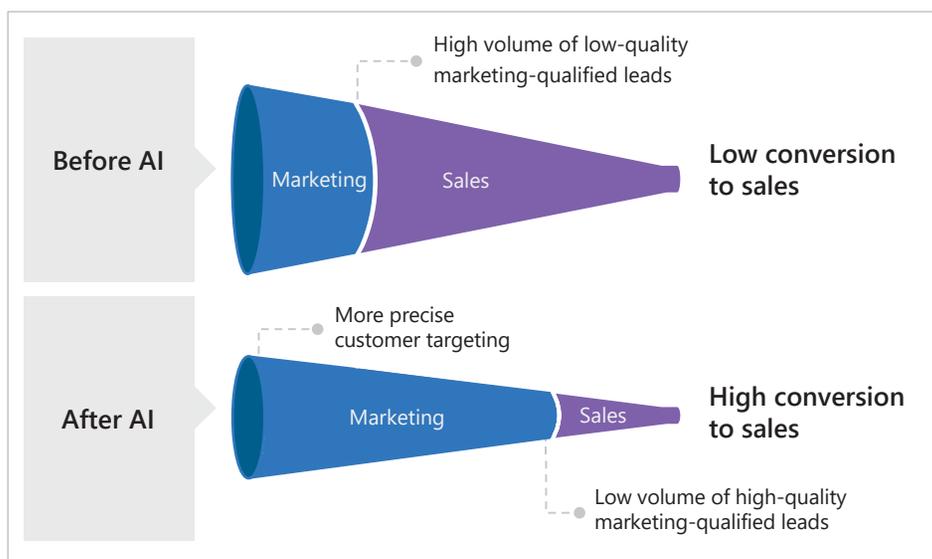


Marketing: Enabling intelligent lead scoring and qualification

Our marketing organisation receives up to 10 million leads per year. While our conversion rates were on par with industry benchmarks, we felt that we could do better.

Given the scale of our customer data, we needed a better way to score leads to reduce the amount of time sellers spend pursuing unproductive ones.

Using AI rather than traditional business rules helped our marketing employees streamline lead qualification and hand off fewer, higher-quality leads to sales teams.



Our AI-based lead scoring platform has helped us more intelligently identify potential customers, ultimately improving conversion rates and marketing ROI.



We combined our marketing employees' understanding of lead quality with the machine learning expertise of our data scientists to create a lead scoring platform. The platform weighs thousands of variables to predict the probability that a lead will convert on any given sales channel. This helps our marketing team send leads to the channel with the best conversion probability and invest more resources in leads with better conversion profiles.



To further qualify leads after they are scored, we created an AI-based lead qualification assistant we call BEAM (Bot-Enabled Augmented Marketing). BEAM emails customers and evaluates their level of interest using natural language processing before we send them to sales. The assistant is very helpful in identifying 'diamonds in the rough', leads that we deemed low propensity but express interest to buy, and 'fool's gold', leads that were scored as hot but were found to be uninterested. We also use the lead qualification assistant to contact expired leads, preventing interested customers from falling through the cracks. Following qualification, marketing hands off the lead to the appropriate sales channel with the entire email conversation, preserving that valuable context.



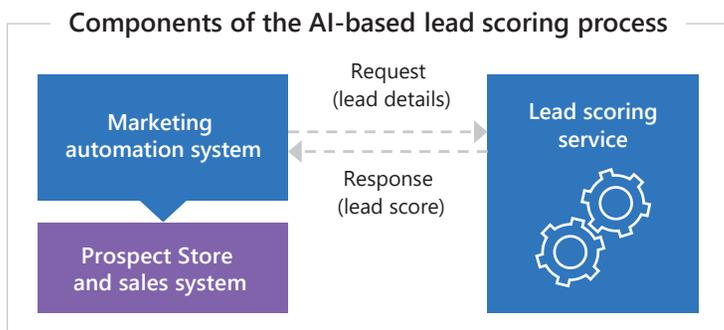
Following these implementations, the conversion rate of marketing-qualified leads to sales-qualified opportunities has more than quadrupled. These solutions were instrumental in helping our sales teams be more successful, thus increasing revenue and improving the customer experience.

AI-based lead scoring and BEAM: Technical background

Our marketing and sales organisation collects leads from content downloads, conferences, webinars, trials, subscriptions and more.

The **lead scoring pipeline** ingests leads from the marketing engine through an API. The system then examines over 200 data points within the context of the lead, including:

- Demographic information, like region or industry
- The lead's role in their organisation
- Prior history with Microsoft, including purchase history and contribution to product reviews and product feedback
- Recent activity like downloading content, responding to a promotional email or attending a Microsoft event



The AI model runs several machine learning algorithms to assign a value to each data point and generate an overall numeric score for the lead. The scored leads are then returned to the marketing engine through the API.

The **BEAM (Bot-Enabled Augmented Marketing)** platform is built on open-source machine learning tools and Microsoft technologies including Microsoft Cognitive Services, Azure Machine Learning and Azure ML Studio. It detects intent and context in customer emails to determine the likelihood that a customer is ready to purchase products or services.

The basic BEAM process consists of the following steps:

1. Send lead from our marketing and sales system an automated email asking how we might help them
2. Ingest email response and clean the data to ensure it's accurate and useable (e.g. convert the email to plain text, run Bing Spell Check and remove private personal information)
3. Use machine learning to detect email context and indicators of the customer's intent to purchase
 - a. Employ several machine learning models: the Language Understanding Intelligent Service (LUIS) from Microsoft Cognitive Services and three models from an open-source tool set. Utilise ensemble methods meta-algorithms to assess each model's discrete result and provide the most accurate results to BEAM
 - b. Leverage natural language processing to perform two assessments of the underlying data:
 - **Intent detection:** classifies intent as hot, not interested, do not contact or none of the above
 - **Context detection:** determines if the lead should continue in the sales pipeline or be redirected (support requests, solicitations, etc.)
4. Place leads back into Dynamics 365 with appropriate intent and scoring information

To learn more, please read our IT Showcase articles [here](#) and [here >](#)

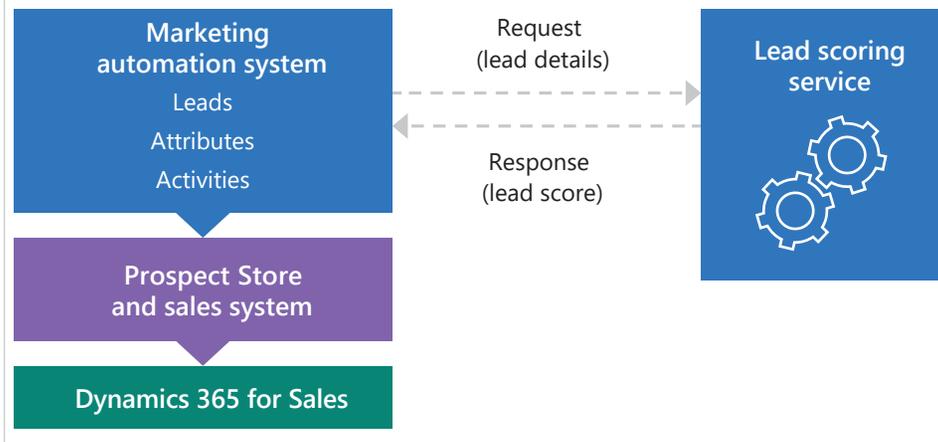
Sales: Increasing conversion rate with Next Best Lead

Improved lead qualification was a great start, but our sales organisation was still receiving almost 10,000 leads every day. We didn't have enough capacity to handle them all in a timely manner, and there were still many that weren't likely to end in a sale.

We wanted to use AI to help sellers on our 'demand response' team prioritise leads. This is the team that handles our traditional call-based business, with customers ranging from managed business accounts to just a single student.

To help sellers pursue the best opportunities, we implemented an AI-based application in Dynamics 365 for Sales called Next Best Lead.

Incorporating Dynamics 365 into the lead scoring process



Next Best Lead empowers sellers with the information they need to prioritise their day's tasks



Upon clicking the 'next best lead' in Dynamics 365 for Sales, the application presents the highest-priority sales lead to sellers.

The application includes the lead score and BEAM conversation history from the marketing organisation so that the seller can understand the full context of the customers' needs before they place a call. Rather than the full list of ranked leads, it uses business rules to provide sellers with a personalised list of leads based on their product focus area, location and language.



Surfacing intelligent insights in our main sales productivity tool creates the best of both worlds for our sellers. Dynamics 365 and AI operate together to give them a view of their prioritised workload based on the highest potential customers in the opportunity pipeline.



Now, sellers have the information they need to maximise their time and revenue. With the new solution, we've increased our conversion rate from 18% to 56%.

Next Best Lead: Technical background

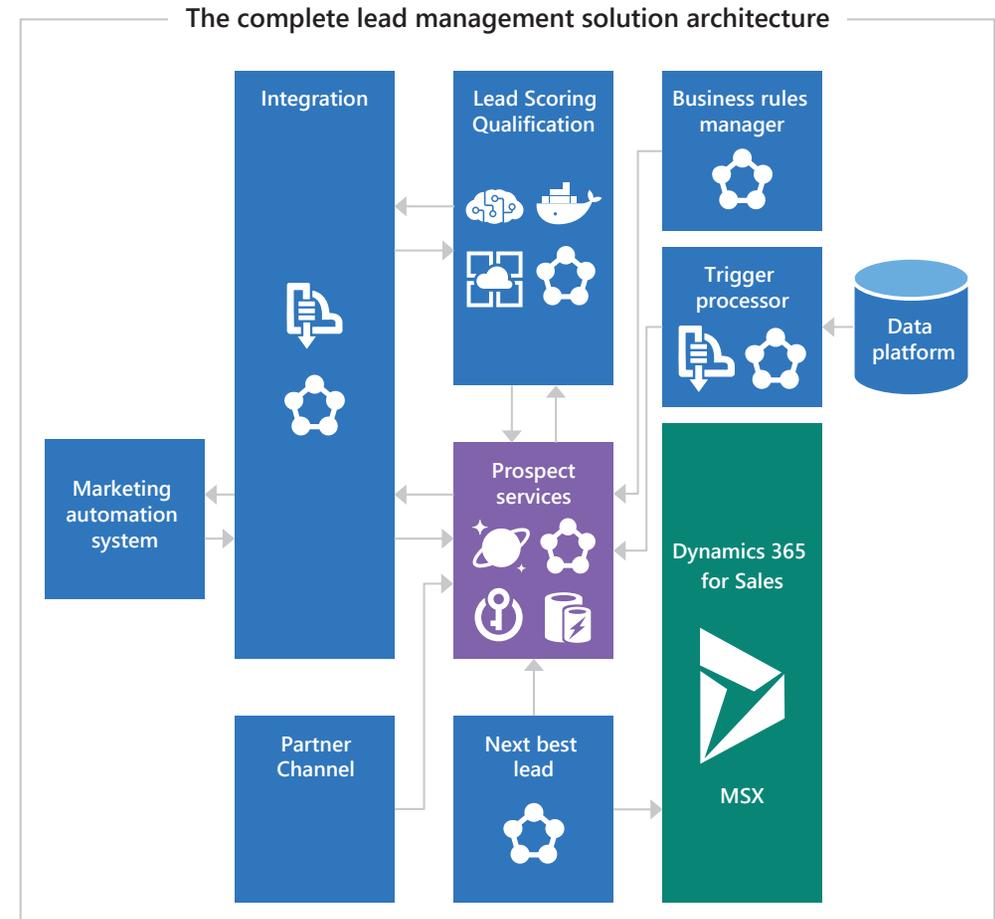
The Prospect Store ingests leads from our marketing system and creates a first in, first out queue containing leads.

The business rules manager directs the scored leads to each seller according to their product focus area, location and language and surfaces those leads through Dynamics 365 for Sales. From there, the seller requests the next best lead by simply clicking 'next best lead'.

The Prospect Store, which allows sellers to customise their settings to pull from only new leads or both new and follow-up leads, creates a lead record for the next best lead in Dynamics 365.

The Prospect Store contains several components that work together:

- **The integration API** interfaces with the internal Prospect Store database and enables insert and update actions for incoming lead data
- **Prospect Services** is built on Azure Service Fabric and provides the core functionality for lead storage and surfacing. Prospect Services includes two components:
 - **The Prospect Store database** is built on Azure Cosmos DB and stores key lead information that the solution manages and surfaces within the Prospect Store
 - **The State service** tracks the state of leads in the Prospect Store to help ensure that only active leads are surfaced in the Prospect Store. Leads can be on hold for various purposes, such as the lead being engaged with another seller for a different Microsoft product or service
- **The Next Best Lead service** surfaces next best leads for sellers in Dynamics 365 for Sales



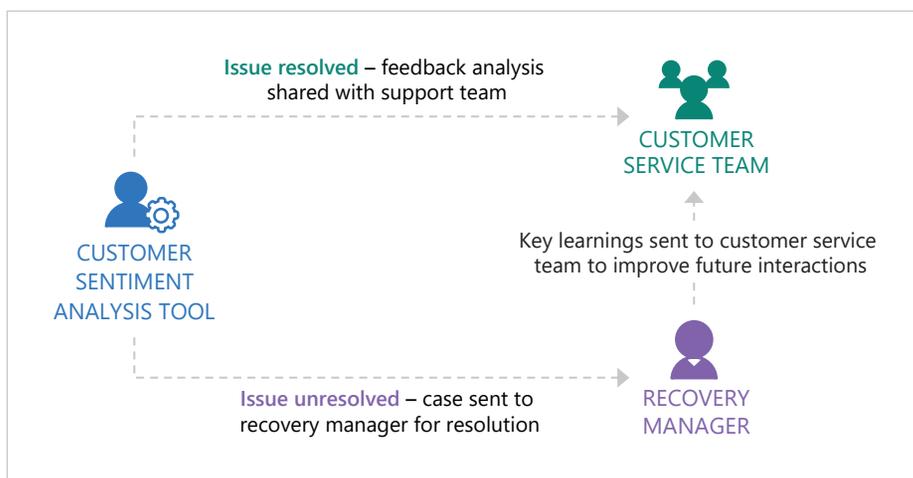
To learn more, please read our [IT Showcase article >](#)

Customer service: Customer feedback analysis

Our customer service organisation relies on customer feedback to help us deliver high-quality experiences. However, as our company has grown, the amount of customer feedback we receive has become far greater than we could process and effectively act on. It would take a human four plus years of work to read a single year's worth of feedback for just a few of our product lines and even longer to act on our customers' input.

To address this challenge, we knew we needed to create a sentiment analysis tool that makes it easier to interpret and act upon post-transaction customer feedback.

By creating an AI-powered customer feedback tool, we are now receiving valuable insights in a fraction of the time and empowering our agents with a new tool they can use to gain relevant information to implement on the frontline.



Our AI-powered customer feedback tool generates valuable insights in a fraction of the time, equipping our agents with the information they need to best serve our customers.



Our Voice of the Customer (VOC) team partnered with data scientists to create a sentiment analysis tool that utilises machine learning to analyse and route feedback. The tool analyses sentiment to identify key factors that drive customer experience and breakpoints where our service didn't live up to expectations.



Our Voice of the Customer team has several processes for acting on customer feedback, including:

1. The **support professional feedback loop** automatically provides weekly sentiment trends and relevant customer feedback to team managers, enabling them to share best practices with their team and address areas of opportunity. Through coaching based on these insights, low-performing agents have received 23% less negative feedback.
2. The **customer recovery feedback loop** is activated when machine learning models detect a lack of issue resolution. The case is automatically routed to the appropriate support delivery team who triages it and engages the customer, as appropriate, to ensure resolution. Our service team then takes the opportunity to learn why the initial interaction did not resolve the issue.



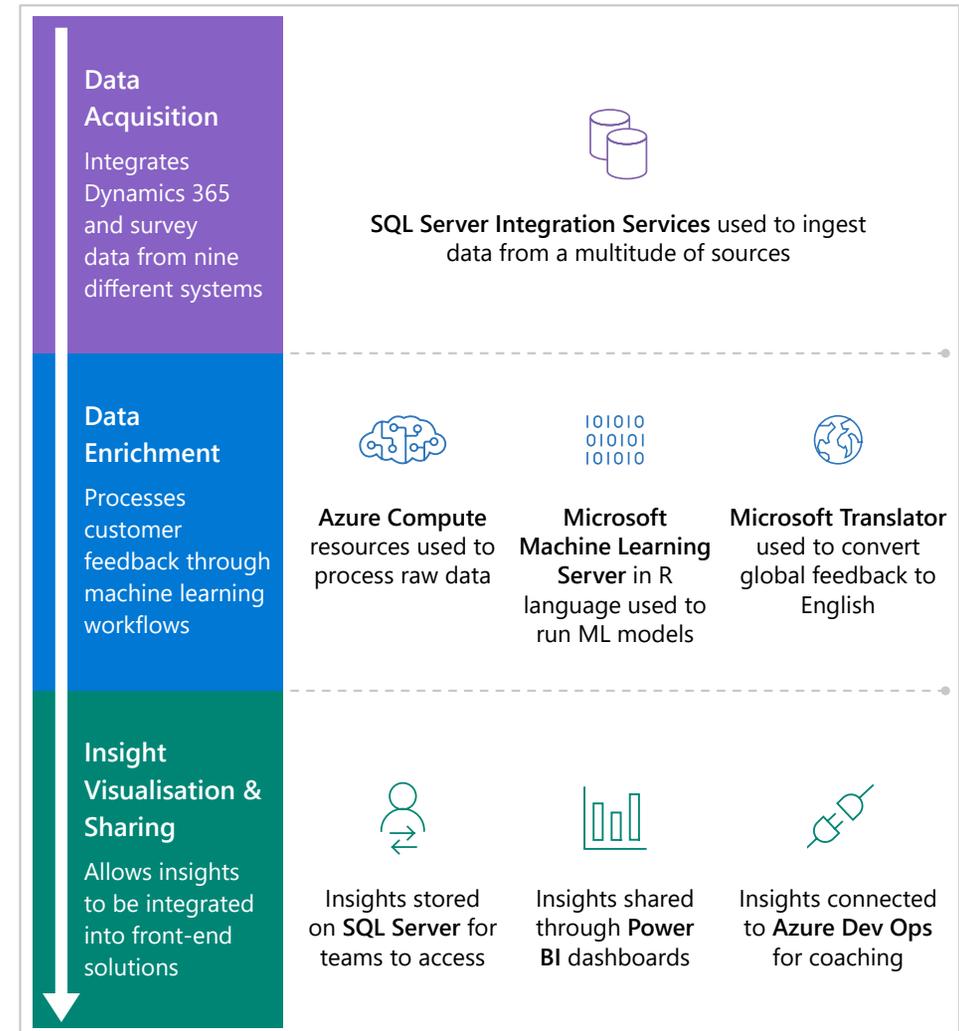
The results have been impressive: our customers are happier, and our agents are able to help them more effectively. For recovered customers, positive sentiment increased on average by 37% and the Customer Satisfaction Score (CSAT) increased by 180% compared to their first interaction with us.

Voice of the customer: Technical background

This AI-based customer feedback analysis process consists of the following steps:

1. Capture and store customer feedback from phone chats, email and IVR surveys
2. Ingest, clean and combine data from nine databases using SQL Server Integration Services (SSIS)
3. Store data on Azure VM
4. Translate all feedback into English using Microsoft Translator, which works for 60 plus languages
5. Run machine learning models over the data using Microsoft Machine Learning Server in R language. We use several different models that are based on N-gram analysis and re-trained every three to six months. These include:
 - a. A model that classifies sentiment as positive, negative or neutral
 - b. A 'breakpoint model' that analyses negative sentiment to identify the cause of dissatisfaction (an agent, a process, a product, etc.)
 - c. A 'resolution model' that identifies if a customer issue is unresolved
6. Store model outputs on SQL server
7. Use Microsoft .NET framework and C# language to pull insights from the SQL server into different front-end solutions including:
 - a. Power BI dashboard: self-service tool refreshed daily that allows users to filter and slice data
 - b. Azure DevOps: tool for tracking the status of a customer recovery
 - c. Email notifications
 - d. Various reports, including a State of the Customer Report that provides detailed information about customer feedback

Finally, these findings are surfaced through our actionability programmes like Customer Recovery Feedback and Support Professional Feedback Loop, which serve as the means by which we act upon insights regularly.



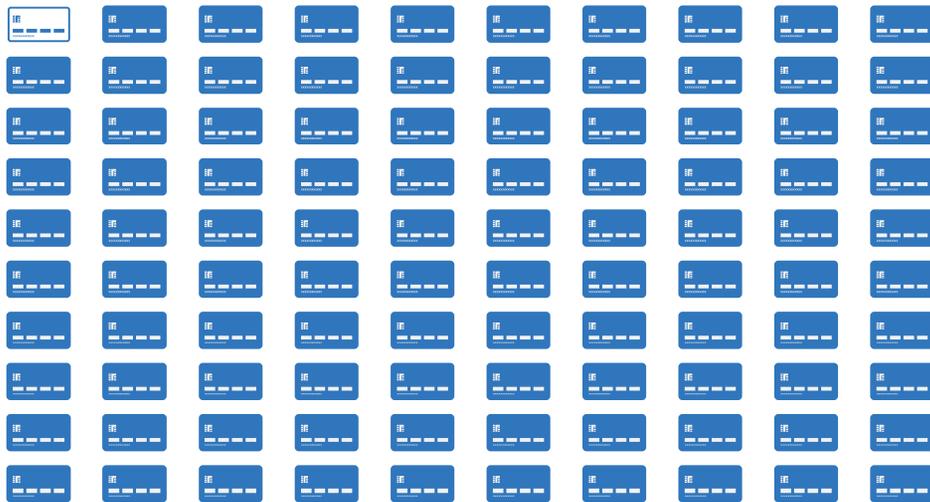
Finance: Predicting late payments

Every year, Microsoft collects more than 100 billion USD in revenue around the world. About 99% of financial transactions between customers and Microsoft involve some form of credit. The company's treasury team manages credit and collections for these transactions. Considering the amount of revenue, you can safely assume that even small improvements in collection efficiency translate to millions of dollars.

Our previous process was to contact 90% of customers with an email reminder about payments.

To reduce that workload and improve the customer experience, we wanted to change that process so that we only contact customers who are likely to pay late.

99% of Microsoft Customer transactions utilise some form of credit



 = Receita DE US\$ 1 bilhão



Our AI-powered late payment prediction tool has reduced the number of customers contacted with payment reminders from 90% to 40%.



The treasury and finance teams partnered with IT to create a model that predicts with over 80% accuracy whether customers are likely to pay late. Managers have a Power BI dashboard with risk scores that indicate the likelihood a customer will pay, which they can use to prioritise actions for their teams. We now only contact about 40% of customers instead of 90%.



In addition to predictions about specific customers, we also get valuable insight into general trends in Power BI, which helps us continuously improve our processes. Examples of trends include:

- Complex invoices are more likely to be late, and contacting customers with complex invoices by phone helps prevent delays
- Some customer types and geographies benefit from phone or face-to-face contact much more than others
- Long-term, high-volume customers and partners are rarely late and can benefit significantly from payment automation

Late payment prediction tool: Technical background

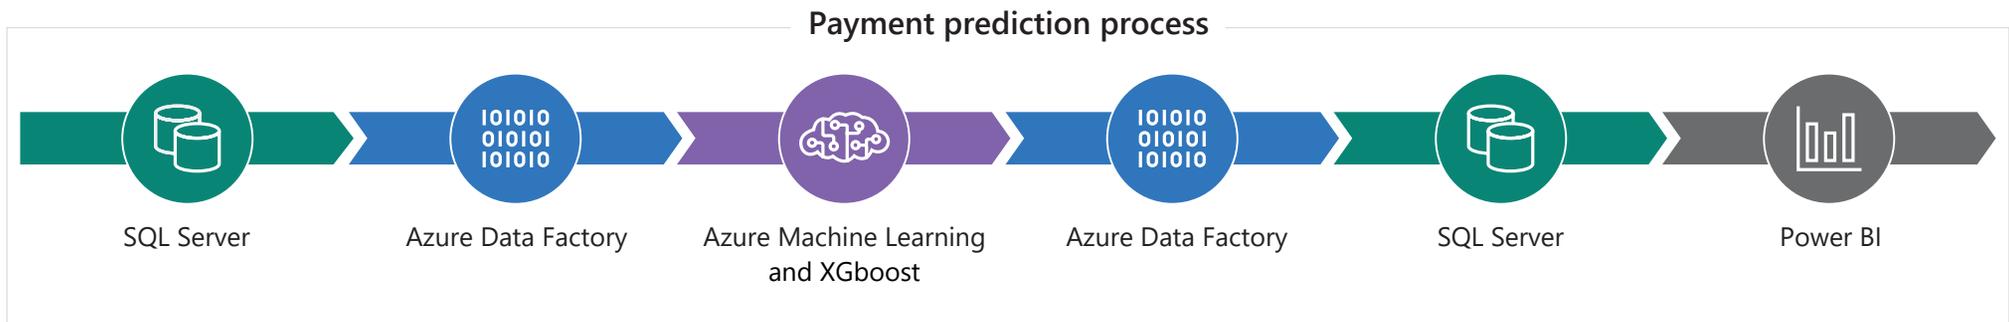
The team used Azure Machine Learning Studio, our cloud-based machine learning service, and a third-party algorithm called XGboost to create the predictive model. To train and continually maintain the model, our process works as follows:

1. To improve the success of the model, we segment our data by variables in a process called feature engineering. Data is drawn from our internal SQL Server data warehouse, which stores 800 gigabytes of data from SAP, Microsoft Dynamics 365, our internal sales system, our credit-management tool and external credit bureaus. In this case, our features included customer type, geography, tenure, invoice amount, products purchased, purchase frequency, whether a due date has been extended, whether there's ever been a billing dispute, and more
2. On an ongoing basis, our engineers identify what data we'll use and then build a pipeline with data from the SQL Server data warehouse to enable the predictive model. Then they provide the data to the data science team
3. Next, the data scientists train the model using Azure Machine Learning Studio. They connect the data to the eXtreme gradient boosting (XGBoost) algorithm, which creates decision trees that produce one of two outcomes:
 - Late = 0 (a customer pays on time)
 - Late = 1 (a customer pays late)

We have more than 1,000 trees, and the largest tree has 100 levels. Azure Machine Learning also gives us a risk percentage score of how likely the customer is to pay on time.

We use Azure Data Factory for moving data between the SQL Server and Azure Machine Learning

4. The scores go into our SQL Server database and are displayed in [Power BI reports](#) to collections teams



To learn more, please read our [IT Showcase article](#) >

Finance: Leveraging chatbots to free up employee time

Like in many businesses, employees from across the company have questions that finance needs to answer, many of which are standard and repetitive.

This bottleneck becomes particularly problematic at the end of each quarter, when our sales force tends to close a lot of deals. During these peak times, our finance team often spends thousands of hours answering questions.

Because most of the queries our finance team receives are routine, this was a prime opportunity to implement AI.

Credit and Collections Chatbot

Today

What is the credit status for Contoso and Fabrikam?


 Joe Smith



Customer name	Number	Country	On hold?	Overdue amount
Contoso	3292889	USA	No	\$0
Fabrikam	1926483	USA	Yes	\$23,500



➤



Intelligent chatbots have saved the finance team hundreds of hours per month and significantly reduced internal support costs.



To meet the demands of our internal teams while keeping costs low, our team has developed and deployed intelligent chatbots to answer routine questions:

1. We created a credit and collections chatbot to enable field sales, operations and collectors to quickly and easily access the latest information about their customers, including credit limit or any overdue balances. As a result, we've eliminated finance as a bottleneck for sales teams and accelerated sales cycles
2. In another scenario, our finance operations group trained a bot called the 'Finance Digital Assistant' to help answer the half-a-billion annual queries about financial systems or processes from Microsoft employees, vendors and partners. This chatbot has sentiment analysis capabilities, enabling us to identify and triage high-risk interactions and prioritise interactions for white-glove service. Recently we enabled Bing translation in the chatbot, so it works seamlessly in a wide range of languages, crucial for a global organisation. The Finance Digital Assistant consolidated 16 discrete finance apps into a single unified interface, creating a much easier experience for our employees. Since implementation, we've seen a massive reduction in support tickets

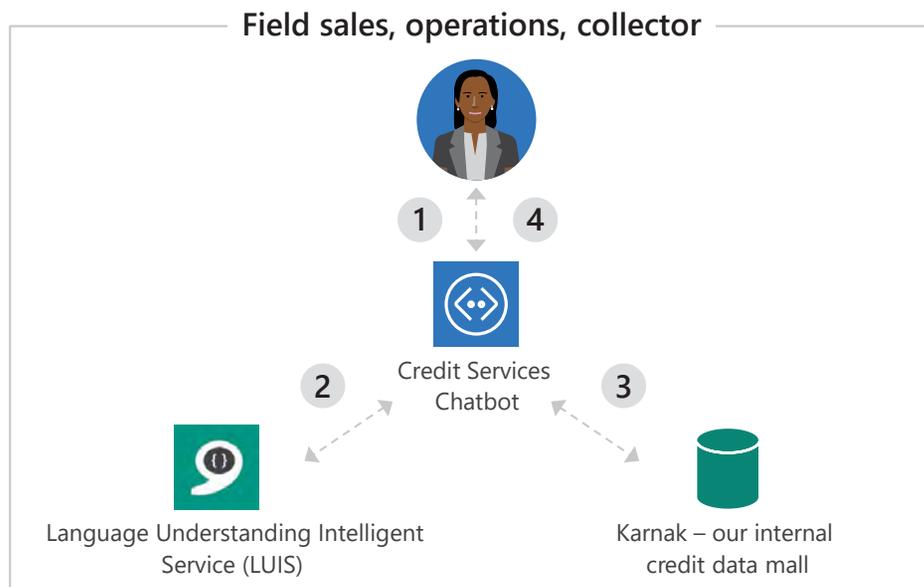


The thousands of hours saved through intelligent automation have enabled our finance team to focus on projects that rely on human intelligence, such as partnering with sales teams to find new revenue opportunities.

Credit and collection chatbot: Technical background

We built our credit and collections chatbot with [Microsoft's Bot Framework](#) and [Azure App Service](#). It works as follows:

1. The user asks a question to the chatbot in plain English
2. The chatbot uses the Language Understanding Intelligent Service (LUIS) to translate the question from plain English to a computer-understandable language
3. The chatbot asks a question to the Azure App Service that connects to our credit and collections data repository
4. The chatbot formats and presents an answer to the user



To learn more, please read our [IT Showcase article >](#)

Finance Digital Assistant: Technical background

The Finance Digital Assistant was built using eight different Microsoft services. It monitors user activity and proactively offers guidance when deemed appropriate.

If the support bot can't quickly resolve the issue, it will shift to a live agent who can interact with the user within the same support window. The live agent has the complete context of the user's query and can start the conversation when the bot fails to answer. A corresponding support ticket is auto-created for the user-bot conversation, allowing the agent to perform offline follow-up or post-call quality analysis.

There were a handful of design considerations we prioritised for the Digital Assistant:

- **Telemetry and analytics:** chatbots log telemetry to collect metrics automatically and unobtrusively. Telemetry will help answer some of the important questions around how our users interact and what options are commonly selected
- **Proactive messaging:** using machine learning models, we trained our bot to understand the most common interactions and present answers proactively
- **Multilanguage support:** using Azure Cognitive Services, the bot detects the user's language and provides answers in that language. When a live agent (who is usually English-speaking) is needed, we use the same solution to present the question in English to the live agent and convert the answer to the user's preferred language

To learn more, please read our [IT Showcase article >](#)

AI drives agility and scale in our products



Bing

In our new deep learning-first Bing stack, we index web documents [using neural networks to improve quality and relevance of search results.](#)

Azure

To optimise the availability and performance of our cloud services, we implemented AI Platform Manager. This set of AI tools uses prediction and big data analysis to mitigate limitations in technologies, enabling huge cost savings.

Bing Ads

We use deep learning across Bing Ads in use cases such as predicting user response and improving how we match ads to queries. Deep learning has increased our ad revenue dramatically and helped us deliver more high-quality clicks for advertisers.

AI powers other vital processes across Microsoft



Smart buildings

We leverage data analytics, IoT and Azure Machine Learning for predictive maintenance, climate control and HVAC optimisation – [keeping our buildings comfortable while minimising our environmental footprint.](#)



Supply chain

We are mining real-time weather data from social media and mapping it to the supply chain footprint (locations of factories and shipping lanes, and where the product is going) [to determine whether there are any potential disruptions to supply chain operations.](#)



HR

We aim to transform HR by [providing employees with what they need in the time it would typically take them to ask for it.](#) For instance, we've created a bot using our [Language Understanding Intelligent Service \(LUIS\)](#) that automates the creation of the 5000–7000 unique travel letters each month.

Turning potential into reality

We believe that every organisation stands to benefit from transforming their business processes with AI. Through our own transformation, we've learnt that effective cross-functional, multidisciplinary collaboration is key to long-term success. By engaging data scientists, analysts and business strategists across typical company lines, we've realised the power of innovation fostered within diverse and inclusive environments.

We offer our experiences with AI not as a one-size-fits-all solution, but as something that our customers can learn from and build on. No two organisations are built or run the same, so before jumping in and adopting the technology, business leaders should think critically about their own organisation's level of AI maturity and about the cultural shifts that may be necessary to make AI initiatives a success. To uncover insights about your organisation's AI maturity, take our [AI Readiness Assessment](#).

To learn more about Microsoft's transformation with AI and consider how you can use it to guide your own initiatives, check out our [AI Business School](#). The AI Business School is a master class series featuring industry experts, case studies, practical guides and tech talks designed specifically to help business leaders address key challenges around implementing an AI-driven approach.

Finally, to take a deeper dive into the technology behind our AI solutions, visit Microsoft.com/itshowcase. The IT Showcase site offers a variety of technical resources, including case studies, white papers and videos.

© 2020 Microsoft Corporation. All rights reserved. This document is provided 'as-is'. Information and views expressed in this document, including URL and other Internet website references, may change without notice. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.