



Persistent

Maximizing Business Value from Data

Persistent Data and Analytics: Preferred Data and AI
Partner for Enterprises worldwide



When it comes to data and analytics, Persistent has a deep heritage. The name Persistent was inspired by working on systems that were persistent, which are databases. Persistent’s expertise in digital engineering and decades of experience across diverse industries enable us to deliver an exceptional portfolio of advanced data and AI offerings.

Our strong team of expert data practitioners, continuous investments in technology, robust partnerships in the data ecosystem and commitment to delivering business outcomes make us a preferred data and analytics partner for global clients across diverse industries, such as Banking, Financial Services, Insurance, Healthcare, Life Sciences, Software and Hi-Tech. We help clients anticipate market changes and adapt quickly by leveraging data to introduce new products and services, optimize costs and streamline processes. Our purpose-built business solutions and AI-based accelerators such as Automated Reporting and Analytics Solution (AURA) deliver results 3x faster than manual methods.

Persistent has been recognized as a Leader in the Everest Group’s Data and Analytics (D&A) Services for Mid-market Enterprises PEAK Matrix® 2023. We have also been named a Leader in the Analytics Services 2023 ISG Provider Lens™ Study. It is an apt validation of our strong capabilities, innovation and investment, vision and strategy, and market adoption in delivering Data and Analytics services to our clients.

This document provides examples of how we transform our clients by helping them maximize business value from their data.


persistent.com/data-and-analytics →

Banking & Financial Services 4

Healthcare & Life Sciences 18

Telecom , CPG, Food & Beverages 38

Banking & Financial Services



The Challenge

To compete in a market where FinTechs offer easy credit and financing options, the client needed to accelerate the mortgage loan process by leveraging automation. They wanted to enable real-time mortgage data and self-service analytics for brokers, underwriters, and lenders.

US mortgage leader accelerates loan processing by 50% through data stack modernization

The client, a Fortune 500 company, is a leading provider of automation software for the US mortgage industry.

They provide a digital mortgage process to help residential loan providers, title agents, homeowners, and mortgage investors close loans faster and increase profit margins.

Banking

Data & Analytics



The Solution

Persistent proposed a solution for data stack modernization using AWS and ensured data monetization through the Insights Platform. This enabled the client to aggregate data from all lenders and scrub it for data privacy. This data could then be presented to each client through a series of analytics, benchmarking, and peer insights targeted for final consumption by business leaders.

The Persistent solution also generated pre-built reports to identify sub-prime

and near-prime lenders across counties. It enabled lenders to use the solution to innovate and roll out competitive loan products.

As the mortgage industry is highly regulated, compliance mandates documents to be shared three days in advance. Persistent deployed a document classification solution that allowed the client to automate the whole process of mortgage documentation depending on the criticality and requirements.

The Outcome

By migrating 60+ servers to the cloud, Persistent ensured 100% regulatory compliance with zero downtime. The client transformed its mortgage automation software and decreased document signing time from weeks to hours. This accelerated loan processing by 50%.

AWS S3

Amazon Kinesis


Snowflake

Enabling 75% faster reporting and cost savings for a major bank merger

The client is one of the leading banks in the US offering banking, investment, and mortgage services and products to over 70 million customers worldwide. With approximately \$1.9 trillion in assets, the bank operates in 35 countries and has a rich history of mergers and acquisitions.

Banking

Data & Analytics



The Challenge

Following a successful merger, the client needed data validation and migration support. They wanted a centralized platform that would ensure a single source of truth. This was critical as the lack of a central data and security framework led to duplication, runaway costs, and data integrity issues.

Moreover, as the data was stored across 100+ disparate banking systems, including credit, deposit, trade services, and treasury management, several business units struggled with report creation.

The Solution

Persistent built an end-to-end data integration platform (iHub), that assimilated data from heterogenous systems including Oracle, SQL Server, and MySQL into a central location. The platform ensured common structural and data standards for each banking system and used a proprietary data model. It also provided data for finance, risk, and regulatory reporting.

The solution strengthened risk data aggregation capabilities and ensured compliance with Basel Committee on Banking Supervision's regulation number 239 (BCBS 239). This enhanced risk management and strategic decision-making processes, resulting in a distinct competitive advantage for the bank.

The Outcome

The solution has harmonized 1.8 trillion data elements distributed across several disparate systems. With 75% faster data mapping and reporting, the bank was able to realize a Total Cost of Ownership saving of \$1.8 million.



80% faster turnaround time and a 75% improvement in response accuracy, improves customer satisfaction for a SaaS tax provider

The Challenge

Our client offers online tax preparation solutions allowing American consumers and tax professionals to electronically file state and federal returns. While this makes life easier for taxpayers, it also presents multiple challenges for the firm. Tax preparation is, by nature, a cyclical business, driven by the annual April filing deadline, as customer queries spike during the peak tax preparation period, leading to scalability issues.

Each year, the firm hires customer service contractors to keep up with surging demand. This leads to knowledge-sharing difficulties; new hires face a steep learning curve, with only limited knowledge repositories and few in-house experts to turn to for assistance. Predictably, response times for customer queries suffer. The CTO and CIO agreed to deploy an AI-powered support solution to improve their customers' experience, reduce costs, and help drive company growth.

The Solution

Persistent crafted a three-phase solution to answer immediate challenges while putting the client on a solid long-term footing. We began by setting up a central knowledge repository to help representatives answer customer queries and, integrate information from multiple repositories that were previously separate and siloed. We developed a conversational AI teams BOT — an intuitive interface for obtaining summarized answers to complex customer queries. Finally, we added an automated data pipeline, i.e., a real-time knowledge feed to ensure

that responses to customer queries were accurate and up to date.

We leveraged Microsoft Cloud, AI, and Generative AI technologies throughout the solution build. This included Azure stack for scalability and security, Azure Open AI models for knowledge capture and integration, and Azure Cognitive Search and BOT Service for the application front-end. The integration of Azure cloud and Generative AI capabilities resulted in a faster and less costly development cycle.

The client is a provider of online tax preparation solutions.

Financial Services

Gen AI



The Outcome

The client achieved 80% faster turnaround time and a 75% improvement in response accuracy. This led to a significantly improved customer experience and ensured scalability. It also provided a 40% saving in FTE cost for the client.

Azure Stack

Azure Open AI Models

Azure Cognitive Search

Azure BOT Service

The Challenge

As customers increasingly demand tailored financial services, retail banks must offer hyper-personalized product offerings at scale. Our client — a successful global Fintech — wanted to differentiate themselves from competitors by understanding customer spending habits and preferences and crafting personalized product and service recommendations. The company wanted a digital solution that would generate personalized marketing and product offers quickly and at scale. Technically, this called for a generative AI-based solution capable of identifying new product recommendation opportunities based on customer behaviors.

Creating explainable and personalized product recommendations at scale for a Fintech leader

The client is a global Fintech firm offering hyper-personalized financial services.

Financial Services

Gen AI



The Solution

Persistent developed a three-part solution designed to understand the client's customers and target them intelligently for increased business. First, we built a large language model (LLM) profile generation engine that produces natural language reports on user spending habits by analyzing transaction data. Second, we added a product recommendation engine, leveraging LLM to generate personalized offers. Finally, we harnessed the power of the LLM to generate and edit persuasive, personalized emails to targeted customers.

AWS cloud and the AWS product set were key to our development effort. The team employed Amazon SageMaker to quickly and easily build and train machine learning models, and then directly deploy them into a production-ready cloud environment. We used AI21 Jurassic 2 Mid, a pre-trained LLM by AI21 Labs that can tackle complex tasks such as question response, summarization, copy generation, and advanced information extraction. We also leveraged some of the ready-to-use models in AWS to accelerate the development effort.

The Outcome

Persistent's solution enabled the client to identify patterns from customer transaction data, generate profile summaries, and then use those summaries for personalized marketing. All the outputs from the

solution are "explainable recommendations," that is, LLM-generated personalized recommendations and marketing content are based on known, observable customer behaviors.

Jurassic 2 Mid

AWS Jumpstart

AWS Sagemaker and AI21

React.js

Fine-tuned

LLAMA2-7B

LangChain

Pandas

Python

PE firm accelerates investment decision-making using Gen AI to reduce deal approval time from 16 hours to 30 minutes



The client is a private equity firm specializing in software, data, and technology-enabled services companies, with a focus on driving growth and innovation

Private Equity

Gen AI

The Challenge

The managing directors of the PE firm, who are responsible for identifying private equity investment opportunities, had to spend many hours creating memos seeking approval for furthering PE investment deals.

A considerable amount of time and effort was invested in the manual identification of risks, challenges, and constraints and in answering follow-up questions. The client needed a solution that would retain every memo creator's writing style and allow for collaboration with stakeholders for reviews.

The Solution

The Persistent solution creates automated memos in user-specific personalized styles by extracting data points from multiple sources. It allows users to select the sources of information such as confidential information packs, meeting transcripts / notes and other data systems. Gen AI is then used to derive facts and figures from these. Gen AI also uses past memos of a specific user to take the derived facts and figures and produce a coherent memo.

The Outcome

The client experienced a 90% reduction in memo creation turnaround time from 16 hours to 30 minutes. The automated solution increased collaboration, enhancing memo creation with a personalized writing style reflected in the output document.

Open Search Vector Db

AWS Claude V2

Langchain framework

AWS Sagemaker

React

Improving productivity with a Generative AI-driven employee collaboration platform for a private equity firm

The client is a private equity firm that invests in mid-market software, data, and analytics companies.

Private Equity

Gen AI

The Challenge

In the complex world of private equity, insights emerge from a team approach, with analysts and decision-makers sharing perspectives on business opportunities, best practices, guidelines, and operational and research reports. Our client — a leading private equity investor in mid-market software, data, and analytics companies — lacked a collaboration platform for internal interactions among its employees and key personnel of its subsidiary companies. One consequence was a lengthy lead time in deriving business insights from high volumes of business data. Further, the company needed to share operational reports with senior executives of its subsidiary companies, i.e., a “role-based” discussion board. The Chief Technology Officer decided that an AI-driven collaboration platform might be the answer.

The Solution

Persistent responded by developing a next-generation employee collaboration and knowledge management app, leveraging the power of various Microsoft technologies. The app delivers actionable business insights through an efficient and intelligent chatbot that accesses AI to analyze structured and unstructured data. Our team selected Microsoft SharePoint and Microsoft Power Apps as the collaboration platform to enable the creation of an intranet to manage and share files with company employees.

The chatbot app is powered by Microsoft Azure Open AI and Azure Cognitive Search used together for intelligent processing and cataloging of large volumes of business information. Azure CognitiveSearch is a fully managed cloud search service that provides a better user experience for customers, offering key phrase extraction, optical character recognition (OCR), and role-based access control (RBAC) that our client required.

Persistent also deployed Azure Application Insights, an analytics service that monitors the performance and usage of live Web applications. We’re also providing managed services support for the entire application and chatbot.



The Outcome

Our AI-driven solution has achieved ~80% accuracy in chatbot-delivered answers to highly complex queries, and has significantly improved productivity, with increased information sharing through a properly structured role-based intranet communications channel. The bottom line is the firm now shares data far more effectively, enabling faster and better investment decisions.

Azure Application Insights


Azure Cognitive Search

Azure OpenAI

Power Apps

SharePoint

Healthcare & Life Sciences



The Challenge

With pandemic-induced hypergrowth, the client was challenged by large amounts of unstructured clinical and research data. Moreover, multiple acquisitions over the years had resulted in several data silos, leading to the creation of 70+ ERP instances for the Analytics Instruments Group.

US biotechnology leader delivers a single source of truth leading to a cost saving of up to \$1 billion on IT spending

The client is a world leader in life sciences solutions and laboratory instruments, operating across 65 worldwide locations. The organization has a rich history of acquisitions and has experienced recent hypergrowth.

Healthcare

Data & Analytics



The Solution

The client wanted Persistent to help them consolidate data from multiple sources, including purchase orders, purchase order receipts, and supplier invoices. Persistent's solution focused on consolidating data from multiple ERPs to provide 360° views and insights using AWS S3-based Data Lake.

Persistent executed the consolidation with a common framework for 10+

ERP systems covering a \$2 billion spending.

With Persistent's support, the client could store, process, and analyze data in a central repository with a cost-effective, secure, and scalable data lake. A data catalog was also developed to enable metadata discovery.

The Outcome

Consolidation of data from all ERPs ensured a single source of truth for the client, leading to a potential cost saving of \$1 billion in IT spending.

The solution gave unprecedented flexibility to business owners and analysts who could now extract rich insights from data.

AWS S3

AWS Lambda

Spark

AWS Glue

AWS DMS



Leading clinical research organization improves time to market with intelligent data platform

The client is the world's second-largest clinical research organization (CRO) conducting clinical trials on behalf of its pharmaceutical customers. The company provides biopharmaceutical services and has helped develop approximately 95% of the 200 top-selling biopharmaceuticals available in the market.

Healthcare

Data & Analytics

The Challenge

From offering regulatory consulting services to decentralized clinical trial experiences, the client's therapeutic, technical, and functional capabilities are characterized by a patient-first culture. The company relies on multiple information management systems and tools to leverage real world data for improving patient journeys.

The client needed quick and easy access to integrated patient data across the clinical trials spectrum. This required eliminating data silos in various systems across an organization and reducing the manual effort to extract insights for data-driven decisions. This was critical for creating a one-stop view of drug verification across studies for different cohorts, locations, phases, and specialties for pharma companies.

The Solution

The client partnered with Persistent to develop a central framework for enterprise data integration and an effective data lake that would enable them to centralize information from disparate data sources and improve accessibility while ensuring security, quality, and integrity of data.

Persistent established end-to-end ownership of clinical trial operations reporting based on Azure Data Platform to deliver a single source of truth. We created dashboards for key metrics for individual studies as well as for studies at the account, portfolio and executive levels across the project lifecycle, financial management, resource management, quality management, compliance, and safety.

The Outcome

The solution enabled the client to consolidate data from 3,000 trials across reporting and platform and create 15 consolidated reports that provided data-driven insights into the progress of trials. Leveraging accelerators, the client was able to speed up time to market for enhancing patient care outcomes.

Azure Data Lake

Azure SQL

Azure Data Factory

SSIS

Health-tech leader doubles patient matching accuracy for advanced cancer research with AI-powered analytics

The client is a US-based, global leader in medical technology, pharmaceutical diagnostics, and digital solutions innovator. It enables healthcare professionals to make faster and more informed decisions through intelligent devices, data, analytics, applications, and services.

Healthcare

Data & Analytics



The Challenge

The client was keen to enhance their Imaging AI platform and enable clinical researchers to accurately identify patients that are most likely to respond positively to new therapies. The client's current solutions based on single imaging

modality were not robust and efficient enough to accurately profile disease types and stages. To achieve these, they needed an experienced technology partner to assist them in integrating and processing multi-omics and radiology data.

The Solution

The healthcare technology provider collaborated with Persistent based on our proven track record of digitally transforming similar organizations.

Persistent considered genomics, radiomics, and clinical data for non-small cell lung cancer available from a public repository, to build training models on individual and various data sources. It was also a critical step in predicting the recurrence and progression of the disease. The idea was to enable accurate prediction for better and personalized patient care.

As part of the solution, Persistent developed a new AI-powered analytics and workflow solution to help the client serve both clinical and biopharma markets. It comprised a HIPPA-compliant SaaS analytics platform built on AWS. We created a data ingestion pipeline to integrate and harmonize data from electronic health records, radiology Information systems, imaging, sequencing instruments, and an integrated multimodal data processing pipeline.

The Outcome

The client has experienced a 2X boost in its efficiency and accuracy in analyzing data from multiple sources. It has allowed them to identify patients who are fit to receive new therapies with

high precision. Additionally, the longitudinal patient data visualization portal assists the client in tracking disease progression to offer personalized care.

AWS

Python

Flask Pandas

NumPy


Scikit-learn NLP

Search

Healthcare

Data & Analytics

LungLife is an American diagnostics company focused on using technology for early lung cancer detection. It provides certainty in the early-stage detection of lung cancer. It is developing a series of companion diagnostics for later-stage treatments while anticipating patients' probable therapeutic responses.



**Saving lives by
improving early-
stage lung cancer
detection rates
by 70%**

The Challenge

It was challenging for LungLife to analyze massive volumes of data of up to 15,000 microscopic images per patient. Manual analysis was time-consuming, leading to delayed cancer detection, higher false positives, and missed circulating tumor cell diagnostics.

LungLife wanted to leverage AI and ML to overcome these challenges. The company was looking to significantly impact patient outcomes and clinical decision-making by detecting lung cancer in the early stage.

The Solution

To develop AI and ML algorithms for LungLife's diagnostic technology, Persistent developed a deep learning-based segmentation model. Persistent deployed annotation tools to ensure enhanced model accuracy of microscopic images. It developed UI-based solutions to help LungLife efficiently verify cells and classify them into circulating tumor cells, single gain cells, single deletion cells, and normal cells.

The Outcome

LungLife significantly reduced diagnosis time by about 70% with the new solution. With 62% less false positive cases, the accuracy rate improved significantly. This accelerated LungLife's efforts to greatly reduce the impact of a disease that claims around 400 lives per day.

AWS

AngularJS

Keras

TensorFlow

The Challenge

The client had distributed patient and doctor data across multiple systems and silos, resulting in inaccurate and outdated information. The inconsistencies in the information published across the enterprise caused operational inefficiencies. Additionally, manual data entry increased the risk of errors, and operational delays increased tenfold. The client aimed to consolidate various data systems and partners to provide unified, global support to doctors and patients.

Improving patient on-boarding by 80% and operational efficiency by 60% for a healthcare provider

The client is a leading US-based healthcare provider supporting primary-care doctors in providing value-based care. They connect doctors with payors and patients through a seamless platform and one-of-a-kind partnership model.

Healthcare

MDM



The Solution

The client selected Persistent as their partner to implement a robust and scalable Master Data Management (MDM) platform utilizing Informatica IICS Customer 360. We profiled patient and healthcare organization data from various data silos to determine the necessary data quality, validation, and de-duplication rules.

Persistent seamlessly connected to different data sources for real-time and batch data extraction by leveraging the IICS and Snowflake technology stacks. This ensured consistent and high-quality data throughout the transformation process. We developed a metadata-

driven ingestion framework that significantly reduced the time required to onboard new source systems.

Persistent embraced an agile approach and MDM best practices to implement a scalable solution that accommodates additional source systems. The matching algorithms were continuously fine-tuned through rigorous testing data from multiple sources. The team also provided comprehensive merge reports, empowering the client with analytical insights, and enabling them to be proactive based on the merged records.

The Outcome

The MDM solution empowered the client to manage data from multiple systems efficiently. Consolidating data into a single platform simplified the process and reduced maintenance needs. The metadata ingestion framework decreased the time required for onboarding new data sources by 80%, eliminating any unnecessary overhead previously associated with the process. Additionally, the matching performance for Patient records

improved by up to 90%, effectively resolving data inaccuracy and duplication issues.

This enhancement enabled the client to provide more accurate clinical recommendations for patients/members. Furthermore, the overall process automation reduced manual efforts, minimized errors, and increased operational efficiency by approximately 60%.

Informatica C360 / R360 SaaS

AWS

Snowflake

Healthcare

MDM

The UK-based client is the world's most extensive global medical technology and equipment manufacturing company that designs, manufactures, and markets orthopedic devices and reconstruction products, sports medicine, wound management, arthroscopy, trauma, and clinical therapy solutions.



Improving accuracy in patient records from 40% to 90% for a global medical technology and equipment company

The Challenge

The client lacked access to core customer information as they relied on a basic Master Data Management (MDM) solution that stored irrelevant information, resulting in data inaccuracies and duplicate customer and product records.

Furthermore, the items were dispersed across multiple systems, resulting in redundant and insufficient item details. This led to substantial

expenses in maintaining data across multiple global markets.

The client needed a more robust mechanism to manage multiple domains within a single repository and efficiently distribute item information to various downstream applications using automated operations.

The Solution

Persistent designed a robust and scalable MDM solution for customer data by developing a unified data platform on AWS. This platform uses Snowflake as a central repository for IQVIA, GPO, SAP, and other internal source systems and contains a copy of customer data that can be distributed to appropriate downstream systems. The platform provided comprehensive match and merge Reports, generating valuable analytical and actionable insights for the client's business.

The implementation enabled the client to use Informatica Product Information Management to establish a reliable MDM database for items, materials, and classifications and publish up-to-date item information across multiple channels. Persistent optimized costs and enhanced operational efficiency by leveraging Snowflake, customizing the existing process, and orchestrating end-to-end workflow.

The Outcome

Data inaccuracies and duplications were eliminated by improving the match performance for customer records to 90% from the previous 40%. This helped the client to propose superior hospital and clinical recommendations for their customers. Migrating the SPH data, to the MDM hub decreased Informatica PIM's licensing cost.

Informatica MDM

Informatica PIM

AWS

Snowflake

**Leading clinical
research
organization
improves
forecasting by 40%
and response time
by 55%**

The client is a US-based biopharmaceutical organization providing contract research and commercial services to pharmaceutical and biotechnology firms. With integrated solutions to accelerate customer success, the company specializes in conducting late-stage clinical trials and promoting healthcare treatments in the market

Life Sciences

MDM



The Challenge

The client faced significant challenges in the volume of data given a surge in data sources. Maintaining consistent and accurate reports with fragmented data in diverse formats and structures was difficult. Repetitive and incomplete data duplicates exacerbated the issue, leading to conflicting information. The lack of reliable

reporting on site and investigator availability added to the complexity. The client missed several opportunities given the delay in getting accurate progress reports for studies and site availability. With stringent regulatory demands, the client had to ensure data accuracy and integrity.

The Solution

Persistent developed a plan to standardize and cleanse the data to improve accuracy and reliability and implemented a scalable Master Data Management (MDM) hub in Reltio MDM to maintain unique data sets related to clinical and commercial aspects. We successfully mastered approximately 2.2 million data points, which established a unified view of both clinical and commercial datasets, providing a comprehensive understanding of both. We used business process management techniques to develop a workflow for data stewardship. The team created an easily manageable system for a low-maintenance data integration framework and a data quality framework for loading data into Reltio.

Persistent configured the matching process to eliminate duplicate entries and enhanced the match rule to accommodate date-locked entries under potential matches. With key business stakeholders, we established a comprehensive enterprise data governance program, ensuring effective reference data management, business glossary, and data catalog management. The implementation established a centralized business-certified data catalog, business glossary, data model, reference data set, and business rules to promote efficiency and consistency.

The Outcome

The client saved over 55% in time and effort by implementing a data validation framework. By utilizing accurate and reliable data, also known as golden data, the client improved forecasting and

analytics by 40%. Additionally, with a streamlined process, onboarding new data sources is much faster and takes less than 6 hours.

Reltio MDM

Oracle EDQ

Pentaho

Azure Data Factory

Collibra

A global biotechnology firm reduces FTEs by 30% with legacy app modernization

The client is a leading biotechnology company that invents and commercializes medicines.

Life Sciences

Gen AI



The Challenge

Our client — a leading biotechnology firm that invents and commercializes life-saving medicines — had come to rely on a legacy application, but struggled to maintain the app. The problem was an outdated codebase that was developed in Ruby on Rails (RoR). The intricacies and obstacles inherent in this platform made it hard to find talent to maintain or update the program. The firm’s Chief Product Officer realized that a scalable app with a modern tech stack was required to drive business growth.

The Solution

The focal point of our modernization initiative was the migration from a monolithic RoR application to a microservices architecture, implemented using Python and Node.js. The adoption of microservices architecture offered scalability, resilience, and performance improvements.

By incorporating Persistent WingMate Generative AI-powered accelerators, we achieved significant reduction in time and cost.

We prioritized the creation of detailed documentation to maintain the app including insights into the architecture and maintenance procedures. This holistic approach guarantees a smooth knowledge transfer process and contributes to the sustained success and adaptability of the modernized technology.

The Outcome

Our client achieved 30% FTE time and cost savings. The solution enabled a faster go-to-market, greater scalability, and cross-platform performance.

Persistent WingMate Accelerators

Azure Open AI GPT4 Model

Python

Node.js


React

US pharma leader accelerates proposal response time by 90% with a Generative AI-based solution

The client is an American multinational healthcare corporation offering a vast portfolio of products and services for the consumer, pharmaceutical, and medical devices markets.

Pharma

Gen AI



The Challenge

The technology service (TS) organization works across 65 workgroups to create 10-12 thousand statements of work and proposal documents yearly. The existing process was manual and inefficient, with dispersed knowledge across multiple unconnected sources. To accelerate the process of designing proposals and solutions for upcoming business requirements, the client wanted to build a Generative AI-based TS CoPilot platform that would automate responses by tapping into past sources of knowledge, service, and application catalogs.

The Solution

The client partnered with Persistent to streamline the overall process for its technology services unit. The CoPilot built by Persistent taps into a variety of data sources including service catalogs, application and component inventories, and other knowledge sources to accelerate the process. Using playbooks, the CoPilot leverages agents powered by LLMs to guide the user in a step-by-step journey from business requirements to proposals.

The CoPilot also interacts with predictive models to estimate timelines, personnel requirements, and costs for the proposed solutions. Using past solution diagrams and solution descriptions, the CoPilot can also generate these for the new business requirement, using multi-modal large models.

The Outcome

The TS CoPilot decreased human touch points by 70-80%, accelerating the proposal response time by 90%. The platform connected disparate technology assets to act as a single source of truth, leading to higher utilization.


Azure Open AI GPT 4

Qdrant Vector Db

Langchain framework

Streamlit UI

Telecom, CPG, Food &
Beverages



Digital communications technology leader leverages AI to reduce operations and support costs by 80%

The client is a global technology firm offering products and services across networking, security, collaboration, applications, and the cloud. The company develops, manufactures, and sells telecommunications equipment, networking hardware, software, and other high-technology products and services

Telecom

Data & Analytics

The Challenge

The client was challenged by an overwhelming volume of support requests that its technical support team had to manage daily.

The issue resolution time was extremely high as only 70% of customers used self-service facilities, and only 9% were able to resolve issues entirely on their own, via self-service. The lack of standalone applications capturing siloed information was a barrier to ensuring successful self-service-based issue resolution.

The Solution

The client wanted to enable customer self-service for more rapid and accurate search and a reduction in support costs. Empowering the customers was expected to achieve improved satisfaction scores and loyalty.

Persistent adopted a design thinking-led approach to meet the client's demands of an AI-driven case deflection with self-service issue resolution. This was to align all end use customer touchpoints with the client's business outcomes. We implemented an outside-in industry perspective, deploying an AI feedback loop model for proactive suggestions and implicit deflections.

The Outcome

The client cut operational costs by 80% and reduced the physical data center footprint by a significant margin. Over 150 reusable components were built, including identity management, PaaS, and security management components, to ensure cost optimization.

The migration of 3,500 apps to Azure cloud helped the client fulfill the commitment of \$2 billion Azure consumption over seven years.



Python

Elastic Search

Swiss Food and Beverage Leader on-boards new markets 75% faster and improves compliance by 80%

Global food and drink processing conglomerate with over 2000 brands and a presence in 191 countries.

CPG, Food & Beverages

MDM



The Challenge

The multinational corporation operates in multiple countries worldwide and adheres to each region's unique nutrition and healthcare standards to maintain regulatory compliance. It offers a wide range of products, including baby food, medical food, bottled water, breakfast cereals, coffee, tea, confectionery, dairy products, frozen food, ice cream, pet food, and snacks. The company's complex inventory management system handles a high volume and variety of products, with the bill of materials managed through a collaborative effort between markets.

The presence of products in various systems resulted in duplication and insufficient product information. This led to inconsistent customer experiences and high costs for maintaining data in multiple markets. Weak data validations, checks, and excessive manual operations hindered regulatory compliance measures. The lack of a centralized repository made it challenging to identify market-specific key performance indicators. Multiple levels of manual collaboration among organizational roles and users caused delays in publishing products downstream and consuming systems.

The Solution

Persistent created a reliable Master Data Management repository for product, material and bill of materials classifications and distributing current item information to channels by regions. This solution ensured data quality according to GS1 standards, and an additional data quality and validation framework to guarantee continuous data refinement through a well-defined governance

process, for each country or market. The data quality rule framework accelerated the data enrichment process by enabling an automatic population of attributes. The solution allowed the customer to onboard new markets for their product data via historical configurations. The deployment strategy utilized an Azure Cloud-based containerized environment on OpenShift.

The Outcome

The result was faster onboarding processes for new markets, with an approximately 75-80% acceleration rate. The solution has reduced data

validation efforts by 80%, ensuring data quality and compliance with GS1, nutrition and healthcare standards, and regulations.

Angular

Red hat OpenShift

Azure

IBM Product Master



persistent.com/data-and-analytics →

Persistent Data and Analytics Practice at a Glance

3,000+ Data Professionals	33+ Years of Delivery Excellence	1,000+ Client Engagements	7 Delivery Centres	80+ Active Clients
Data Strategy & Consulting <ul style="list-style-type: none"> / Data strategy, governance and compliance / Master data management, data quality management 	Enterprise Data Management <ul style="list-style-type: none"> / Data engineering including new development and migration / Data as a Service 	BI & Visualization <ul style="list-style-type: none"> / BI tool / reports assessment and rationalization / Dashboards & reports for CXO & business users / Self-service BI enablement / Regulatory & compliance reporting 	AURA: Data Solutions Powered by AI and ML <ul style="list-style-type: none"> / Data and Documents Insights / Data Ops: Data Quality, Data Reconciliation / Platforms: Data Foundry, Gen AI Hub 	

Why Persistent?

AI Expertise Recognized as a Gen AI Leader by HFS Research	Persistent Data Foundry Pre-built ready to use data modernization platform	360° Industrial Presence BFSI, HCLS, Software and Hi-Tech, Telecom, Media, Retail	Rapid Innovation Hackathons, studios & Persistent University
--	--	---	--

We are large enough to drive global scale but small enough to be nimble and caring

Persistent: Partnership Overview (Ecosystem Orchestrator of Choice)

 360° Partnership Snowflake Premier Tier Partner	 Databricks — Global SI and PS Partner	 360° Partnership Direct Cloud Services Provider Reseller and Solution Partner Designation for Six Solution Areas
 360° Partnership Google Cloud Premier Partner	 Product Engineering Partner for Appflow, Kendra, SageMaker and CodeGuru	

Persistent's Differentiators

Prebuilt Accelerators for Operations Management	Multiple use cases and domains. Eg., Marketing, Revenue, Commerce, Patient 360, & Genome Sequencing	Data Quality Assessment, "Cost - Monitoring Dashboard", Reports and Predictive Algorithms	Data Security and Access Control Management
---	---	---	---

**See
Beyond.
Rise
Above.**

About Persistent

With over 23,000 employees located in 20 countries, Persistent Systems (BSE & NSE: PERSISTENT) is a global services and solutions company delivering Digital Engineering and Enterprise Modernization. We work with the industry leaders including 14 of the 30 most innovative companies as identified by BCG, 8 of the top 10 largest banks in the US and India, and numerous innovators across the healthcare and software ecosystems. As a participant of the United Nations Global Compact, Persistent is committed to aligning strategies and operations with universal principles on human rights, labour, environment, and anti-corruption, as well as take actions that advance societal goals.

USA

Persistent Systems, Inc.
2055 Laurelwood Road
Suite 210, Santa Clara
CA 95054
Tel: +1 (408) 216 7010

India

Persistent Systems Limited
Bhageerath, 402
Senapati Bapat Road
Pune 411016
Tel: +91 (20) 6703 0000

