



Persistent

# Generative AI – **The Persistent Way**





\*Image generated by Dall-E

# A changing world: **Generative AI**

Software consumed the world.  
Now Generative AI is doing the same

## Message from CEO

As a pioneer in Digital Engineering, we have always believed in staying ahead of the disruptive technology curves. With Generative AI, we're once again seizing the opportunity to build solutions that revolutionize digital engineering & enterprise modernization and accelerate value creation for our global clients. Based on our 150+ customer conversations through this year, we have increasing validation of our assets and capabilities and nowhere is our market advantage clearer than with our investments in Generative AI.

Generative AI is proving to be a powerful ally for Digital Engineering. As the tools and capabilities evolve, it will dramatically reduce development times, and will allow developers to concentrate on what matters the most – ideation and innovation. Generative AI is helping businesses design newer products & services and create customized experiences all along the customer journeys. Through our early engagements in Generative AI, our customers are already unlocking new levers for business performance in terms of revenue growth and improved operational efficiencies.

With the evolution multi-modal, purpose-built models of Generative AI and the immense possibilities around them, our customers need a reliable partner who can help them with the right implementations. At Persistent, we've built an integrated team of experts trained on major AI platforms who are focused on developing solutions for our clients which not only generate cost savings and accelerate time to market, but also uphold the highest levels of data security, compliance, and privacy.

We are constantly expanding our partnerships with global AI leaders such as Microsoft, Google, AWS, IBM, Salesforce and promising start-ups to develop and deliver innovative solutions to the market. Our proactive approach and deep collaboration with our partners have enabled us to provide game-changing Generative AI solutions for our clients in software, banking, financial services, healthcare, and other industries.

I am ambitious about the business value that Generative AI can unlock and its ability to power our journey towards \$2 billion and beyond. I am looking forward to team up with each one of you and make Generative AI the new way of working at Persistent!

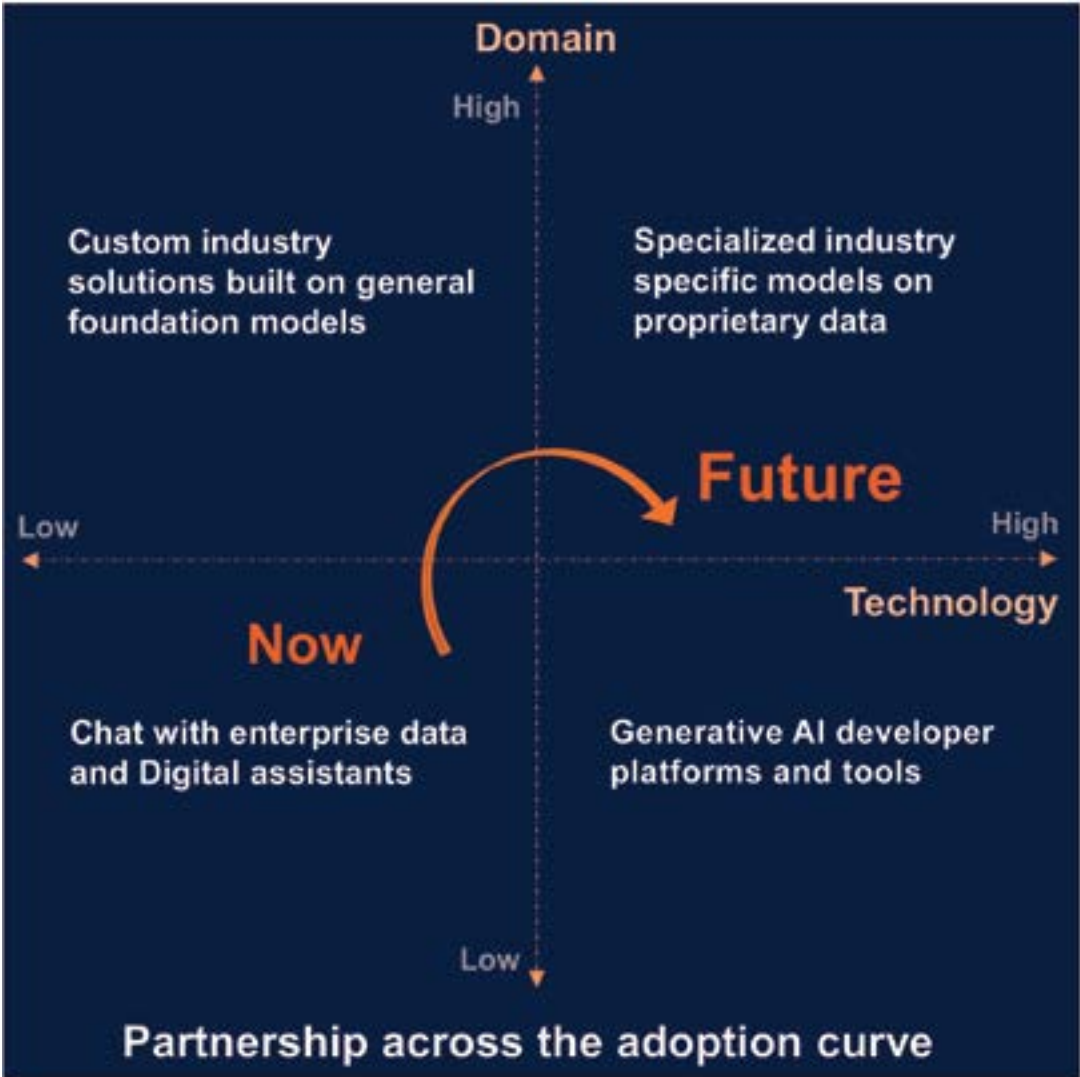


**Sandeep Kalra**

Chief Executive Officer and Executive Director



# Re-imagining Digital Engineering with **Generative AI**



## Custom industry solutions on foundation models

A generation of Generative AI inspired ideas is leading to more innovative and effective designs across functions, processes and industries. Partnering with Persistent, our customers can train their models with their company specific information to create purpose-built applications.




## Specialized industry specific models on proprietary data

As the Generative AI implementations become pervasive, companies will insist on proprietary models to ensure data privacy, and exclusive access to business benefits that come from marrying their specific context with next generation of software powered by Generative AI.



## Chat with enterprise data and Digital Assistants

Generative AI are enabling better conversational interfaces, either through chatbots or digital agents, for customer service, employee engagement and other applications. It can produce original content in response to queries, drawing from previous user interactions in style unique to that business.



## Generative AI developer platforms and tools

Product Engineering is ushering a new age of productivity, efficiency and reliability. Co-pilot programs, where Generative AI based tools and digital assistants work alongside software product engineering teams, will help in sustaining existing applications and built modern ones more efficiently with higher security.



# Generative AI at Persistent

We have always been at the forefront of digital innovation and have systematically bolstered our capabilities in software product development, application modernization, cloud enablement, enterprise security, intelligent automation, and data and AI.

There are many ways for enterprises to take advantage of Generative AI for real business benefits, but they must work with an experienced and trusted partner that can determine the best use cases while managing risk and avoiding unnecessary delays. We are in market today with offerings and solutions that can bring the power of Generative AI to your business.


## Our focus: Drive growth and turbo charge productivity across our clients



- \\ Industry Specific Solutions Banking, Insurance, Healthcare, communications, Media, Retail
- \\ Customized solutions



- \\ Modernize rule based chatbots
- \\ Content summarization
- \\ Content Generation – Media Content, Blogs



- \\ Natural Language Search for enterprise stakeholder and external users
- \\ Complex data (structured and unstructured) made intuitively accessible



- \\ SDLC Productivity
- \\ Code Migration
- \\ Proactive software maintenance– upgrade paths and remediation

**Trends and Emerging Use Case Patterns from over 150+ clients**



As the Digital Engineering leader, we know Generative AI is a top priority for CIOs and CTOs across ISVs and Hi-Tech.

We leverage Generative AI’s transformative capabilities to reduce manual operations, lower costs, spur faster revenue growth and enhance customer experiences.

**Generative AI Use Cases for Software and Hi-Tech**

Our solutions guide customers through their Generative AI use cases end-to-end — from design, user experience and data modernization to customization, automation and deployment.

**Engineering**

- \ WingMate and ExtenSURE Frameworks
- \ Code Generation
- \ Code Migration (legacy tech to new tech)
- \ Test Case Generation
- \ Document Generation
- \ Code Vulnerability Scanning / Identification
- \ Infrastructure Deployment Automation

**Customer and IT Support Generative AI Agents**

- \ Support Tickets Automation
- \ Self-service Troubleshooting
- \ Agent Assist
- \ Case Notes Summarizations

**Human Capital Management**

- \ Assisted Authoring for job descriptions
- \ Resume Filtering and Candidate Screening
- \ Employee Onboarding and Engagement
- \ Vendor Contracts Creation
- \ Contract Renewal and Terminations

**Supply Chain**

- \ Exception Processing and Dispute Handling and Resolution
- \ Fraud Prevention by refining models with synthetic data generated using large language models

# Generative AI in Software and Hi-Tech





# Generative AI in BFSI

Generative AI is transforming BFSI at incredible speed, with the potential to unlock new digital services, enable operational scale and create groundbreaking user experiences. Our expertise across banking, insurance, wealth management and other related industries, combined with our Digital Engineering skills and data offerings, allows clients to unleash new growth and unlimited scale with Generative AI.

## Vertical specific Use-cases for Banking and Financial Institutions

Banking	Payments
<ul style="list-style-type: none"><li>\ Product Recommendation and FAQ</li><li>\ Contact center optimization with solutions around voice recognition, call deflection and content search</li><li>\ Customer onboarding functions like authentication, AML-KYC checks and document verification</li><li>\ Regulatory compliance automation with use of synthetic datasets</li></ul>	<ul style="list-style-type: none"><li>\ Exception processing and handling of disputes</li><li>\ Fraud prevention by refining models with synthetic data generated using LLMs</li></ul>
Wealth and Investments	Insurance
<ul style="list-style-type: none"><li>\ Robo-Advisory – FAQ led wealth product recommendation and advice</li><li>\ Customer portfolio reporting with interactive dashboards</li><li>\ Sales Pitchbook generation</li></ul>	<ul style="list-style-type: none"><li>\ Policy comparison and insights for quicker sales cycle</li><li>\ Customer service assistant to search large data sources for queries, documents and policy servicing related content</li><li>\ Underwriting process improvements with precise insights and data points</li></ul>



We deliver positive patient outcomes, improve diagnostics and analytics, and enhance employee productivity through our Generative AI use cases.

We have deep expertise working with healthcare providers and life sciences companies to develop secure solutions combining the power of cloud, data, automation and Generative AI that push the boundaries of scientific discovery and create scalable value.

**Generative AI Use Cases for Healthcare and Life Sciences**

Our team of tech and domain experts help companies and providers take the guesswork out of creating Generative AI solutions that deliver tangible business benefits in a responsible way.

Diagnostic Imaging	Multimodal Analysis	Diagnostic Imaging	Diagnostic Imaging
Preclinical and Clinical Research	Regulatory / Compliance Document Generation	Revenue Cycle Management	Clinical and Policy Document Management
Authorization, Complex Care Management and Patient Personalization		Screening, Triage and Decision-Making Assistants	

# Generative AI in Healthcare and Life Sciences





# Why Persistent is the Generative AI Partner of Choice





# Partnerships

We collaborate closely with global AI leaders and promising start-ups to develop and deliver Generative AI solutions to the market. Our software engineering and product development experience gives us a distinct advantage versus competitors in co-developing new offerings with our partners.



- Part of Microsoft AI Partner council, a program for Microsoft partners with demonstrated leadership in Generative AI and leading the Open AI services Implementation
- Amongst select partners for the EAP program on M365 Co-pilot
- Microsoft Viva Practice Integrated with Generative AI for EX transformation



- Generative AI Strategic Collaboration Agreement (SCA) in place
- Approved as a partner in the Generative AI Co-Sell Motion at AWS
- Early access and engineering partner for Amazon Generative AI services
- Featured in and presented in multiple AWS events
- Hosted joint customer roundtables



- Participate in Build, Sell, and Service engagements
- TTP (Trusted Testing Partner) for Generative AI
- Working with Google Generative AI PDMs to co-sell our services through Google Sellers to end customers



- IBM has provisioned dedicated Generative AI co-innovation sandbox to enable us to create solutions and capabilities for clients across industries.
- Dedicated Watsonx CoE at Persistent



With a staff that has attained more than 5,000 Salesforce certifications, we're building unique Generative AI-powered solutions on top of Einstein GPT for vital vertical industry functions, such as fraud detection and automated loan underwriting in financial services, and risk prediction and personalized triage in healthcare.

## Persistent named Generative AI Market Leader in HFS Horizons: Generative Enterprise™ Services, 2023 Report



Note: All service providers within a “Horizon” are listed alphabetically



Growing list of Persistent Generative AI Accelerators

### WingMate

Generative AI Powered SDLC companion, code modernization

### Offer Wise

Hyper-personalized recommendations with explanations

### Knowledge Platform

Platform to Create, Compile, Query and Analyze Knowledge

### Content Assist

Generate enterprise relevant content on the fly

Persistent Digital Engineering Platform powered by Generative OS

AI Assisted Product Takeover			AI Assisted Product Engineering		AI Assisted Product Modernization			AI Assisted Product Sustenance and Support		
Interventions/ pain points	Code Review/ Auditing	Test Case Generation	Migration (Code, App)	Code Health	Technical Debt Reduction/ Architecture Review	Tech Due Diligence	DevSecOps	Knowledge Engineering	Guided Upgrade and Remediation	AI Powered Sustenance
	Data generation, provisioning, masking	Cloud dev/ Migration	AI-based Digital Assistants							
Assets/Platform Modules	Persistent WingMate	Persistent PiCloud	Agent Framework	Persistent ExtenSURE.ai					Persistent Open Source Sustenance Platform	
	Persistent Generative Enterprise OS									
Backbone	Playground Connect data and models		Skills Framework Support for new programming languages, testing libraries, etc.		Evaluation Framework Responsible AI driven evaluation of models and RAG pipelines			Knowledge-base Knowledge graphs and Vector DB with patterns on code modernization		
	Governance		User Management		Observability		Cost Management		Security	
	HYPERSCALERS – AWS, Azure, Google, IBM, Salesforce									

### Smart Clinical Trials

Patient Cohort Identification using Generative AI

### LLM Evaluation framework

Automated evaluation of RAG

### AURA

Automated reporting and analytics solutions

### Contracts Intelligence

Ingest contracts and enable intra and inter document search.







## AI Requires Data, and We're Data Experts

Our company's beginnings are rooted in data systems. With decades of experience, we are the ideal partner for companies looking to utilize Generative AI through connections to a variety of enterprise data sets for scalable solutions that adhere to all security, privacy and governance requirements.



## Digital Engineering Leadership

Software is core to our business and Generative AI allows us to extend our unmatched software engineering leadership into new Generative AI engagements. We work with clients across industries to use the technology to turbocharge software development, employee and customer services, and workplace processes.



## IP, Accelerators and Investments

We're ambitious about the efficiencies, outcomes and value that Generative AI can generate for clients. It's reflected in our ongoing IP investments, our extensive staff training on partner platforms, and our suite of Generative AI accelerators to improve workforce productivity, streamline app modernization and accelerate software engineering.



## Deep Domain Knowledge

We work with premier brands across financial services, life sciences and software. Working with our clients, we apply our industry expertise to develop Generative AI use cases to harness new data insights, create industry buyer personas, test new models, and provide employee and customer-facing digital agents.

## Our Commitment to Responsible AI

Generative AI presents tremendous possibilities, yet it also requires working with a provider that is a proven and trusted partner of choice for enterprises. Persistent is building on its data heritage to further enhance our extensive AI experience — because AI requires a full and comprehensive knowledge of how to utilize and maintain data in all aspects to deliver value. We have partnerships with major AI players, ongoing investments and training, a suite of accelerators, and industry-specific use cases for verticals such as financial services, healthcare and life sciences. Most importantly, we're committed to maintaining data security, privacy, governance and transparency in all of our client engagements.

### Responsible AI at the core of all our Generative AI implementations



#### Accountable

- Human-in-the loop
- Bias and Fairness



#### Reproducible

- MLOps pipelines
- Data Catalog, Model Registry



#### Transparent

- Explainability
- Interpretability
- Data/Concept Drift



#### Secure

- Error cohorts' analysis
- Confidential computing



#### Private

- Federated Learning
- Differential Privacy





# Our **Generative AI** Client Success Stories



# Improved Customer Experience for a Provider of SaaS Tax Software



Financial Services

Contact Center

Customer experience

The client is a provider of online tax preparation solutions.

## The Challenge

Our client offers online tax preparation solutions, allowing American consumers and tax professionals to electronically file state and/or federal returns. While this definitely 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, which 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 that deploying an AI-powered support solution would 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, integrating information from multiple repositories that were previously separate and siloed. Next, 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 base 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, including 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, as well as a better end result.

## The Outcome

80%

Faster turnaround times

75%

Improvement in response accuracy

40%

Savings in total FTE cost

Azure Stack

Azure Open AI Models

Azure Cognitive Search

Azure BOT Service



# Personalized recommendation engine for a Fintech Leader



Financial Services

Digital Engineering

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

## 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 enable them to generate personalized marketing and product offers quickly and at scale. Technically, this called for a Generative AI-based analytical solution capable of identifying new product recommendation opportunities based on customer behaviors – and our client saw this approach as the key to engaging and retaining customers.

## 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)-based 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. And finally, we harnessed the power of Large Language Models (LLMs) 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 (ML) 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 Jumpstart to accelerate the development effort.

Other products and technologies employed in this innovative build included the LLAMA2-7B LLMs from Hugging Face (a collection of pre-trained and fine-tuned generative text models); LangChain’s use cases for rapid application creation; and a Pandas DataFrame (a Python library for working with data sets). The entire solution was coded in Python and React.js.

## The Outcome

Persistent’s multi-phase Generative AI/ML solution enabled the client to identify patterns from customer transaction data, generate profile summaries, and then use those summaries for personalized marketing. All the output from the solution can be described as “explainable recommendations,” — the LLM-generated personalized recommendations and marketing content are clearly based on known, observable customer behaviors. In short, Generative AI-driven data insights now enable Fintech clients to transform transactional data into interpretable human profile summaries and attractive personalized product offers that drive new business.

AWS Sagemaker and AI21

Jurassic 2 Mid

AWS Jumpstart

Fine-tuned

LLAMA2-7B

LangChain

Pandas

Python

React.js



# Modernization of a Legacy App for a Global Biotechnology Firm



Life Sciences

Software Product Engineering

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

## The Challenge

One of the hallmarks of 21<sup>st</sup>-century technology is the reliance on digital engineering across all the sciences, and not just the IT industry. Our client – a leading biotechnology firm that invents and commercializes life-saving medicines – had come to rely on a legacy mobile application, but struggled to maintain the app. The problem was an outdated codebase — the app was developed in Ruby on Rails, and the intricacies and obstacles inherent in this platform made it hard to find talent to maintain or update the program. In addition, the lack of test cases and documentation posed a significant challenge in comprehending the legacy codebase, and the app’s monolithic architecture restricted its ability to scale. The firm’s Chief Product Officer realized that a scalable app with a modern tech stack was what the firm needed to drive business growth.

## The Solution

Persistent delivered an innovative microservices-based application architecture developed in Python, Node.js, and React. This modern stack has significant advantages, compared with the legacy platform: Node.js provides a cross-platform, open-source server environment that runs on Windows, Linux, Unix, and macOS, and React offers an open-source front-end JavaScript library for quickly building user interfaces based on components. Furthermore, our software engineering team leveraged Docker for containerization, simplifying the development and delivery of the updated application.

We addressed the lack of test cases and documentation by producing unit tests using two code-specific utilities. The team used Pytest for the Python codebase – leveraging a testing framework that originated from the PyPy project – while Jest was utilized as a testing framework for the app’s JavaScript codebase. And everything needed to maintain the app was fully documented in a new User Manual.

From the standpoint of “time and resources,” the project was completed more quickly and at a lower cost thanks to the use of Persistent WingMate Generative AI-powered accelerators in combination with API Gen, for generation of API client libraries; Google Convertor to help convert units of volume, weight, length, and temperature; and AI-enabled image creation.

## The Outcome

**Greater scalability**

and cross-platform performance

**Faster**

Go-to-market times

**30%**

FTE time and cost savings

Persistent WingMate Accelerators

Azure Open AI GPT4 Model

Python

Node.js

React



# Building a Generative AI-Driven Employee Collaboration Platform for a Private Equity Firm



Private Equity

Software Product Engineering

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

## 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 a highly 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 internal intranet to manage and share files with company employees. SharePoint is designed for creating Web pages and document libraries, while presenting videos, news, and updates; Power Apps enables developers to quickly build automated workflows and dashboards, connected to multiple data sources.

The chatbot app is powered by Microsoft Azure Open AI and Azure Cognitive Search – employed 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 (including the chatbot) on an ongoing basis.

## The Outcome

Our AI-driven intranet 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. Thanks to the easy app-building tools in Power Apps, the new solution has an enhanced UX with a modern look and feel and greater ease of use. The bottom line, the firm now shares data far more effectively, enabling faster and better investment decisions.

SharePoint

Power Apps

Azure OpenAI

Azure Cognitive Search

Azure Application Insights



# Getting Started

As the Generative AI trend continues apace, many companies will struggle with determining the proper starting point. Many possess data, process and business expertise, but need help determining a plan to apply those skills in the appropriate ways to create Generative AI value.

Every company is different in its approach to these types of strategic projects, its IT infrastructure and its security procedures. Nonetheless, any Generative AI deployment needs to take into consideration several factors:

## Outcomes

Companies should first assess and agree the business or IT challenges that would benefit from the generative capabilities of Foundation Models using corporate data. Strong benchmarks and metrics (e.g., speed, efficiency, customer satisfaction) are required for current operations, processes and services. These can be used to set reasonable and achievable expectations and goals with a Generative AI deployment.

## Data

Generative AI is intrinsically dependent on data, and in most organizations, data lives in different places, in various formats, and in assorted states of quality. Depending on the use case, a company needs to identify what data dependencies are required to utilize Generative AI and be mindful of how that information is presented to users for maximum explainability and transparency.

## Cloud

Like data, companies need a scalable cloud infrastructure to power Foundation Models. A Generative AI deployment is an ideal time for a company to consider modernizing and remediating existing cloud models — whether private, public or hybrid — to ensure that it possesses the scale needed at an affordable price point.

## Security

Generative AI can expose an organization to new unforeseen levels of risk if not properly deployed. Companies must assess their current security readiness and resiliency, and add proper data safeguards and guardrails to ensure the highest levels of security, privacy, and proper governance.

**Persistent engages with customers to plot a Generative AI plan around these and other functional areas, allowing us to ensure that the potential benefits of a Generative AI deployment quickly come to fruition.**

**Our Generative AI Workshop can help you get started. We start with an initial discussion to understand your business needs and data availability. Within 60 days we will implement a Generative AI data pipeline followed by fine-tuning. From there we agree with you on the next steps: engage with Persistent for implementation, continue the work in-house or decommission the use case, data and outcomes.**



# FastStart



## Week 1

### Generative AI Deep Dive Ideation Workshop

Immersive ideation workshop for stakeholders to build the AI adoption roadmap.

**Attendees:** Product, Business, Tech

**Time commitment:** 2 days.

## Week 2

### Solution Design Best Practices

Design in context of business needs and deployment needs including data governance, privacy, UX, MLOps, CI/CD.

**Attendees:** Tech, Infosec, Data

**Time commitment:** 2 days.

## Week 3

### First Generative AI deployment

Build out the first use-case for Generative AI based on the AI adoption roadmap. End to end integration with upstream and downstream systems.

**Attendees:** Product, Business, Tech

**Time commitment:** 2 days.

## Week 4

### Path to production. Next Steps:

Demo day. Review the Generative AI adoption roadmap. Review next steps.

**Attendees:** Product, Business, Tech

**Time commitment:** 1 day.

## Workshop Takeaways



### Awareness

Gain a deep understanding of Generative AI and its potential impact on your organization. Dive into the concepts, techniques, and applications that will shape your organization's AI journey.



### Ideation


Develop an AI adoption roadmap and architecture tailored to your specific needs. Chart a clear path forward, identifying the key milestones and resources required for successful implementation.



### First Step

Take the first tangible step towards AI transformation. Integrate a pre-trained LLM (Language Model), connect a data source, incorporate a user-friendly front-end interface, and establish a pathway to production.





**Persistent**

See Beyond, Rise Above

[www.persistent.com](http://www.persistent.com)