

# Generative AI for Enterprises

---

*Essential insights for decision makers*

---

Vishal Anand



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

First Edition 2024

Copyright © BPB Publications, India

ISBN: 978-93-55516-978

*All Rights Reserved.* No part of this publication may be reproduced, distributed or transmitted in any form or by any means or stored in a database or retrieval system, without the prior written permission of the publisher with the exception to the program listings which may be entered, stored and executed in a computer system, but they can not be reproduced by the means of publication, photocopy, recording, or by any electronic and mechanical means.

### **LIMITS OF LIABILITY AND DISCLAIMER OF WARRANTY**

The information contained in this book is true to correct and the best of author's and publisher's knowledge. The author has made every effort to ensure the accuracy of these publications, but publisher cannot be held responsible for any loss or damage arising from any information in this book.

All trademarks referred to in the book are acknowledged as properties of their respective owners but BPB Publications cannot guarantee the accuracy of this information.

To View Complete  
BPB Publications Catalogue  
Scan the QR Code:



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

Kup ksi k

**Dedicated to**

*My beloved parents  
and  
Mother nature*

## Forewords

It is with great pleasure that I introduce *Generative AI for Enterprises*, a comprehensive guide to building enterprise-class Generative AI (GenAI) solutions, authored by my esteemed colleague, Vishal Anand. In this book, Vishal leverages his extensive architectural experience to bring a nuanced perspective on the future of GenAI and its transformative potential for modern enterprises.

Vishal's work stands out not only because of his technical acumen but also due to his thought leadership in describing the future trajectory of GenAI. He offers a forward-looking vision that is rooted in practical, timeless architectural insights. This dual focus ensures that while GenAI will undoubtedly evolve, the foundational principles and best practices outlined here will remain relevant, providing a robust framework for enterprise application.

One of the key strengths of this book is its in-depth exploration of the unique needs of enterprises. Vishal thoroughly understands that enterprise environments come with their own set of challenges and requirements. By addressing these complexities, he provides readers with a rich, actionable guide to implementing GenAI in a way that aligns with corporate goals, enhances operational efficiency, and drives innovation.

Throughout the book, Vishal meticulously details the rise of GenAI in enterprises, covering advanced concepts such as operating models, orchestration platforms, and the use of landing zones for building scalable applications powered by language models. His insights into model selection, deployment patterns, and cost optimization strategies are invaluable for decision-makers, IT professionals, and anyone looking to harness the transformative power of Generative AI.

Moreover, the book delves into critical areas such as ethical dimensions, model governance, and the practicalities of integrating GenAI into existing enterprise ecosystems. These discussions are vital for ensuring that AI technologies are implemented responsibly and effectively, augmenting human capabilities without causing inadvertent harm.

In summary, Vishal's book is an essential resource for anyone looking to lead the implementation of Generative AI within an enterprise context. It combines visionary thought leadership with practical, detailed guidance, making it an indispensable tool for driving successful AI transformations. I am confident that readers will find this book not only informative and insightful but also deeply inspiring as they navigate the exciting frontier of Generative AI.

***Blaine Dolph***

*IBM Fellow*

*CTO Assets, Offerings and Industries,  
IBM Consulting*

**Disclaimer:** Views presented in the book are entirely his own and not necessarily those of IBM.

In the rapidly evolving landscape of technology, few innovations have sparked as much excitement and transformation as Artificial Intelligence (AI). At the precipice of a new era, AI promises to revolutionize industries, redefine business paradigms, and reshape our everyday lives. For organizations, particularly for Chief Information Officers (CIOs) and Chief Technology Officers (CTOs), understanding the architecture and design considerations that underpin this powerful technology is imperative. The strategic deployment of AI can be a game-changer, offering competitive advantages and driving innovation. However, this requires a deep understanding of AI's potential, capabilities, and limitations.

As we venture deeper into the AI era, advancements in this field leapfrog each other, propelling us into uncharted territories with unprecedented speed. Technologies that took years to evolve now do so in months if not weeks. This constant flux can be daunting, especially for executives steering their organizations through these turbulent waters. The notion of beta has become a permanent fixture, with systems and solutions perpetually in a state of refinement and improvement. Therefore, the need for a comprehensive architecture and design guide for AI solutions has never been more pressing.

This book serves as an essential guide to AI, addressing challenges head-on and providing a structured understanding of its intricacies and applications. It emphasizes the strategic considerations necessary to effectively harness AI's potential, serving as both a compass and a roadmap for executives navigating the maze of AI technologies, frameworks, and methodologies. In a world where data is the new oil, the ability to seamlessly combine and analyze vast amounts of information from different domains is paramount. This book delves into the architectural principles that facilitate such integration, offering insights into building robust, scalable systems capable of handling the demands of modern AI applications. Finally, as AI systems become more autonomous and influential and are embedded into everything we do and experience, ensuring that they operate within

ethical boundaries and align with societal values is crucial. This book provides a thoughtful discourse on these issues, underscoring the need for transparency, accountability, and fairness in AI design and implementation.

Having collaborated with Vishal on numerous projects, I have witnessed firsthand his deep technical skills and expertise in applying technology in an enterprise setting. This book reflects this pragmatic approach, distilling complex concepts into visionary insights and practical advice. It is a testament to Vishal's dedication and brilliance, offering invaluable guidelines for building and industrializing AI solutions that deliver meaningful business outcomes.

Combining cutting-edge research with real-world knowledge, this book provides readers with a rich tapestry of actionable, forward-thinking expertise. Each chapter is meticulously crafted, offering insights into AI advancements while grounding them in practical, executable strategies. This book will serve as a beacon of knowledge and inspiration to help you unlock AI's full potential and lead your organization with vision and confidence.

***Srini Koushik***

*President*

*AI, Technology, and Sustainability -  
Rackspace Technology Global Leader,  
Foundry for AI by Rackspace (FAIR)*

In a world where technology is advancing at an unprecedented pace, enterprises find themselves at a pivotal juncture, faced with the challenge of harnessing cutting-edge innovations to drive growth and competitiveness. Among these innovations, Generative AI stands out as a transformative force, poised to revolutionize the way businesses operate, innovate, and scale.

Generative AI for Enterprises, authored by Vishal Anand, is a timely and indispensable guide for leaders and professionals navigating the complex landscape of Generative AI. With over two decades of experience in IT and a distinguished career at IBM, Vishal brings a wealth of knowledge and practical expertise to this comprehensive work.

This book is not merely a technical manual; it is a strategic blueprint for decision-makers aiming to leverage Generative AI to its fullest potential. Vishal meticulously explores the intricate dynamics of building enterprise-class Generative AI solutions, emphasizing the importance of understanding the multifaceted requirements of modern enterprises. From the transformative impact of Generative AI across industries to the nuanced challenges of deploying and fine-tuning models, each chapter offers valuable insights and practical guidance.

Vishal's deep dive into advanced concepts such as operating systems, orchestration platforms, and landing zones highlights the critical infrastructure needed for robust and scalable applications powered by language models. His exploration of prompt engineering, fine-tuning, and orchestration of Generative AI workflows provides a detailed roadmap for integrating these technologies seamlessly into enterprise environments.

Moreover, the book addresses essential ethical dimensions, ensuring that the deployment of AI technologies augments human capabilities responsibly and transparently. By emphasizing principles such as accountability, fairness, and security, Vishal lays a foundation for the ethical integration of Generative AI into the corporate world.



For decision-makers and experienced professionals alike, this book offers a wealth of practical wisdom and best practices. It equips readers with the knowledge and skills to become proficient leaders in the field of modern enterprise development using Generative AI. Vishal's expert perspective and hands-on demonstrations make complex concepts accessible, enabling readers to implement effective and sustainable AI solutions.

In an era where the ability to innovate and scale can determine the success or failure of an enterprise, Generative AI for Enterprises is an essential resource. It provides the strategic insights and technical expertise needed to navigate the challenges and seize the opportunities presented by Generative AI. I am confident that this book will serve as a valuable guide for leaders and professionals seeking to drive their organizations forward in the digital age.

With this comprehensive guide in hand, you are well-equipped to embark on the journey of implementing Generative AI in your enterprise, ensuring success beyond scaling. Embrace the insights and strategies shared by Vishal Anand, and lead your organization into a future where AI-powered innovation is at the forefront of business growth and transformation.

***AB Vijay Kumar***

*Master Inventor and Global  
CTO for IBM Consulting*

**Disclaimer:** Views presented in the book are entirely his own and not necessarily those of IBM.

When Vishal Anand, a close collaborator at work, reached out and shared with me his new book, *Generative AI for Enterprises*, I was intrigued and wanted to learn more as we have witnessed a massive and rapid change in the capabilities of AI technologies in recent years. **Generative AI for Enterprises** book comes out at the right time, as over 80% of enterprises are adopting Generative AI<sup>1</sup>, yet 60% of leaders have yet to deploy Generative AI for commercial activities<sup>2</sup>. There are enormous opportunities for leaders to learn about and adopt the technology and practices in this book. So, if you are reading this foreword, you are on the right track.

As new large language models come out on almost a daily basis, with different properties and implications, Mr. Anand discusses key considerations for model selection as well as implications on the production environment requirements. This book covers the entire stack of decisions that leaders need to consider when devising their Generative AI strategy, from platform and runtime design to model deployment, serving and optimization; and operational challenges. Mr. Anand introduces different strategies for an enterprise to be an AI value creator, from contextually crafted prompts to fine tuning, and explores the value of retrieval-augmented generation techniques.

As the models are becoming more powerful, available in a multitude of domains beyond language such as code, time series, and vision, they take the shape of multi-modal solutions. The book ends with a viewpoint of the next wave of upcoming innovations namely (multi) agent-centric frameworks, where we will witness the interplay of agents to compose, plan, execute and introspect the crafted outcomes.

Vishal Anand brings his extensive technical expertise and experiences in client engagements and delivers a practical guide to *Generative AI for Enterprises*.

---

1. Scale Zeitgeist: AI Readiness Report, a survey of more than 1,600 executives and ML practitioners, 2023

2. McKinsey & Co, AI-powered marketing and sales reach new heights with Generative AI, 2023

I strongly recommend Mr. Anand's book as a go to resource for any decision makers on Generative AI strategy. I hope you enjoy this journey that introduces core elements of Generative AI in the context of enterprise solutions and architectural considerations at scale.

*Maja Vukovic*

*AI for Application Modernization,  
IBM Research*

**Disclaimer:** Views presented in the book are entirely his own and not necessarily those of IBM.

## About the Author

**Vishal** has 21+ years of ingenuity experience in the field of IT performing various technical leadership roles across the globe spanning from delivery, architecture, solutioning, offerings, product engineering, and hybrid cloud transformation. He works with IBM as Global Chief Technologist for hybrid cloud and Generative AI transformation, serving global clients across the industries. He is an IBM Master Inventor with many inventions patented. He is also an Open Group Certified Distinguished Architect, IBM Certified Thought Leader Architect, Fellow of British Computer Society and Top Voice on AI.

**Disclaimer:** Views presented in the book are entirely his own, and not necessarily those of IBM.

## About the Reviewers

- ❖ **James Radtke** is a Senior Solutions Architect working for AWS, helping strategic accounts optimally deploy their workloads in the cloud. James brings decades of experience in several roles working with influential companies such as Sun Microsystems and Red Hat, primarily focused on infrastructure and automation. James has published or contributed to a number of AWS docs, guides, and projects involving Amazon Elastic Kubernetes Service, as well as related open source and CI/CD components. James has a passion for Linux, Hybrid/Edge cloud, and deploying resilient solutions anywhere there is a network. That passion has led to exploring AI/ML solutions at the Edge and with Kubernetes.
- ❖ **Rajiv Avacharmal** is a leading expert in the field of AI/ML risk management, with a particular focus on Generative AI. With a distinguished career spanning over 13 years, Rajiv has held senior leadership roles at several multinational banks and currently serves as the Corporate Vice President of AI and Model Risk at a leading Life Insurance Company. Rajiv's research interests lie at the intersection of AI/ML, risk management, and explainable AI.

## Acknowledgement

I want to express my deepest gratitude to my family for their unwavering support and encouragement throughout this book's writing, especially my parents.

I am also grateful to BPB Publications for their guidance and expertise in bringing this book to fruition. It was a long journey of revising this book, with valuable participation and collaboration of reviewers, technical experts, and editors.

I would also like to acknowledge the valuable contributions of my colleagues, clients, and mentors during many years working in the tech industry, who have taught me so much and provided valuable feedback on my work time to time.

Finally, I would like to thank all the readers who have taken an interest in my book and for their support in making it a reality. Your encouragement has been invaluable.

# Preface

Building enterprise class Generative AI solutions is a complex task that requires a comprehensive understanding of the enterprise ecosystem, latest technologies, and understanding of multi-dimensional requirements. Generative AI has become a strong pillar for growth in the modern enterprises.

This book is designed to provide a comprehensive guide to deliver Generative AI based solutions with robust decision making. It covers a wide range of topics, including the rise of Generative AI in enterprises, advanced concepts such as operating system, orchestration platform, operating model and the use of the landing zone for building robust and scalable applications powered by the language models.

Throughout the book, you will learn about the key requirements associated with Generative AI and language models, and how to use them for enterprise applications that are efficient, reliable, and easy to govern. You will also learn about best practices and deployment patterns for building enterprise class solutions and will be provided with broader practical wisdom to help you understand the concepts for enterprises.

This book is intended for decision makers who are the leaders and want to lead how to successfully implement Generative AI for enterprises. It is also helpful for experienced professionals who want to expand their knowledge of these technologies and improve their skills in building robust and reliable applications powered by Generative AI.

With this book, you will gain the knowledge and skills to become a proficient leader in the field of modern enterprise development using Generative AI. I hope you will find this book informative and helpful.

**Chapter 1: The Rise of Generative AI in Enterprises** - focusses on the transformative impact of Generative AI technology across various industry sectors. We will explore how businesses can integrate the same to enhance creativity, automate processes, and drive innovation.

Through an expert view, the chapter touches upon the background, philosophy, and prospects of Generative AI in the business world.

**Chapter 2: Complex Needs of Production** - dives into the complexities, exploring the multifaceted needs from model and data governance to transformation needs, that organizations must address to effectively harness the power of Generative AI. By unpacking these critical elements, we aim to provide a comprehensive guideline that helps enterprises not only implement but also thrive with Generative AI at the core of their production strategies, predominantly from the large language models' aspect.

**Chapter 3: Model Selection for Enterprises** – considers a range of critical factors and criteria that align with the strategic objectives and operational requirements. This chapter highlights those key considerations to ensure it meets business needs, crucial for maintaining credibility, utility in business applications, and its integration capabilities with the ecosystem, allowing for seamless adoption and minimal disruption. These criteria collectively ensure that the chosen LLM not only meets immediate needs but also adapts to future challenges and opportunities in the business landscape.

**Chapter 4: Model Deployment for Enterprises** - dives into deploying Generative AI large language models within the enterprise sector, a task both intricate and rewarding. We will go through the various deployment patterns. Each pattern offers distinct advantages and challenges. This chapter not only maps out these patterns but also critically examines their pros and cons, aiding decision-makers in positioning them effectively within their strategic IT landscapes.

**Chapter 5: Operating System for Enterprises** - provides a pivotal framework for robust decision making to redefine the boundaries of technical architectures to achieve organizational goals. As part of the operating system in the context of GenAI, it is imperative to not only understand the boundaries but also the requirements within the corporate sphere. As we go through this chapter, the focus shifts to those practical boundaries of Generative AI in deployment contexts. Moreover, we will address the critical challenges and considerations



that enterprises must navigate when integrating Generative AI into their landing zones.

**Chapter 6: Prompt Engineering for Enterprises** - explores the fundamental principles of prompt engineering and its strategic deployment within the enterprise application development landscape. Initially, the concept is introduced in its most simplistic form, laying the groundwork for understanding how these prompts serve as a critical interface between human operators and sophisticated AI systems. As we transition from theory to practice, the focus shifts to the intricate process of integrating prompt engineering into the enterprise environment.

**Chapter 7: Fine-tuning for Enterprises** - explores the transformative practice of fine-tuning language models to meet specific business needs. Following the theoretical exploration, the chapter provides a practical, hands-on demonstration on how it looks like in real enterprise environment. This chapter delves into the intricate process of fine-tuning, a method where a pre-trained model is further trained - usually on a smaller, task-specific dataset - to refine its capabilities for specialized tasks. Through fine-tuning, these models, which initially acquire a broad understanding of language from extensive data, are meticulously tailored to exhibit improved performance in specific contexts, thereby bridging the gap between general linguistic ability and specialized knowledge.

**Chapter 8: Orchestration of Generative AI Workflows** - delves into the sophisticated realm of automated orchestration platform that streamlines the integration of data, model related tasks, ecosystem, and continuous improvement. As we explore, we will uncover what capabilities the platform should bring to simplify the complexities inherent in managing AI projects and enhance the efficiency and effectiveness of generative models. Through a detailed description, we will see how these systems facilitate seamless interaction between various AI components and applications, thereby fostering a conducive environment for innovation and scalability – across landing zones.

**Chapter 9: Six Ethical Dimensions for Enterprises** - accountability, fairness, transparency, human-orientation, security, and lawfulness, emerge as crucial factors in the development and deployment of Generative AI solutions. These principles serve not only as a moral compass but also as foundational guidelines that ensure AI technologies augment human capabilities without causing inadvertent harm. Together, these principles shape the scaffolding upon which the trustworthiness and utility of Generative AI rest, paving the way for its ethical integration into society for enterprises.

**Chapter 10: Designing a Target Operating Model** - an effective operating model for Generative AI is crucial for businesses that wish to harness this technology's full potential. An adeptly crafted operating model serves as a blueprint that guides the integration of AI with core business processes and strategies. It ensures that AI technologies are implemented in a way that aligns with organizational goals, enhances performance, and drives competitive advantage. A well-designed model addresses the technical, ethical, and governance challenges that come with AI deployment, thereby facilitating a sustainable and responsible adoption of AI within the enterprise.

**Chapter 11: Cost Optimization Strategies** - explores the cost optimization strategies for large language models within enterprise settings, exploring technological pre and post deployment approaches to maximizing the return on investment. Through a comprehensive analysis of cost factors, from computational demands to scaling efficiencies, this discussion aims to provide businesses with actionable insights on optimizing operational costs while harnessing the full potential of large language models.

**Chapter 12: Retrieval-augmented Generation for Enterprises** - explores the framework of RAG, its comparison with traditional AI, and its scalability within enterprise applications. This discussion highlights how RAG at scale can transform industries by providing AI systems that not only generate desired responses but also adapt to evolving informational landscapes.

**Chapter 13: Model as a Service for Enterprises** - represents a transformative approach in the deployment and consumption of large (or small) language models. By providing models over the cloud platform, MaaS enables enterprises to access and leverage sophisticated GenAI capabilities without the need for extensive in-house infrastructure or expertise. Key advantages include cost-efficiency, as it reduces the need for upfront capital investment in hardware and software, scalability, allowing companies to adjust usage based on demand, and agility, enabling rapid deployment and iteration of models to respond to evolving business needs.

**Chapter 14: Confidential AI** - explores the multifaceted challenges and pioneering solutions associated with securing sensitive data as it is processed across various layers. It explores the needs and solution, such as trusted execution environments (TEEs) for both virtual machines and graphical processing units, which provide stringent security measures essential for protecting data during processing phases.

**Chapter 15: Latency in Generative AI Solutions** - delves into the multifaceted nature of latency in Generative AI solutions, exploring its origins, implications, and the various strategies employed to minimize it. Understanding the sources and their impact is essential for developers and engineers striving to create more efficient and responsive AI systems. This chapter will explore various latency factors such as input network latency, hardware latency, model inference latency, integration point latency and output network latency.

**Chapter 16: Multi-modal Multi-agentic Assistant Framework for Enterprises**- delves into the evolving landscape of how sophisticated AI agents, equipped with multi-modal capabilities and collaborative intelligence, are set to revolutionize the way companies function. As these technologies advance, they promise to enhance efficiency, personalization, and decision-making processes, offering enterprises a competitive edge in an increasingly complex and dynamic market. By examining the potential of AI agents and multi-agent frameworks, this chapter provides a comprehensive insight into the impending shifts and opportunities that lie ahead for businesses willing to embrace the next wave of Artificial Intelligence.

## Coloured Images

Please follow the link to download the  
*Coloured Images* of the book:

**<https://rebrand.ly/kbgp9r1>**

We have code bundles from our rich catalogue of books and videos available at **<https://github.com/bpbpublications>**. Check them out!

## Errata

We take immense pride in our work at BPB Publications and follow best practices to ensure the accuracy of our content to provide with an indulging reading experience to our subscribers. Our readers are our mirrors, and we use their inputs to reflect and improve upon human errors, if any, that may have occurred during the publishing processes involved. To let us maintain the quality and help us reach out to any readers who might be having difficulties due to any unforeseen errors, please write to us at :

**[errata@bpbonline.com](mailto:errata@bpbonline.com)**

Your support, suggestions and feedbacks are highly appreciated by the BPB Publications' Family.

Did you know that BPB offers eBook versions of every book published, with PDF and ePub files available? You can upgrade to the eBook version at [www.bpbonline.com](http://www.bpbonline.com) and as a print book customer, you are entitled to a discount on the eBook copy. Get in touch with us at :

**[business@bpbonline.com](mailto:business@bpbonline.com)** for more details.

At **[www.bpbonline.com](http://www.bpbonline.com)**, you can also read a collection of free technical articles, sign up for a range of free newsletters, and receive exclusive discounts and offers on BPB books and eBooks.

## Piracy

If you come across any illegal copies of our works in any form on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at **business@bpbonline.com** with a link to the material.

## If you are interested in becoming an author

If there is a topic that you have expertise in, and you are interested in either writing or contributing to a book, please visit **www.bpbonline.com**. We have worked with thousands of developers and tech professionals, just like you, to help them share their insights with the global tech community. You can make a general application, apply for a specific hot topic that we are recruiting an author for, or submit your own idea.

## Reviews

Please leave a review. Once you have read and used this book, why not leave a review on the site that you purchased it from? Potential readers can then see and use your unbiased opinion to make purchase decisions. We at BPB can understand what you think about our products, and our authors can see your feedback on their book. Thank you!

For more information about BPB, please visit **www.bpbonline.com**.

## Join our book's Discord space

Join the book's Discord Workspace for Latest updates, Offers, Tech happenings around the world, New Release and Sessions with the Authors:

<https://discord.bpbonline.com>



# Table of Contents

<b>1. The Rise of Generative AI in Enterprises .....</b>	<b>1</b>
Introduction.....	1
Structure.....	1
Objectives.....	2
Background of Generative AI.....	2
Philosophy of Generative AI.....	4
Potential use cases across enterprises.....	5
Architectural overview .....	7
Conclusion.....	9
Key terms.....	9
<b>2. Complex Needs of Production.....</b>	<b>11</b>
Introduction.....	11
Structure.....	12
Objectives.....	12
Challenges of model-specific scaling.....	12
Dimensions of model at scale .....	13
<i>Business objectives</i> .....	14
<i>Model sourcing</i> .....	16
<i>Open GenAI LLMs</i> .....	16
<i>Closed-source GenAI LLMs</i> .....	16
<i>Data management</i> .....	17
<i>Model selection</i> .....	19
<i>Model evaluation</i> .....	21
<i>Model optimization</i> .....	23
<i>Model fine-tuning</i> .....	24
<i>Model orchestration</i> .....	25
<i>Model observability</i> .....	26
<i>Model governance</i> .....	28
LLMOps.....	29
Conclusion.....	30
Key terms.....	31
<b>3. Model Selection for Enterprises.....</b>	<b>33</b>
Introduction.....	33

Structure.....	33
Objectives.....	34
Importance of model selection .....	34
Criteria for selecting the model for enterprise needs.....	36
Conclusion.....	44
Key terms.....	44
<b>4. Model Deployment for Enterprises.....</b>	<b>45</b>
Introduction.....	45
Structure.....	45
Objectives.....	46
Deployment patterns for LLMs.....	46
Pros and cons of deployment patterns.....	51
Positioning of deployment patterns .....	53
Deployment strategies of AI applications powered by LLM...	54
Conclusion.....	57
Key terms.....	57
<b>5. Operating System for Enterprises.....</b>	<b>59</b>
Introduction.....	59
Structure.....	59
Objectives.....	60
Crux of operating system .....	60
Demarcation of operating system .....	63
Conclusion.....	70
Key terms.....	70
<b>6. Prompt Engineering for Enterprises.....</b>	<b>71</b>
Introduction.....	71
Structure.....	71
Objectives.....	72
Glimpse of prompt engineering .....	72
Enterprise view of prompt engineering .....	74
<i>Single input scenario</i> .....	75
<i>Multiple inputs scenario in terms of scaling</i> .....	78
Conclusion.....	82
Key terms.....	82
<b>7. Fine-tuning for Enterprises.....</b>	<b>83</b>
Introduction.....	83

Structure.....	84
Objectives.....	84
Crux of fine-tuning.....	84
Practical demonstration of fine-tuning .....	88
Conclusion.....	93
Key terms.....	94
<b>8. Orchestration of Generative AI Workflows .....</b>	<b>95</b>
Introduction.....	95
Structure.....	96
Objectives.....	96
Background .....	96
Orchestration platform .....	99
Orchestration platform and landing zone .....	102
<i>Cloud perspective</i> .....	103
<i>On-premises perspective</i> .....	105
Conclusion.....	106
Key terms.....	108
<b>9. Six Ethical Dimensions for Enterprises .....</b>	<b>109</b>
Introduction.....	109
Structure.....	110
Objectives.....	110
Correlated dimensions.....	110
Importance of the dimensions .....	113
<i>Responsible infusion</i> .....	114
<i>User centricity</i> .....	114
<i>Guardrails</i> .....	115
<i>Governance</i> .....	115
<i>Communication</i> .....	117
Conclusion.....	119
Key terms.....	120
<b>10. Designing a Target Operating Model.....</b>	<b>121</b>
Introduction.....	121
Structure.....	122
Objectives.....	122
Holistic operating model.....	122
Seven layers.....	124
<i>Layer 1: Actor layer</i> .....	125



<i>Layer 2: Requirements layer</i> .....	126
<i>Layer 3: Consumption layer</i> .....	129
<i>Layer 4: Platform layer</i> .....	130
<i>Layer 5: Execution layer</i> .....	130
<i>Layer 6: Landing zone layer</i> .....	132
<i>Layer 7 - Applicability layer</i> .....	133
<i>Feedback loop</i> .....	134
Conclusion .....	135
Key terms .....	136
<b>11. Cost Optimization Strategies</b> .....	<b>137</b>
Introduction .....	137
Structure .....	137
Objectives .....	138
Background of cost optimization opportunities .....	138
Inside periphery .....	140
Conclusion .....	145
Key terms .....	145
<b>12. Retrieval-augmented Generation for Enterprises</b> .....	<b>147</b>
Introduction .....	147
Structure .....	147
Objectives .....	148
GenAI with RAG versus traditional AI .....	148
Retrieval-augmented generation .....	150
Retrieval-augmented generation at scale .....	151
<i>Infrastructure</i> .....	152
<i>Data</i> .....	152
<i>Integration</i> .....	153
<i>Observability</i> .....	153
<i>Performance</i> .....	154
<i>Prompt and fine-tuning</i> .....	154
<i>Technology selection</i> .....	155
Conclusion .....	159
Key terms .....	159
<b>13. Model as a Service for Enterprises</b> .....	<b>161</b>
Introduction .....	161
Structure .....	161
Objectives .....	162

Model as a Service .....	162
<i>Architecture</i> .....	162
<i>MaaS Studio</i> .....	163
Decision quadrants .....	166
Implications and mitigations .....	168
Conclusion .....	170
Key terms .....	171
<b>14. Confidential AI .....</b>	<b>173</b>
Introduction .....	173
Structure .....	173
Objectives .....	174
Background of confidential AI .....	174
Technicality under the hood .....	179
Conclusion .....	184
Key terms .....	185
<b>15. Latency in Generative AI Solutions .....</b>	<b>187</b>
Introduction .....	187
Structure .....	187
Objectives .....	188
Background .....	188
Latency factors .....	189
<i>Factor 1: Input network latency</i> .....	190
<i>Factor 2: Hardware latency</i> .....	192
<i>Factor 3: Model inference latency</i> .....	193
<i>Factor 4: Integration point latency</i> .....	195
<i>Factor 5: Output network latency</i> .....	197
Conclusion .....	197
Key terms .....	198
<b>16. Multi-modal Multi-agentic Assistant Framework for Enterprises ..</b>	<b>199</b>
Introduction .....	199
Structure .....	199
Objectives .....	200
Future of AI agents .....	200
Multi-modal multi-agentic assistant framework .....	201
Conclusion .....	204
Key terms .....	204
<b>Index .....</b>	<b>205-209</b>

# CHAPTER 1

# The Rise of Generative AI in Enterprises

## Introduction

In this chapter, we will focus on the transformative impact of Generative AI technology across various industry sectors. We will explore how businesses can integrate the same to enhance creativity, automate processes, and drive innovation. Through an expert view, the chapter touches upon the background, philosophy, and prospects of Generative AI in the business world.

## Structure

In this chapter, we will explore the following topics:

- Background of Generative AI
- Philosophy of Generative AI
- Potential use cases across enterprises
- Architectural overview

## Objectives

Readers will be able to visualize the underlying philosophy of Generative AI technology and explore potential use cases. The objective of this chapter is to provide an exploration for the readers on Generative AI within the enterprise sector, examining its philosophical underpinnings, architectural overview, and the potential use cases it presents. This aims to equip readers with an understanding of how Generative AI technologies have emerged as transformative tools in business and the profound impact they are anticipated to have across various industries. It will set the stage for providing further valuable insights for the industry leaders and technologists looking to harness these tools for strategic advantage.

## Background of Generative AI

**Generative Artificial Intelligence (GenAI)** is rapidly gaining prominence within enterprises, reshaping how businesses operate and innovate. Industry leaders today anticipate that Generative AI will transform their organizations within the next few years and express confidence in their overall strategy and technology infrastructure for Generative AI adoption. Generative AI is no longer a niche trend. It has become a top priority for the C-suite across various sectors, including healthcare, life sciences, legal, financial services, and the public sector. One of the top technological and consulting firms says more than 80% of enterprises will have used Generative AI APIs or deployed generative AI-enabled applications within this decade. The rise of Generative AI in enterprises is driven by technological advancements, business needs, and the promise of transformative impact across various sectors.

Generative AI, a branch of Artificial Intelligence focused on generating new content, be it text, images, audio, video, or code, has undergone significant evolution, reshaping the landscape of enterprise applications and capabilities. This technology, powered by sophisticated Machine Learning models, has transitioned from a novel concept to a pivotal tool for businesses across various sectors. Its latest evolution has not only enhanced the efficiency and creativity within enterprises but has also introduced novel ways of problem-solving and customer engagement.

At the heart of this transformation lies the progression of ML models, particularly deep learning algorithms, which have become adept at understanding and generating complex patterns. These models are

trained on vast datasets, enabling them to produce desired outputs for human and machines' consumption. For enterprises, this capability has opened up new avenues for content creation, from auto-generating marketing materials to developing more engaging and personalized customer experiences.

Generative AI can produce a wide range of content, including articles, reports, and marketing copy, tailored to the specific tone and style of a brand. This not only streamlines content creation processes but also ensures a consistent brand voice across all communications. For example, in customer service, AI-generated responses can handle various inquiries, providing instant, 24/7 support to customers. This level of responsiveness and personalization enhances customer satisfaction and loyalty, crucial metrics for any enterprise.

Moreover, Generative AI has revolutionized product design and development. By harnessing the power of these AI models, companies can explore a broader spectrum of design possibilities in a fraction of the time it would take through traditional methods. This rapid prototyping accelerates the innovation cycle, enabling companies to respond more swiftly to market trends and consumer preferences. For industries such as automotive and fashion, where design plays a pivotal role, this capability is invaluable.

In addition to enhancing creativity and efficiency, Generative AI is also redefining data analysis and decision-making processes within enterprises. AI models can generate simulations and predictive models based on existing data, allowing businesses to anticipate market changes, understand consumer behaviour, and make informed decisions. This foresight can be a significant competitive advantage, enabling proactive rather than reactive strategies.


However, the integration of Generative AI into enterprise operations is not without its challenges. Issues related to data privacy, ethical considerations, and the potential for generating misleading or harmful content necessitate a careful and considered approach to deployment. As such, enterprises must establish robust guidelines and ethical frameworks to govern the use of these technologies.

The latest evolution of Generative AI presents both remarkable opportunities and significant challenges for enterprises. By harnessing its capabilities, businesses can unlock unprecedented levels of creativity, efficiency, and insight, positioning themselves at the forefront of innovation. However, this journey requires a careful balance between

leveraging the potential of AI and maintaining ethical standards and human oversight. As Generative AI continues to evolve, it will undoubtedly remain a key driver of enterprise transformation, shaping the future of industries in ways we are only beginning to imagine.

# Philosophy of Generative AI

Before we go further into the details, it is imperative to visualize how it works. The philosophy of Generative AI explains the fundamental reality of its real-world applications and integration. The following diagram illustrates the dynamic workflow of Generative AI, showcasing how input prompts are processed to produce various types of output completions, which can then be seamlessly integrated into broader business processes to enhance operational efficiency and innovation:



*Figure 1.1: Input and output*

The process begins with an input prompt. This can be a question, a partial sentence, or any form of textual, visual, or aural input. The input prompt serves as the seed for the generative process.

The central black box represents generative AI. It embodies the power of algorithms, neural networks, and deep learning models. Generative AI takes the input prompt and generates novel content based on learned patterns and representations.

The output completion points to different types of data.

Generative AI can create human-like text, generate stories, and even compose poetry. It can generate code snippets, complete functions, and assist in software development. Generative AI can create visual content, such as art, architecture diagrams, and imaginative scenes. It can simulate video sequences, generate animations, and enhance multimedia experiences. It composes music, generates speech, and enhances audio content.