
Contents

Preface	xv
Devasis PRADHAN	
Acknowledgments	xvii
Devasis PRADHAN, Hla MYO TUN, David BANJERDPONGCHAI and Prasenjit CHATTERJEE	
Chapter 1. Real-Time Big Data Processing in the Cloud: Scalable, Cost-Efficient, AI-Driven Solutions for Financial Analytics	1
Saravanan Thirumazhisai PRABHAGARAN, Ishu Anand JAISWAL and Hina GANDHI	
1.1. Introduction	1
1.2. The evolution of big data in finance	2
1.3. Architectural components of real-time analytics	3
1.4. Real-time ingestion and stream processing	4
1.5. Cost optimization strategies in cloud infrastructure	6
1.5.1. Avoiding the “always-on” trap	6
1.5.2. Using storage tiers wisely	6
1.5.3. Picking the right compute model	7
1.5.4. Watching out for data transfer fees	7
1.5.5. Letting AI watch the bills	7
1.6. Case studies and applications in financial analytics	8
1.6.1. Case study 1: real-time fraud detection	8
1.6.2. Case study 2: personalized customer offers.	8
1.6.3. Case study 3: faster trading decisions	8
1.7. How data enters the pipeline	9
1.8. Processing in real-time	9
1.9. Storing the results	9
1.10. Scaling without blowing the budget	10
1.11. A few examples	10

1.12. Additional implementation stories	11
1.12.1. Example: a mid-sized bank modernizing batch systems	11
1.12.2. Example: a serverless fintech startup	11
1.12.3. Example: cross-region replication gone wrong	12
1.13. Common pitfalls and how to avoid them	14
1.14. Lessons learned.	14
1.15. Where the model sits in the pipeline	15
1.16. Example: fraud model in production	15
1.17. How teams deploy models.	16
1.18. Retraining and model drift.	16
1.19. Latency considerations.	17
1.20. A few more lessons	19
1.21. Why people do it anyway	19
1.22. What usually goes wrong	19
1.23. How some teams made it work	20
1.24. A few stories	21
1.25. Conclusion	23
1.26. References	23

Chapter 2. Scalable Microservices and Data Integration for Modern eCommerce and Enterprise Systems. 27

Aneeshkumar P. SUNDESWARAN, Dilip Prakash VALANARASU
and Anaswara Thekkan RAJAN

2.1. Introduction	28
2.2. Literature review.	32
2.3. Methodology	35
2.3.1. Illustration of the result.	39
2.3.2. Design and experimental measures	41
2.4. Results	42
2.4.1. Comparison of graphs and table in throughput versus concurrency.	42
2.4.2. Graph and table survey of latency percentiles	42
2.4.3. Deep-seated conditions that influence throughput	43
2.4.4. Latency improving forces	44
2.4.5. Business and eCommerce (SSI)..	45
2.5. Conjunction with former work	45
2.6. Future direction	46
2.6.1. AI-powered orchestration and predictive autoscaling	46
2.6.2. End-to-end governance automation by policy-as-code	46
2.6.3. Lightweight integration and edge-native microservices	47
2.6.4. Data integration low-code/no-code	47
2.6.5. Scalable benchmarking and full-featured observability suites	48

2.7. Conclusion	48
2.8. References	51

Chapter 3. Modern Data Engineering on GCP: Scalable Pipelines, Real-Time Processing and Governance. 55

Jagadeesh THIRUVEEDULA and Souvari Ranjan BISWAL

3.1. Introduction	55
3.1.1. GCP and modern data engineering	56
3.2. Designing scalable pipelines	57
3.2.1. Serverless architecture	57
3.2.2. Event-driven pipelines	57
3.2.3. Partitioning and sharding.	58
3.2.4. Auto-scaling and load balancing.	58
3.2.5. Cloud composer to orchestrate workflows	58
3.3. Real-time data processing.	58
3.3.1. Pub/Sub + dataflow streaming.	58
3.3.2. Windowing and triggers	58
3.3.3. Stateless versus stateful processing	58
3.4. Data governance and compliance	59
3.4.1. Data lineage and metadata management	60
3.4.2. Identity and access control	60
3.4.3. Data protection and encryption	60
3.4.4. Compliance certifications	61
3.5. Case study: real-time fraud detection	61
3.6. Advanced patterns and best practices	61
3.6.1. Lambda versus Kappa architectures.	62
3.6.2. Data lakehouse integration	62
3.6.3. Microservices and data mesh	62
3.7. Monitoring and observability	63
3.7.1. Pipeline monitoring.	63
3.7.2. Data quality monitoring	63
3.8. Cost optimization strategies.	63
3.8.1. Storage management	64
3.8.2. Resource right-sizing	64
3.8.3. Query optimization	64
3.9. ML integration	64
3.9.1. Inference of real-time ML	64
3.9.2. MLOps and retraining	64
3.10. Looking ahead	64
3.11. Final thoughts.	65
3.12. References	65

Chapter 4. AI for Personalization and Security in Digital Travel Platforms	67
Kartheek DOKKA, Saravanan Thirumazhisai PRABHAGARAN and Maheswari GOVINDARAJU	
4.1. Introduction	67
4.2. Evolution of AI in travel and digital services	70
4.2.1. Early developments: rule-based systems and static models	70
4.2.2. Expansion through machine learning and data availability	70
4.2.3. Emergence of deep learning and contextual AI	71
4.2.4. Real-time and predictive intelligence	72
4.2.5. Integration with the Internet of Things and edge computing	72
4.2.6. Convergence of personalization and security.	73
4.3. AI-driven personalization techniques	74
4.3.1. Collaborative and content-based filtering.	74
4.3.2. Deep learning models.	74
4.3.3. RL for real-time adaptation	75
4.3.4. NLP and conversational AI	75
4.3.5. Contextual and predictive personalization	75
4.4. AI in platform and transactional security	76
4.4.1. Behavioral biometrics and user profiling	76
4.4.2. AI for anomaly and intrusion detection	76
4.4.3. Payment fraud prevention	77
4.4.4. Bot and automation attack detection	77
4.4.5. Identity and document verification	77
4.5. Integration of AI in travel platforms	78
4.5.1. Dynamic pricing and revenue optimization.	78
4.5.2. AI-powered customer support and chatbots	79
4.5.3. Personalized travel planning and recommendations	79
4.5.4. Fraud-resistant booking systems.	80
4.5.5. AI-enhanced identity verification	80
4.6. Challenges and ethical considerations	81
4.6.1. Data privacy and consent management	81
4.6.2. Algorithmic bias and discrimination	82
4.6.3. Lack of explainability	82
4.6.4. Security vulnerabilities and adversarial threats.	83
4.6.5. Regulatory uncertainty and compliance burdens.	83
4.7. Future outlook and strategic implications	85
4.7.1. Hyper-personalization through predictive intelligence	85
4.7.2. Federated learning and privacy-centric architectures	85
4.7.3. Real-time identity and zero-trust security models	86
4.7.4. Cross-platform ecosystem integration.	86
4.7.5. Governance, transparency and responsible AI	86

4.8. Conclusion	87
4.9. References	88

**Chapter 5. Modernizing Legacy E-commerce Platforms:
Cloud Migration, DevOps and API-First Strategies.** 91

Dilip Prakash VALANARASU and Arnab KAR

5.1. Introduction	92
5.2. Understanding legacy e-commerce architectures	94
5.2.1. Characteristics of legacy systems	94
5.2.2. Operational constraints	94
5.2.3. Integration challenges	95
5.2.4. Scalability and performance bottlenecks	95
5.2.5. Organizational impact	96
5.2.6. Technical debt and business inertia	96
5.3. Cloud migration strategies for legacy systems	97
5.3.1. Drivers for cloud migration	97
5.3.2. Common cloud migration models	98
5.3.3. Migration frameworks and toolchains.	99
5.3.4. Security and compliance considerations	99
5.3.5. Challenges in cloud migration	100
5.3.6. Strategic best practices	100
5.4. DevOps as a catalyst for e-commerce agility	101
5.4.1. The DevOps philosophy in e-commerce	101
5.4.2. Tooling and automation	102
5.4.3. DevOps and legacy system constraints	102
5.4.4. Organizational benefits and metrics.	103
5.4.5. DevSecOps and security automation	103
5.5. API-first strategy: modularizing e-commerce architectures	105
5.5.1. Fundamentals of API-first design	105
5.5.2. Enabling decoupled architecture.	106
5.5.3. Integration with ecosystem services.	106
5.5.4. Improving developer and operational efficiency	107
5.5.5. Challenges in API-first modernization	107
5.5.6. Strategic implementation practices	108
5.6. Security, observability and compliance in modern platforms	108
5.6.1. Security in decentralized, API-driven architectures	108
5.6.2. Observability and platform intelligence.	109
5.6.3. Compliance and data governance in cloud platforms	110
5.6.4. DevSecOps and security automation	110
5.6.5. Security metrics and governance models	111
5.7. Strategic roadmap and future directions	111
5.7.1. Phased modernization strategy.	111
5.7.2. Platform intelligence and hyper-personalization	112

5.7.3. Edge computing and low-latency architectures	112
5.7.4. Future of compliance and digital trust.	112
5.7.5. Ecosystem participation and composability	113
5.7.6. Talent, skills and cultural shifts	113
5.8. Conclusion	113
5.9. References	114

Chapter 6. Building Scalable and Resilient Stream Processing Systems with Kafka and Serverless Tools 119

Vamsi Krishna KOGANTI and Naveen Saikrishna PUPPALA

6.1. Introduction: the evolution of stream processing in cloud-native systems	119
6.2. Apache Kafka: core principles and its role in stream ingestion	121
6.3. Serverless computing paradigms in stream processing.	123
6.4. Architecting scalable stream processing pipelines	126
6.5. Resilience and fault tolerance in distributed stream systems.	129
6.6. Performance optimization, observability and cost management.	132
6.7. Conclusion: future directions for serverless and Kafka-based architectures	134
6.8. References	136

Chapter 7. AI-Driven Automation for Scalable Media, Data and Search Systems 139

Aatishkumar DHAMI and Maheswari GOVINDARAJU

7.1. Introduction	139
7.2. The rise of automated data pipelines	141
7.3. Scaling media processing without losing control	142
7.4. The art and trouble of real-time search systems.	143
7.5. Building robust, automated data pipelines.	147
7.6. Machine learning in analytics and personalization	148
7.7. Monitoring and keeping systems trustworthy	149
7.8. Training methods for sorting	151
7.9. Checks and alerts	152
7.10. Manual controls and overrides	152
7.11. Monitoring and cost controls	154
7.12. Fast updates and small delays	155
7.13. Checks that made a difference	155
7.14. When costs rose fast	156
7.15. Adjustments made during high-volume events	156
7.16. Conclusion	158
7.17. References	158

Chapter 8. Scalable and Cost-Efficient Sap Cloud Migration and Operations for Large Enterprises 163

Ankit Kumar GUPTA

8.1. Introduction	163
8.1.1. Context and importance	164
8.1.2. Challenges in SAP cloud migrations	164
8.1.3. Objectives	165
8.2. Background and related work.	166
8.2.1. SAP system landscape evolution	166
8.2.2. Cloud computing paradigms for SAP	167
8.2.3. Review of existing migration strategies.	167
8.2.4. Gaps in current approaches	169
8.3. Drivers and benefits of SAP cloud migration	170
8.3.1. Scalability demands in large enterprises	170
8.3.2. Cost-efficiency motivations	171
8.3.3. Security and compliance considerations	171
8.3.4. Competitive advantages	171
8.4. Migration planning framework	172
8.4.1. Assessment phase	172
8.4.2. Readiness analysis phase	173
8.4.3. Cloud platform selection phase	173
8.4.4. Cost-benefit analysis (CBA) phase	173
8.5. Migration execution strategies	174
8.6. Post-migration optimization	175
8.7. Case studies	179
8.7.1. Siemens AG: transforming global SAP landscapes	179
8.7.2. Coca-Cola European Partners: accelerating analytics in the beverage industry	180
8.7.3. BP: digitally transforming energy operations.	180
8.8. Challenges, risks and future directions.	181
8.9. Conclusion	184
8.10. References	184

Chapter 9. Integrative Morphological Assessment of Ayurvedic Haritaki (*Terminalia chebula* Retz.) Varieties: A Comparative Study Across Indian Regions and the Western Ghats. 191

Komal TAYADE, Elingbam SINGH, Varsha GHATE,
Suresh JAGTAP and Rakshanda PATIL

9.1. Introduction	192
9.2. Materials and methods	194
9.2.1. Study design and conceptual framework	194
9.2.2. Literature review and classical correlation	195

9.2.3. Sampling site and collection of specimens	196
9.2.4. Morphological characterization and analysis	196
9.2.5. Data integration and interpretation	196
9.3. Results and discussion.	197
9.3.1. AI-based image classification	202
9.4. Conclusion	205
9.5. References	205

Chapter 10. Smart Analytics in SAP: Streamlining Big Data for Faster, Smarter Business Decisions 207

Vaidheyar Raman BALASUBRAMANIAN

10.1. Introduction	208
10.1.1. Technological architecture of smart analytics in SAP	209
10.1.2. SAP BusinessObjects (BOBJ)	209
10.1.3. SAP analytics cloud (SAC).	209
10.1.4. SAP S/4HANA embedded analytics.	210
10.1.5. Data integration tools	210
10.1.6. Data governance and security	212
10.1.7. Openness and interoperability	212
10.2. Applications and use cases across industries.	212
10.2.1. Manufacturing: predictive maintenance	212
10.2.2. Retail: real-time customer intelligence	213
10.2.3. Healthcare: patient-centered analytics	213
10.2.4. Financial services: risk and fraud analysis	213
10.2.5. Public sector: smart governance	213
10.2.6. Energy sector: grid optimization	213
10.2.7. Communication networks: decision analytics for intelligent operations.	214
10.3. Tools, functionalities and AI/ML integration in SAP smart analytics	215
10.3.1. SAP analytics cloud (SAC): unified intelligence platform.	215
10.3.2. SAP HANA predictive analysis library (PAL)	215
10.3.3. SAP automated predictive library (APL)	216
10.3.4. SAP AI Core and SAP AI Launchpad	216
10.3.5. Embedded AI in SAP S/4HANA.	218
10.3.6. SAP digital assistant and natural language processing	218
10.3.7. Explainable AI (XAI) and model governance.	218
10.4. Implementation challenges and best practices for SAP smart analytics	219
10.4.1. Key implementation challenges	219
10.4.2. Best practices for successful implementation	221
10.4.3. Organizational roles and responsibilities	223
10.4.4. Summary table: challenges and solutions	223

10.5. Future scope	223
10.5.1. Integration of generative AI (GenAI)	223
10.5.2. Real-time, AI-powered ERP systems	224
10.5.3. Augmented analytics and decision intelligence	224
10.5.4. Unified data layer with SAP DataSphere	224
10.5.5. Edge Analytics and Industry 4.0	225
10.5.6. Sustainability analytics and ESG reporting	225
10.6. Conclusion	226
10.7. References	226
Chapter 11. AI-Powered Feedback Loops in Digital Advertising	229
Divij PASRIJA	
11.1. Introduction	229
11.2. Understanding feedback loops in digital advertising	232
11.2.1. The concept of feedback loops	234
11.2.2. Components of a feedback loop in digital advertising	234
11.2.3. Traditional versus AI-powered feedback loops	235
11.2.4. Real-time bidding and programmatic advertising	235
11.2.5. Personalization through feedback loops	236
11.2.6. User engagement metrics as feedback	236
11.3. Role of AI in enhancing feedback mechanisms	237
11.3.1. From automation to intelligence	237
11.3.2. Deep learning for granular insights	238
11.3.3. Reinforcement learning for real-time optimization	238
11.3.4. Natural language processing in ad feedback	239
11.3.5. Computer vision for creative optimization	239
11.3.6. Multi-modal learning and fusion models	240
11.3.7. Predictive analytics and forecasting	240
11.3.8. User segmentation and behavioral clustering	241
11.4. Architectures and models for AI-driven ad feedback loops	241
11.5. Real-world applications	244
11.6. Challenges, ethical considerations and future directions	248
11.7. Conclusion	251
11.8. References	251
List of Authors	257
Index.	261