Autonomous Networks

Using AI to Advance Operations in SP and Enterprise Domains



STEFAN-ALEXANDRU MANZA
JOSH HALLEY

ciscopress.com

FREE SAMPLE CHAPTER

Autonomous Networks: Using AI to Advance Operations in SP and Enterprise Domains

Stefan-Alexandru Manza, CCIE No. 15687
Josh Halley, CCIEx3 No. 11924

Cisco Press

Autonomous Networks: Using AI to Advance Operations in SP and Enterprise Domains

Stefan-Alexandru Manza, Josh Halley

Copyright© 2027 Cisco Systems, Inc.

Published by:
Cisco Press
Hoboken, New Jersey

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms, and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit https://www.pearson.com/global-permission-granting.html.

No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

Please contact us with concerns about any potential bias at www.pearson.com/en-us/report-bias.html.

$PrintCode

Library of Congress Control Number: 2026938972

ISBN-13: 978-0-13-547336-8
ISBN-10: 0-13-547336-5

Warning and Disclaimer

This book is designed to provide information about achieving network autonomy in enterprise and service provider domains, by leveraging automation, observability, and AI. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied.

The information is provided on an “as is” basis. The authors, Cisco Press, and Cisco Systems, Inc. shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it.

The opinions expressed in this book belong to the authors and are not necessarily those of Cisco Systems, Inc.

Trademark Acknowledgments

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Cisco Press or Cisco Systems, Inc., cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Feedback Information

At Cisco Press, our goal is to create in-depth technical books of the highest quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the unique expertise of members from the professional technical community.

Readers' feedback is a natural continuation of this process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at feedback@ciscopress.com. Please make sure to include the book title and ISBN in your message.

We greatly appreciate your assistance.

Please contact us with concerns about any potential bias at https://www.pearson.com/report-bias.html.

Head of IT & Professional Learning, Enterprise Learning and Skills: Julie Phifer

Alliances Manager, Cisco Press: Caroline Antonio

Executive Editor: James Manly

Managing Editor: Sandra Schroeder

Development Editor: Ellie C. Bru

Senior Project Editor: Mandie Frank

Copy Editor: Chuck Hutchinson

Technical Editors: Nico Hambauer, Jochen Heidl

Designer: Chuti Prasertsith

Composition: codeMantra

Indexer: Timothy Wright

Proofreader: Barbara Mack



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

About the Authors

Stefan-Alexandru Manza, CCIE No. 15687, is a Distinguished Architect in the office of the CTO at Cisco Systems. In his current role, he focuses on strategic adoption of new technology trends such as observability, automation, and AI within Cisco and by Cisco's key enterprise and SP customers.

Over the years Stefan has specialized in leading large and complex technology transformations with customers within the EMEA region, covering a diverse set of technical domains like licensed radio access, IP transport, private data center, and Mobile Packet Core. His efforts include ensuring the associated process and people skills evolutions are properly executed to ensure the successful adoption of the new technology.

Stefan has a PhD in systems engineering from The University Politehnica of Bucharest and an MBA from Vlerick Business School in Leuven; he also has completed postgraduate studies in big data and analytics at KU Leuven and has an active C-level presence at international forums such as Cisco Live, Mobile World Congress, and TM Forum DTW.

Josh Halley, CCIEx3 No. 11924, is a Principal Architect in the office of the CTO at Cisco Systems, where his current role focuses on domains related to up-and-coming and burgeoning new technology trends and innovation, engaging in large and complex strategic negotiations, and working with C-Levels to set technology strategy and direction within their organizations. Over the years Josh has worked in many differing roles at organizations, ranging from technology companies to banking and finance and management consulting.

Within Cisco, Josh has worked side by side with many of the company's business units, product managers, and technical marketing engineers and has been directly involved in the creation and deployment of many new software features across multiple technology domains seen in Cisco's products and portfolio today.

Today Josh maintains his focus on technology, driving innovation with customers in the field and actively participating in the creation of patents, whitepapers, and new product and service offerings for Cisco. Further to his work within Cisco, Josh is also an open-source software advocate coleading the Cloud Native Computing Foundation (CNCF) Artificial Intelligence Technical Community Group, and is one of the coauthors of the Cisco Virtual Kubelet CNCF project, providing him an opportunity to share his knowledge and expertise with the wider open-source community.

About the Technical Reviewers

Nico Hambauer is a PhD candidate at the University of Regensburg, Germany, part of a government-funded research group on White-Box AI. His work focuses on benchmarking and designing novel, explainable machine learning models. He studies explainable approaches such as generalized additive models or cluster-then-predict methods. The presumption that models that are inherently explainable must sacrifice predictive performance is challenged and remedied by his work. His applications target high-stakes domains such as healthcare, customer analytics, and industrial business contexts. His research has been published in renowned outlets, including the European Conference on Information Systems, the Conference on Wirtschaftsinformatik, the Hawaii International Conference on System Sciences, and Business & Information Systems Engineering.

Jochen Heidl is a Principal Architect within Cisco's Customer Experience organization for service providers in the EMEA. In this capacity, he leads the advancement of next-generation operational technologies, guiding innovations from initial incubation and early-stage trials through to full-scale deployments. Jochen has partnered with leading operator accounts worldwide, delivering transformational projects in automation and assurance system integration. His current focus is driving the evolution toward autonomous networks, with a particular interest in translating AI concepts into practical AIOps solutions that address specific customer needs and deliver measurable value. Jochen is based in Frankfurt, Germany.

Dedications

Stefan: I would like to dedicate this book to my wife, Ann; and children, Pauline, Julia, and Alexandru, who supported me through my book-writing journey, including holidays, weekends, and late evenings.

Josh: I would like to dedicate this book to my loving wife and children. Thank you for your extra support and patience during my first book, this one…and maybe another one yet to come.

Acknowledgments

Stefan: Thank you, Josh, for your partnership and seamless collaboration, which were key for this book's success. Thank you, Jochen and Nico, for your detailed reviews and significant time spent to make the quality of the content better. Thank you, Oleksandr Vystoropskyi, for your key AI algorithms and implementation details in Chapter 4. Thank you, Roque Gagliano, for your feedback and suggestions on Chapter 8, and Martin Conway for your valuable input and improvement suggestions on operational process improvements in Chapter 10. Thank you, Michael Kaemper, Adele Trombetta, Pablo Steiner and Brian Meaney for your continued support. A special thank you to the TM Forum organisation who contributed with key material ensuring this book alignment with industry standards and trends. Finally, I would like to thank my colleagues and the customers with whom I collaborated in implementing network transformation programs over the past few years: Working with you has given me incredible insight, which has made this book possible.

Josh: First and foremost, I would like to send a big THANK YOU to Nico and Jochen for going to extended lengths well beyond your day jobs to carefully review and share feedback on the content. Stefan, my coauthor, thank you for your time, patience, discussion, and collaboration in bringing everything together. James, Ellie, and the whole Pearson team, thanks for translating what I sent over to you in indecipherable Australian into eloquent English. I would also like to send a special thanks to Morten Skriver for your observability input in Chapter 3, Eleni Grosdouli for your cloud native automation input in Chapter 6, Lukas Marschhausen for your perspectives on generative AI concepts in Chapter 5, and Gary Palmer and Ariel Leza for your Web3 insights in Chapter 9. Furthermore, I would like to thank my leadership team, Michael Kaemper and Adele Trombetta, for their support.

Contents at a Glance

Introduction xxiii

Chapter 1 Network Autonomy 1

Chapter 2 Regulations Impact on Autonomous Networks 59

Chapter 3 Observability and Data Sources 115

Chapter 4 Conventional and Task-Specific AI/ML Models Usage for Autonomous Networks 167

Chapter 5 Generative AI in the Context of Autonomous Networks 221

Chapter 6 Automated Execution and Testing 281

Chapter 7 Closed-Loop Command and Control 345

Chapter 8 Communication Service Provider Autonomous Network Vision 401

Chapter 9 Enterprise Domain Vision 459

Chapter 10 Autonomous Network Considerations—Organizational Change 509

Index 575

Reader Services

Register your copy at www.ciscopress.com/title/ISBN for convenient access to downloads, updates, and corrections as they become available. To start the registration process, go to www.ciscopress.com/register and log in or create an account*. Enter the product ISBN 9780135473368 and click Submit. When the process is complete, you will find any available bonus content under Registered Products.

*Be sure to check the box that you would like to hear from us to receive exclusive discounts on future editions of this product.

Contents

Introduction xxiii

Chapter 1 Network Autonomy 1

Introduction 1

State of Network Automation Today 2

Network Autonomy 3

Business Drivers 4

Observation 8

Decision 9

Execution 11

MLOps 11

Criticality Over Time 12

Automation Options 16

Automation Methods 20

Automation Tools 24

Data Model-Driven Automation 25

Continuous Integration (CI) and Continuous Delivery (CD) 27

Industry Frameworks and Trends 27

Service Providers 29

Enterprise 32

Effectively Managing and Dealing with Data 34

Challenges in Achieving Autonomy 39

Platforms 40

Vendor-Specific Orchestrators vs. Open-Source Tools 41

Multivendor Deployment 42

Hosted Architectures for Inference and Training 43

People and Processes 45

System Integrations 53

Summary and Future Considerations 58

Chapter 2 Regulations Impact on Autonomous Networks 59

Introduction 59

Short History of Data Privacy 59

Pre-Internet Era 60

1948—UN "Right to Privacy": A Human Right 60

1972—US Supreme Court: Legality of Domestic Intelligence Surveillance 60
1974—Family Educational Rights and Privacy Act 61
1974—Code of Fair Information Practice by US Federal Institutions 61
1980—OECD Recommendations for Personal Data Privacy and Transfers 61
1991—Privacy Against Automated Telephone Soliciting 62
Early Internet Days 62
1995—EU Data Protection Directive 63
1996—HIPAA Health and Medical Privacy 63
1998—COPPA: Protecting Children's Online Privacy 64
1999—GLBA: Financial Sector Data Privacy and Security 64
Privacy Versus Security Debate: USA Patriot Act of 2001 65
Digital Era 65
2002—E-Government Act: Government Entering the Internet Age 66
2003—California State Data Breach Notification Law 66
2010—Fighting Against Identity Theft with Red Flags Protection 66
2012—Right to Erase Personal Data: EU Right to Be Forgotten 67
2018—Comprehensive Privacy Protection with GDPR 68
2020—Ensuring Consumer Privacy with CCPA 69
Start of the AI Era 70
2024—Regulating AI Usage in EU with EU AI Act 70
2025—US State Laws Addressing AI Risks 71
Overview of EU Digital Regulations 72
EU Declaration of Digital Rights: Key Principles Supporting an Effective Digital Strategy Implementation 73
EU Data Strategy 74
The Wave of EU Digital Regulation 74
The State of AI Regulation in the Rest of the World 80
USA AI Regulation 81
Latin America: Brazil AI Regulation 83
China AI Regulation 84
Japan AI Regulation 85
South Korea AI Regulation 85
Australia AI Regulation 86
State of AI Regulation Conclusions 86

How Companies React to Existing AI and Data Regulation 88
The Cisco Approach to Data Privacy and Responsible AI 88
Safeguarding Data Privacy: Cisco's Global Privacy Policy 89
Deutsche Telekom Approach to Data Privacy and Responsible AI 91
Softbank Approach to Data Privacy and Responsible AI 94
Regulation Impact on Data and AI Projects and Mitigation Measures 96
Organizational Setup 97
Security Control Measures 99
Impact of Regulations on Autonomous Network Implementations 103
Examples of Regulations Impacting Autonomous Network Implementations 107
Challenges and Opportunities Created by Regulations 108
Summary and Future Considerations 109
References 110

Chapter 3 Observability and Data Sources 115
Introduction 115
Network Data Sources 117
Syslog 120
SNMP 122
Model-Driven Telemetry (MDT) 122
TACACS+ 124
RADIUS 125
NETFLOW/IPFIX 125
DHCP 126
DNS 126
Configuration Management Database 127
ITSM 127
Certificates (PKI/CA) 127
Collaboration Tools 127
Granular Data on Devices 128
Core Dumps 128
Enterprise Orchestrators (Collecting Telemetry) 128
Model-Driven Telemetry 131
End-User Quality of Experience 133
Observability and Tracing Principles 135
Observability Traces 137

Observability Metrics 138
Observability Logging 140
Instrumentation 141
Baggage 142
Profiling 143
Observability 2.0 145
Where Is It Used Today in Service Provider and Enterprise Domains? 146
Audit Trails in Autonomous Networks 147
Data Quality Principles 148
Data Mesh 149
Collecting and Correlating Disparate Traces 150
Baggage or No Baggage 151
Instrumentation 153
Network Observability Architecture Patterns 154
Observability Use Cases 156
Network Automation with Application Performance 156
AI Operations (GPU Activity) 157
AI Operations (Task Completion Rate/Quality) 159
AI Operations (Model Comparison) 159
AI Operations (Security Observability) 162
Summary and Future Considerations 165
References 165

Chapter 4 Conventional and Task-Specific AI/ML Models Usage for Autonomous Networks 167

Introduction 167
Short History of Artificial Intelligence 168
The Beginnings 168
First AI Winter (1970s) 169
Expert Systems (1980s) 169
Second AI Winter (Late 1980s–Early 1990s) 169
Machine Learning (After 1990s) 169
Deep Learning (After 2010s) 170
Generative AI (After 2017) 171
Current Status and Future Evolution of AI/ML for Autonomous Networks 171
Machine Learning Overview 172

Supervised Learning 173
Linear and Logistic Regression 174
Ensemble Models 176
Bagging: Random Forest 176
Boosting: AdaBoost, Gradient Boosting Machines, XGBoost 177
AdaBoost 178
Gradient Boosting Machines 179
Stacking 180
Support Vector Machine, Support Vector Classifier, and Support Vector Regression 182
Semi-Supervised Learning 184
Unsupervised Learning 186
Clustering 187
K-Means 188
DBScan 189
Dimensionality Reduction 189
Principal Component Analysis 189
Association Rule Learning 190
Reinforcement Learning 191
Value-Based Algorithms 192
Policy-Based Algorithms 193
Actor-Critic Algorithms 194
Model-Based Algorithms 195
Deep Learning and Neural Networks 196
Supervised Deep Learning: LSTM 198
Graph Neural Networks 199
Unsupervised Deep Learning: Autoencoders 200
Reinforcement Learning: Deep Q-Networks 201
Data Preprocessing 202
Data Cleaning 203
Data Transformation 203
Feature Engineering 204
Handling Imbalanced Data 204
Data Splitting 205
Model Evaluation, Validation, and Monitoring Metrics 206
Major Model Performance Metrics 207

Classification 207
Regression 209
Unsupervised Learning 210
Reinforcement Learning 212
Deep Learning 212
Ethical Considerations 214
MLOps 215
Challenges and Limitations of AI in Autonomous IP Network Operations 217
Lack of Data Quality and Label Scarcity 217
Lack of Interpretability 218
Robustness and Failure Amplification 218
Scalability and Computational Constraints 218
Limited Causal Reasoning and Long-Term Planning 219
Summary and Future Considerations 219
References 220

Chapter 5 Generative AI in the Context of Autonomous Networks 221
Introduction 221
Large, Small, and Vision-Driven Language Models 224
Multimodality 235
Prompting 236
Zero-Shot, Single-Shot, and Multi-Shot Prompting 237
Tokenization 238
Chain of Thought/Tree of Thought 240
Retrieval-Augmented Generation 241
Model Fine-Tuning 246
Beyond Language 249
Agentic Workflows 252
Goal-Based Operations 255
Practices for Building Reliable Agents 257
Agent Infrastructure 262
Federated Learning 264
Infrastructure for xLMs 268
Networking Use Cases 276
Summary and Future Considerations 279
References 280

Chapter 6 Automated Execution and Testing 281

Introduction 281

Industry Automation Frameworks and Standards 284

Automation Tools 287

Sources of Truth 289

The Data 289

Data Processing and Transformation Tools: 290

Processing Network Telemetry Data 291

Cluster Level A/B Testing 294

Automation Tools to Deal with the Data Itself 296

Stream Processing 297

Batch Processing 298

Dealing with Streaming Data 299

Using GitOps 301

Continuous Integration 302

Continuous Deployment 304

Container Creation and Best Practices 305

Multistage Building 308

Multi-Architecture Builds 310

Tags and Labels 310

WebAssembly 312

Cisco Virtual Kubelet 314

Network Test Automation 315

Test Definition 316

Testing AI Systems 316

Functional vs. Nonfunctional 317

Unit Test 319

Integration Test 319

System Test 319

Acceptance Test 320

Scale/Property Test 320

Service Test 320

Extended Testing Considerations 321

Test Tools 324

Network System Interaction 326

Using AI to Author and Conduct Network Testing 328

AI Agent and Model Performance 329
Chaos Engineering 331
Domain vs. Cross-Domain Automation Architecture 334
Network Automation Use Cases 337
Use Case 1: Predictive Maintenance 337
Use Case 2: Trend Analysis 338
Use Case 3: NOC/SOC Operations 339
Service Provider Focus Use Cases 340
Summary and Future Considerations 343
References 343

Chapter 7 Closed-Loop Command and Control 345

Introduction to Closed-Loop Systems 345
Origins of Control Theory 345
Closed Loop vs. Open Loop 347
Closed Loop: Negative and Positive Feedback 349
Why Closed Loop Matters in Networking Today 350
Evolution of Network Operations 351
Traditional Manual and Reactive NetOps 351
Shift to Automation and Programmability (SDN, NFV) 353
Emergence of Intent-Based Networking and AIOps 354
Core Concepts of Closed-Loop Network Operations 360
Observe, Orient, Decide, Act Cycle 360
Feedback Loops and Continuous Adaptation 362
Closed Loop at Different Network Operation Layers (Infrastructure, Services, Applications) 365
Role of AI/ML and Analytics in Automation 367
Hierarchies of Control Loops (Local, Domain, End-to-End) 369
Closed-Loop Operations Architectural Framework 371
Key Building Blocks 371
Data Collection and Telemetry 371
Data Processing and Analytics 373
Policy and Decision Engine 374
Orchestration and Integration with SDN 376
Standards and Reference Models 377
ETSI ZSM 377
ONF 379

TM Forum 379
ONAP 381
3GPP 383
High-Value Use Cases in Autonomous Network Operations 384
Use Case 1: Predictive Network Maintenance and Self-Healing 386
Use Case 2: Dynamic Service Assurance and Closed-Loop Optimization 389
Use Case 3: Change Management 391
Challenges and Limitations in Implementing Closed-Loop Operations 395
Data Quality and Real-Time Observability 395
Policy and Intent Translation 395
Closed-Loop Stability Through Reliable Decision Automation 396
Integration with Legacy and Multivendor Environments 396
Operational Trust, Governance, and Security 397
Closed-Loop Future Evolution 397
Summary and Future Considerations 398
References 399

Chapter 8 Communication Service Provider Autonomous Network Vision 401
Background and Industry Context 401
Evolution of Telecommunications Toward Digital Service Ecosystems 402
The Growing Complexity of Network and Service Operations 405
Definition of Network Autonomy Within Telecom Service Provider Context 407
End-to-End Autonomous Network Architecture 411
Core Architectural Principles 411
Conceptual Overview of an Autonomous Network Architecture 413
Key Elements 413
Observability 413
Common Messaging Data Bus 416
Automation Workflow and Service Catalog 418
Deploying CI/CD-Based Delivery Processes 418
Having a Scalable Agentic Platform 420
Comprehensive AI Governance 422
Reference AIOps Architecture 425
Typical Use-Case Deployment for Level 3+ 428
Service Impact Analysis 429

Predictive Maintenance 431
Root Cause Analysis 432
Capacity Planning and Forecasting 433
Typical Use Case Deployment for Level 4+ 435
Agentic Incident Resolution Framework for CSP Networks 437
Alert Triage Agent 438
Root Cause Analysis (RCA) Agent 439
Resolution Planning Agent 439
Resolution Execution (Orchestration) Agent 440
Validation and Learning Agent 440
Agentic Change Management Using Network Digital Twin 442
GameX Agentic Service Experience Assurance 445
Business Benefits 449
Levels 0–2: Moving from Manual to Rule-based Automation 449
Level 3: Human-Assisted Autonomy and Advanced Analytics 450
Level 4: Closed-Loop Autonomy 450
Dependencies and Obstacles 452
Summary and Future Considerations 455
References 455

Chapter 9 Enterprise Domain Vision 459
Introduction 459
The Function of Today's Networks 464
Usage Ratios for Artificial Intelligence Methods 466
Security Considerations 467
Web3 468
Decentralized Identity 468
Smart Contracts 472
Security of Model Weights 474
Shared Data Center Networks 475
Confidential Computing 478
Security of the Data Pipelines 480
Federation 482
SPIFFE Runtime Environment 482
End-to-End Autonomous Network Evolution Roadmap 484
Alignment with Core Business Applications 489
Planning and Deployment 494

Integration with Existing Systems 495
Technical Challenges 499
Summary and Future Considerations 506
References 507

Chapter 10 Autonomous Network Considerations—Organizational Change 509

Short History of Network Operations 509
Phase I: The Traditional NOC (1960s–Late 1990s) 510
Interval 1: The Mainframe/Analog Era (1960s–Mid-1970s) 511
Interval 2: The Digital Shift and Start of Decentralized Compute (Mid-1970s–Mid-1980s) 511
Interval 3: Birth of Standard Frameworks (Mid-1980s–Late 1990s) 512
Phase II: The Service Operation Center (SOC) and Service Assurance (Late 1990s–2010s) 512
Phase III: The Integrated Operations Center (2010s–Present) 513
Phase IV: The Autonomous and "Dark" NOC (Near Future) 515
Core Network Operation Concepts and Frameworks 516
COBIT 517
ISO 519
ITIL 522
eTOM 524
FCAPS 527
DevOps 529
Platform Engineering 532
Site Reliability Engineering 534
Conclusion 539
Typical Network Operations Center Setup for CSP 540
Typical NOC Setup for an IT Enterprise 548
Organizational Culture Shift 553
The Mindset Shift—From Heroes to Experience Assurers 553
Killing the "Bespoke" Culture 553
Democratizing Data (Breaking the "Not My Fault" Silo) 554
The Trust Gap—Eliminating the AI Black Box Syndrome 554
The Explainable AI Narrative 554
The Build vs. Buy Stress Test 555
The Identity Evolution—From CLI Wizard to Network Architect 555
Reskilling vs. Replacing (The DevNet Career Path) 555

Organizational Change: Managing the "Frozen Management Middle" 556

Ensuring Humans Can Still Gain Expertise in an Autonomous Network World: "The Glass Cockpit" Effect 556

New KPIs Being Defined to Better Reflect Autonomous Network Performance 557

Key Skills Required for AN-Based Operations 560

Programming Languages 560

Data Fluency 563

Understanding the Telemetry Transport and Ingestion Mechanics 563

Understanding Data Payloads and Data Navigation 563

Understanding Temporal Data Mechanics 564

Mastering Graph-Based Topological Representation 564

Understanding Metadata 565

Performing Data Validation 565

Data Science Fluency 565

Using Prompt and Context Engineering 567

Using Vibe Coding for Fast Software Prototyping 569

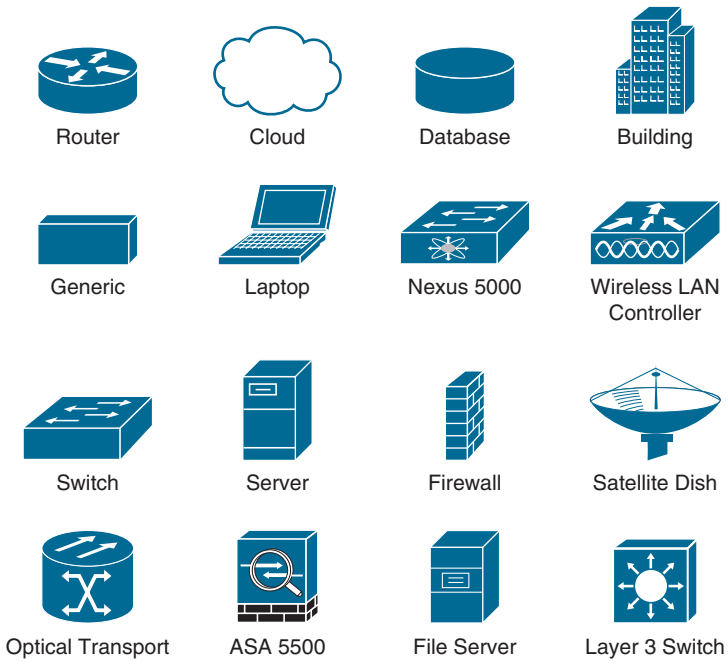
New Skills Adoption 570

Summary and Future Considerations 572

References 573

Index 575

Icons Used in This Book



Command Syntax Conventions

The conventions used to present command syntax in this book are the same conventions used in the IOS Command Reference. The Command Reference describes these conventions as follows:

- **Boldface** indicates commands and keywords that are entered literally as shown. In actual configuration examples and output (not general command syntax), boldface indicates commands that are manually input by the user (such as a **show** command).
- *Italic* indicates arguments for which you supply actual values.
- Vertical bars (|) separate alternative, mutually exclusive elements.
- Square brackets ([ ]) indicate an optional element.
- Braces ({ }) indicate a required choice.
- Braces within brackets ([{ }]) indicate a required choice within an optional element.

Introduction

This book aims to address an actual trend and answer the question: "How do you adopt and combine artificial intelligence, data management, and automation to ensure network, system, and application teams can reduce human hands-on involvement and enhance operational autonomy?" The idea of this book came while both authors were deeply engaged in transformational projects and initiatives supporting both enterprise and communication service provider (CSP) operational transformation and adoption of autonomous network frameworks.

While the practical implementation is still in its initial stages, the authors believe the industry and readers will benefit from understanding the key frameworks, technologies, regulatory frameworks, new processes, and skills that are making this transformation possible. This book makes each of the presented concepts real by presenting concrete examples where the concepts can be applied, relying on their personal experience in leading similar transformations with various customers.

After finishing this book, readers should have a practical understanding of all the elements that support the successful implementation of the operational autonomy vision, such as data management, AI/ML models, key network automation frameworks, relevant industry frameworks, and the regulatory environment as well as the key processes and new workforce skills.

Who Should Read This Book?

This book is aimed at all stakeholders involved in defining, implementing, and using the latest technology trends to enhance the level of network operational autonomy. This includes operations, implementation, and design engineers; data, automation, site reliability, and AI architects; data scientists; leaders; and executives who want to understand the key regulatory, technical landscape, key industry frameworks, and operational aspects they need to be aware of when planning their service provider or enterprise company evolution toward network autonomy. Given the authors' extensive industry experience, the theoretical concepts are paired with concrete real-world implementation and use-case examples to ensure practical adoption.

How This Book Is Organized

Chapter 1, "Network Autonomy": A history of network automation and key factors that need consideration for making the leap toward autonomy.

Chapter 2, "Regulations Impact on Autonomous Networks": An overview of the key regulatory trends across the world that govern AI and data domains.

Chapter 3, "Observability and Data Sources": An evaluation of network versus application observability, application performance monitoring, and the data needed to reach key outcomes.

Chapter 4, "Conventional and Task-Specific AI/ML Models Usage for Autonomous Networks": An overview of the key non-LLM AI/ML models and implementation methods that are relevant for supporting network deployments.

Chapter 5, "Generative AI in the Context of Autonomous Networks": An exploration of key concepts with generative and neuro-symbolic AI and related components such as AI agents and how these concepts are relevant for network autonomy.

Chapter 6, "Automated Execution and Testing": An overview of automation and test concepts, within both network and application domains, ranging from DevOps to GitOps, and cloud native orchestration.

Chapter 7, "Closed-Loop Command and Control": An overview of the closed-loop concepts and their usage in key high-value autonomous network use cases.

Chapter 8, "Communication Service Provider Autonomous Network Vision": An overview on how the core network autonomy concepts and key use cases are applied within the communication service provider environment.

Chapter 9, "Enterprise Domain Vision": An exploration of key concepts that are relevant in enterprise domains to support AI systems, including key security components and concepts.

Chapter 10, "Autonomous Network Considerations—Organizational Change": An overview of the key network operations center/service operations center (NOC/SOC) process changes and new skills that are required to support a successful adoption of autonomous network frameworks.

Credits

Figure 1-2	From TM Forum 2024
Figure 1-3	From Segment Routing for Service Provider and Enterprise Networks by Florian Deragisch, Leonir Hoxha, Rene Minder, Matthys Rabe, Kateel Vijayananda Published Dec 18, 2024 by Cisco Press.
Figure 1-5	John Boyd's OODA loop by Mark Lambertz. https://commons.wikimedia.org/wiki/File:OODA_Deutsch.png
Figure 1-6	Datadog Trace View Documentation
Figure 1-9	https://de.wikipedia.org/wiki/Datei:Internet_users_per_100_inhabitants_ITU.svg
Figure 1-12	https://commons.wikimedia.org/wiki/File:Diminishing_Returns_Graph.svg
Figure 1-20	From "Orange sets ambitious autonomous network operation goals by Joanne Taaffe, 23 Feb 2023. Used by permission of TM Forum.
Figure 1-22	Cloud Native computing foundation. https://landscape.cncf.io/
Figure 1-23	Adapted from Crossing the Chasm, 3rd Edition: Marketing and Selling Disruptive Products to Mainstream Customers by Geoffrey A. Moore, adapted by the Cloud Native Computing Foundation (CNCF).
Figure 1-27	From Networks with intelligence: Why and how the telecom sector should accelerate its autonomous networks journey, Capgemini Research Institute, 21 Feb 2024.
Figure 1-28	https://www.openconfig.net/
Figure 1-31	Exhibit from "How banks can supercharge technology speed and productivity", [Month Year], McKinsey & Company, www.mckinsey.com. Copyright (c) 2026 McKinsey & Company. All rights reserved. Reprinted by permission
Text 2.1	https://www.un.org/en/about-us/universal-declaration-of-human-rights
Text 2.2	https://gdpr-info.eu/art-4-gdpr/
Text 2.3	https://openfuture.eu/observatory/digital-decade/
Text 2.4	https://ksapa.org/eu-digital-rights-declaration-a-human-rights-milestone/
Text 2.5	https://www.eu-data-act.com/Data_Act_Article_3.html
Text 2.6	https://artificialintelligenceact.eu/article/3/
Text 2.7	https://www.meti.go.jp/shingikai/mono_info_service/ai_shakai_jisso/pdf/20240419_9.pdf
Text 2.8	https://api.telekom.me/api/fs/download/Tenderi/2020/9102020-Prilog %205.pdf?root=/Tenderi/2020/9102020-Prilog%205.pdf#:~:text=5.6%20 Data%20Security%2C%20Data%20Protection%20and%20AI,shall%20 be%20non%2Ddiscriminatory%20and%20transparent%20towards%20users

Text 2.9	https://www.high-tech-tele.com/code-of-conduct
Text 2.10	http://static.hrvatskitelekom.hr/webresources/tht/pdf/investitori/politike/en/Supplier_Code_of_Conduct_ENG_version_for_publication.pdf? utm_source=chatgpt.com
Figure 2-2	From Organizational Digital Governance Report 2025. copyright©The International Association of Privacy Professionals, Inc. Used by permission.
Figure 2-3	Adapted from CMS
Figure 2-4	Adapted from Diconium
Figure 2-6	National Institute of Standards and Technology. U.S dept of commerce. https://airc.nist.gov/airmf-resources/playbook/
Figure 2-8	https://www.telekom.com/en/company/digital-responsibility/details/ house-of-digital-responsibility-1008292
Figure 2-12	Adapted from Iapp
Figure 2-13	Adapted from Iapp
Figure 2-15	https://www.robustintelligence.com/ai-security-taxonomy
Figure 3-1a	https://www.etsi.org/newsroom/blogs/technologies/entry/the-recently-released-etsi-zsm-specification-on-closed-loop-enablers-provides-a-significant-boost-for-enabling-full-end-to-end-network-and-service- automation
Figure 3-1b	https://www.etsi.org/newsroom/blogs/technologies/entry/the-zsm-story-the-power-to-transform#:~:text=The%20ZSM%20framework%20supports%20the%20separation%20of,assurance%20(data%20collection%2C%20analytics%20and%20intelligent%20automation)
Figure 3-2	From Whitepaper of Catalyst Project Autonomous Networks Hyperloops – Enabling Self-X Digital Services for Smart-X Verticals. Used by permission of Tm Forum
Figure 3-3	From Whitepaper of Catalyst Project Autonomous Networks Hyperloops – Enabling Self-X Digital Services for Smart-X Verticals. Used by permission of Tm Fourm
Figure 3-14a	Marisol Palmero https://www.linkedin.com/in/mpalmero/
Figure 3-17	https://github.com/open-telemetry/opentelemetry-ebpf-profiler apache 2 licensing
Figure 3-18	steve flanders
Text 4.1	https://doi.org/10.1007/s12525-021-00475-2
Text 4.2	https://www.sciencedirect.com/science/article/pii/S0377221719307581
Text 4.3	https://christophm.github.io/interpretable-ml-book/

Figure 4-2	From “Long-Term Data Traffic Forecasting for Network Dimensioning in LTE with Short Time Series” by Gijón et all. Electronics 10, no. 10: 1151. https://doi.org/10.3390/electronics10101151
Figure 4-5	From Realizing Low-Energy Classification Systems by Implementing Matrix Multiplication Directly Within an ADC. IEEE Transactions on Biomedical Circuits and Systems, (), 1–1 by Wang et all. doi:10.1109/TBCAS.2015.2500101
Figure 4-6	Adapted from Implementation of XGBoost (eXtreme Gradient Boosting) by geeksforgeeks
Figure 4-7	From “Solving the multicollinearity problem to improve the stability of machine learning algorithms applied to a fully annotated breast cancer dataset” by Katrina I. Sundus et all. https://doi.org/10.1016/j.imu.2022.101088
Figure 4-12	Adapted from Data teach note
Figure 4-19	From Mobile Traffic Prediction from Raw Data Using LSTM Networks by Hoang Duy Trinh et all. 10.1109/PIMRC.2018.8581000
Figure 4-20	From “Anomaly Detection of Water Level Using Deep Autoencoder” by Isack Thomas Nicholaus et all. https://doi.org/10.3390/s21196679
Figure 4-21	From Deep Reinforcement Learning Based Resource Allocation For Narrowband Cognitive Radio-IoT Systems by K.F. Muteba et all. https://doi.org/10.1016/j.procs.2020.07.046
Figure 4-23	From Non-linear Dimensionality Reduction-based Intrusion Detection using Deep Autoencoder by Sreenivasa Chakravarthi & Jagadeesh Kannan. 10.14569/IJACSA.2019.0100822
Figure 4-24	From Short-Term Traffic Volume Prediction for Sustainable Transportation in an Urban Area by Habio Mei et all. https://doi.org/10.1061/(ASCE)CP.1943-5487.0000316
Figure 4-25	From Selecting the number of clusters with silhouette analysis on KMeans clustering. © Copyright scikit-learn developers (BSD License).
Figure 4-26	From Deep Learning Methods for Drum Transcription and Drum Pattern Generation, Nov 2018 by Richard Vogl. Used by permission
Figure 4-27	From Operationalizing AI in Future Networks: A Bird’s Eye View from the System Perspective by Qioong Liu et al
Text 5.3	https://www.anthropic.com/engineering/building-effective-agents
Text 5.4	https://breakingdefense.com/2024/12/how-the-pentagon-domesticated- the-chatbot-2024-in-review/
Figure 5-5	Redrawn by Josh Halley - adapted from https://arxiv.org/pdf/1706.03762
Figure 5-10	Taken from web using openai tokenizer

Figure 5-11	From Chain-of-Thought Prompting Elicits Reasoning in Large Language Models by Jason Wei et al. https://arxiv.org/pdf/2201.11903
Figure 5-12	From Tree of Thoughts: Deliberate Problem Solving with Large Language Models by Shunyu Yao et al. https://arxiv.org/pdf/2305.10601
Figure 5-29	Adapted from Reco vs. LLM from The Next Platform.
Figure 5-39	From AN JOURNEY GUIDE Level 4 industry blueprint high-value scenarios Networks, JUNE 2024. Used by permission of TM Forum.
Text 6.1	https://x.com/solomonstre/status/1111004913222324225
Figure 6-1	From Autonomous Networks Framework v2.0.0 (IG1218F). Used by permission of TM Forum
Figure 6-6	TM Forum's Digital Maturity Model. Used by permission of TM Forum
Figure 6-26	Re-created from https://medium.com/@shivraj.jadhav82/webassembly-wasm-docker-vs-wasm-275e317324a1
Figure 6-28	L isa Crispin and Janet Gregory, Agile Testing: A Practical Guide for Testers and Agile Teams.
Figure 6-31	https://twitter.com/teenybiscuit// status/707670947830968320 (Accessed March 2021).
Figure 6-37	https://www.devopsschool.com/blog/what-is-chaos-toolkit-and-use-cases- of-chaos-toolkit/
Figure 6-39	From Autonomous Networks – Reference Architecture v1.0.1 (IG1251). Used by permission of TM Forum
Text 7.1	https://www.gartner.com/en/information-technology/glossary/aiops-artificial-intelligence-operations
Figure 7-15	Adapted from Sketchplanations
Figure 7-16	From Choosing a closed loop pattern and logical architecture. Used by permission of TM Forum
Figure 7-20	From Back to the future: NGOSS, autonomous networks and 'ODA Contracts'. Used by permission of TM Forum
Figure 7-25	From What is a telemetry pipeline? Understanding and building effective telemetry data pipelines By Tamás Burtics. Used by permission of NXLog. Ltd
Figure 7-26	From Agentic AI Architecture: A Practical, Production-Ready Guide by Monoj Kanti Saha, Aug 30, 2025. Used by permission.https://medium.com/agenticai-the-autonomous-intelligence/agentic-ai-architecture-a-practical-production-readyguide-2b2aa6d16118
Figure 7-27	From AI Agents Under Threat: A Survey of Key Security Challenges and Future Pathways by Zehang Deng et all. https://arxiv.org/abs/2406.02630
Figure 7-29	© ETSI 2026. All rights reserved.

Figure 7-30	From ODA blog series: Taking autonomous networks from theory to reality. used by permission of TM Forum
Figure 7-31	https://docs.onap.org/en/jakarta/guides/onap-developer/architecture/onap-architecture.html
Figure 7-32	https://www.onap.org/architecture
Figure 8-4	From An Analysis of Telco Opex: Minimizing Costs and Maximizing Efficiency by Adam Mackenzie. Used by permission of Omdia
Figure 8-5	From The Nature of Telecom Capex – a 2023 Update. Used by permission of Kim Larsen
Figure 8-6	From Autonomous Networks Technical Architecture v1.1.1 (IG1230). Used by permission of TM Forum
Figure 8-7	From Autonomous Networks Technical Architecture v1.1.1 (IG1230). Used by permission of TM Forum
Figure 8-15	https://docs.agntcy.org/#technical-objectives
Figure 8-21	From IG1401 TM Forum AN Journey Guide: Level 4 industry blueprint high-value scenarios v6.0.0. Used by permission of TM Forum
Figure 8-25	https://www.tmforum.org/catalysts/projects/M25.0.824/Game-X-?_gl=1*7lwo7j*_gcl_au*OTU5OTkwODQxLjE3NjI3MDE5OTk.
Figure 8-26	https://www.tmforum.org/catalysts/projects/M25.0.824/Game-X-?_gl=1*7lwo7j*_gcl_au*OTU5OTkwODQxLjE3NjI3MDE5OTk.
Figure 8-28	From Networks with intelligence Why and how the telecom sector should accelerate its autonomous networks journey. Used by permission of Capgemini Research Institute
Figure 8-29	From Networks with intelligence Why and how the telecom sector should accelerate its autonomous networks journey. Used by permission of Capgemini Research Institute
Figure 8-30	From ACHIEVING NEXT-LEVEL NETWORK AUTONOMY. Used by permission of STL partners
Figure 8-31	From Networks with intelligence Why and how the telecom sector should accelerate its autonomous networks journey. Used by permission of Capgemini Research Institute
Text 9.1	https://www.melconway.com/Home/pdf/committees.pdf
Table 9-2	https://ieeexplore.ieee.org/document/9150191
Figure 10-1	Everett Collection/Shutterstock
Figure 10-3	From Open Digital Architecture (ODA). Used by permission of TM Forum
Figure 10-4	Adapted from ISACA, COBIT 2019 Framework: Introduction and Methodology. Generated via ChatGPT by the author

Figure 10-6	https://yasm.com/wiki/en/img/yasm-frameworks/itil/itil-4-practices-itil-v3-processes-comparison.jpg
Figure 10-7	https://en.wikipedia.org/wiki/Business_Process_Framework_(eTOM)
Figure 10-10	From CNCF Platforms White Paper. https://tag-app-delivery.cncf.io/whitepapers/platforms/
Figure 10-11	From Site Reliability Engineering (SRE): A Comprehensive Guide. Used by permission of Spoclearn. https://www.spoclearn.com/blog/what-is-site-reliability-engineering-sre/
Figure 10-12	https://sre.google/resources/practices-and-processes/ incident-management-guide/
Figure 10-22	From "The Top Programming Languages 2025 Does AI mean the end for the Top Programming Languages? by Stephen Cass 23 Sep 2025. Used by permission.

Chapter 2

Regulations Impact on Autonomous Networks

In this chapter, you will learn the following:

- The history of data privacy legislation
- European Union digital regulations in the past 10 years
- The approaches the rest of the world takes to AI regulations
- How companies react to existing data and AI regulations
- The existing regulation impact on data and AI projects

Introduction

Data is key to the system observability and AI-driven decision-making process that enables network autonomy. In addition, the task of collecting, storing, and processing data has been increasingly regulated over time, highlighting its growing importance as a valuable insight-providing asset and its potential to harm if used improperly.

In this chapter, we will present an overview of the major data privacy and AI regulatory frameworks across the world, with a special focus on EU data and AI governance. We will examine how companies are addressing these regulations by implementing sometimes even stricter internal standards and share several best practices on how projects implementing autonomous networks elements could ensure compliance with existing and future regulations.

Short History of Data Privacy

Privacy is the ability of an individual or group to seclude themselves or information about themselves and thereby express themselves selectively.

The concept of privacy has existed since ancient Greece, where a key distinction was made between public places (or *polis*), where issues could be debated freely and openly, and the private area of the home.

Over the years, the issue of ensuring privacy was closely linked with technological advancements. In the 14th century, opening somebody else's letters was the main privacy concern, but during the 19th century, collecting and publishing somebody's private data in a journal article became a subject of focus. The addition of photography, which could visually capture situations deemed private, added a new dimension to the list of technology-enabled privacy concerns. As technology evolved, video, audio, and even more complex topics such as user health, education, and behavioral data gradually became part of the general topic of privacy. With the introduction of big data technologies and artificial intelligence algorithms that can consume large amounts of personal data and act to influence a person's behavior, sharing personal data willingly or unwillingly and the general topic of privacy sit naturally at the forefront of regulatory concerns.

Pre-Internet Era

In the era before the Internet, laws concerned with privacy in general focused on the government's ability to collect, store, and process an ever-increasing amount of data without legal oversight. Guiding principles for data privacy were being set.

1948—UN "Right to Privacy": A Human Right

After the second world war, in the light of the horrors the population had been exposed to, with governments collecting and using citizens' private data for their own military or political purposes, the newly created United Nations addressed privacy as a fundamental human right as part of the Universal Declaration of Human Rights (UDHR).

In Article 12, the UDHR states that privacy is an inalienable human right: "No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honour and reputation. Everyone has the right to the protection of the law against such interference or attacks." [UDHD article 12]

This was the first time privacy was addressed at the international level, and this declaration went above and beyond what nation states' constitutions guaranteed at the time.

While the Universal Declaration of Human Rights was not legally binding, it set the stage for the formal and legal recognition of 'the personal right to privacy' as a human right across the world, such as article 8 of the European Convention of Human Rights adopted in 1953.

1972—US Supreme Court: Legality of Domestic Intelligence Surveillance

As electronic surveillance and wiretapping techniques evolved, there was an increased need to ensure they follow a legal due process. One landmark case (*US vs US District Court 407 US 297,* also known as the Keith case) enabled the US Supreme Court to reconfirm that any domestic intelligence surveillance is governed by a pre-existing legal warrant.

1974—Family Educational Rights and Privacy Act

The aftermath of the Watergate scandal highlighted the need to protect citizens against the dangers of government data gathering and the abuse of personal files.

At the same time, the need for a student-specific privacy law was recognized because

- Attending school is mandatory; hence, students must have their data processed.
- Schools process data regarding children, who can be considered a vulnerable party.
- Data processing is a mandatory part of school processes; hence, opt-outs are generally not feasible.

The Family Educational Rights and Privacy Act (FERPA) establishes the importance of personally identifiable information (PII) within student records, grants parents access to their children's records, and allows information amendments by record owners. It also introduces the concept of consent, which is required to be provided by a student (after turning age 18) for any private information disclosure.

The law applies to all schools, colleges, universities, and other institutions that receive funds from the US Department of Education.

FERPA was one of the earliest privacy laws passed in the United States, and it acted as a blueprint for other privacy laws introduced later across the world.

1974—Code of Fair Information Practice by US Federal Institutions

Like FERPA, the Privacy Act of 1974 was created in the aftermath of the Watergate scandal, which highlighted the need to establish a set of best practices governing the collection, maintenance, use, and dissemination of personally identifiable information of individuals that is maintained by US federal agencies. The act prohibits disclosure of information from a system of records without the written consent of the subject individual (barring certain exceptions). It also grants individuals the power to access, amend, and find out if their records have been disclosed.

While the Privacy Act provisions were an early admission of the risks posed by the increasing power of technology, allowing an ever-growing amount of data to be collected, stored, and shared at scale, they also highlight the sector-specific focus of the US privacy legislation laws (they cover only data stored by federal agencies, leaving out data held by courts, executive components, or non-agency government entities), which has shaped the overall philosophy of the US privacy legislation up to the current times.

1980—OECD Recommendations for Personal Data Privacy and Transfers

To help create a comprehensive data protection system throughout Europe, the Organization for Economic Cooperation and Development (OECD) has issued a set of seven guiding principles governing the protection of privacy and trans-border flows of

personal data that set the basis of the EU's data protection legislation, building on existing worldwide practices:

- **Notice:** Data subjects should be given notice when their data is being collected.
- **Purpose:** Data should be used only for the stated purpose.
- **Consent:** Data should not be disclosed without the data subject's consent.
- **Security:** Collected data should be secured.
- **Disclosure:** Data subjects should be informed who is collecting their data.
- **Access:** Data subjects should be allowed to access their data and make amendments.
- **Accountability:** Data subjects should have a method available to hold data collectors accountable for not following the aforementioned principles.

While nonbinding in nature, the OECD recommendations set the overall strategy for later EU data privacy implementations. While the US has endorsed OECD recommendations, it did not prioritize their adoption within a unified federal approach, preferring to continue with a sector-specific data privacy legislation strategy.

1991—Privacy Against Automated Telephone Soliciting

The Telephone Consumer Protection Act (TCPA) was designed to address privacy protection against the evolving technology that allowed mass-scale automated calling for telemarketing or general purposes.

This limits companies from calling potential customers using automated dialing systems, artificial or prerecorded voice messages, Short Message Service (SMS), and fax machines. It also enforced the soliciting entity to provide the necessary information allowing its identification by the called customer.

The TCPA also allows customers to request not to be called for soliciting purposes, their choice being recorded either in company specific do-not-call lists (valid for five years) or in a national Do Not Call Registry, which has broader applicability.

This is a good example of a law regulating the right of consumer privacy against technology advances (mass automated calling for soliciting purposes) and adding one of the earlier forms of "opt-out" rights through the Do Not Call Registry implementation.

Early Internet Days

The early days of the Internet saw a hybrid form of the laws, covering various forms of data transfer and protection, which include but are not generally focused on the Internet, following the OECD and UN Right to Privacy principles previously set.

1995—EU Data Protection Directive

With the growth of the Internet, data exchange began to gain significance for individuals and businesses; however, the lack of a unified EU data privacy law prevented a truly free flow due to differences of data privacy interpretation across various EU member states.

Building on top of the 1980 OECD data privacy principles and article 8 of the EU Convention of Human Rights, which granted a "right to respect for one's private and family life and his correspondence," the EU proposed the Data Protection Directive, regulating the processing of personal data, regardless of whether this was done automatically or not.

The directive centers around the definition of *personal data*, which is any information relating to an identifiable natural person, with any means necessary (direct, indirect, or by referring to one or more factors that are unique to this person's identity).

Data protection rules apply not only when the controller is established within the EU but also for any entity that uses equipment located within the EU.

The directive also sets three principles under which data should be processed:

- **Transparency:** This directive summarizes the OECD's Notice and Disclosure principles: data subjects have the right to be informed when and who is processing their data.
- **Legitimate Purpose:** This directive adds legal legitimacy to the OECD's Purpose principle.
- **Proportionality:** This directive aims to limit the amount, details, and time storage of data beyond a reasonable legitimate purpose.

Another important directive focus area is the transfer of personal data to third countries; this transfer is allowed only if the target country provides an "adequate" level of protection of the data. This directive has caused multiple negotiations related to the transfer of EU PII data, especially within the US.

The EU Data Protection Directive has shown that the EU approach to data privacy is comprehensive and exhaustive, contrary to the targeted sectorial US legislation in this matter.

1996—HIPAA Health and Medical Privacy

The Health Insurance Portability and Accountability Act (HIPAA) was created to streamline the flow of healthcare information and publicize standards for the electronic exchange, privacy, and security of health information.

As part of the HIPAA provisions, a special section describes the "Standards for Privacy of Individually Identifiable Health Information," also known as the "Privacy Rule," which for the first time established a set of national standards for the protection of key health information. The regulation strikes a balance that permits important use of health information while protecting the privacy of people who seek healthcare.

The protected health information (PHI) protects all individually identifiable health information in any form or media (electronic, paper, or verbal).

Just like in the EU Data Protection Directive, identifiable health information is data that can be used to identify the individual, showing the importance of the concept in both the EU and US.

1998—COPPA: Protecting Children's Online Privacy

Due to the significant increase in Internet-based electronic commerce, the Children's Online Privacy Protection Act (COPPA) was created to address various concerns regarding data collection practices and the general impact of Internet commerce on children (any individual under the age of 13).

COPPA was considered the first US privacy law written especially for the Internet age. As part of its specification, COPPA mandates

- Websites must require verifiable parental consent using an approved method before the collection or use of any under-13 information (the Consent principle).
- There must be a clear guideline of the data collection policy, including the mandatory display of the policy when data is being collected (the Notice principle).
- Websites must include what responsibilities the website operator legally holds regarding under-13 privacy and safety (the Accountability principle).
- Website operators must allow parents to review their children's personal information (the Access principle).

1999—GLBA: Financial Sector Data Privacy and Security

The Gramm-Leach-Bliley Act (GLBA), or the Financial Modernization Act, was created to address concerns about data security and privacy within the financial sector. GLBA was one of the first US data privacy laws to impose specific data privacy constraints on businesses.

The regulation aims to protect consumers' financial information and prevent sensitive data exposure by imposing several best practices to the financial organizations handling user data:

- **Financial Privacy Rule:** Organizations must provide a clear description on how personal data is collected, used, and shared (the Notice and Disclosure principles).
- **Safety Rule:** Organizations must maintain robust data security practices to safeguard user data against unauthorized access (the Security principle).
- **Pretexting Rule:** It is illegal for anyone to obtain or disclose financial customer information under false pretenses (which is an early form of protection against identity theft).

Privacy Versus Security Debate: USA Patriot Act of 2001

Following the terrorist attacks of 2001, the US Congress enacted the USA PATRIOT Act (commonly known as the Patriot Act), which rolled back a significant portion of the data privacy laws provisions that existed at the time.

Citing national security as a primary concern, the act allows government agencies to gather "foreign intelligence information" from both US and non-US citizens. This included

- Expanding domestic and international wiretapping
- Expanding authority to access electronic communications
- Allowing secret "sneak-and-peak" searches (without user consent or notification)
- Removing privacy protections to allow federal agencies to share more information

Although the USA Patriot Act provisions were temporary in nature and eventually eliminated, this law highlights the delicate balance between privacy and security, which increases in relevance in the AI/big data age.

This is the first known situation in recent US history when privacy regulations were rolled back due to a single but significant security incident, and it raises significant questions on how this precedent will be used (or not) for future similar situations, especially given that technology currently has the means to collect, process, and interpret population-scale data.

Digital Era

In the digital era, regulations are directly concerned with digital data that is being generated, stored, or transmitted over the Internet, acknowledging digital communication and data as being the biggest privacy concern of the 21st century.

The exponential growth of data being generated, which is at the center of defining the need for additional regulation, is illustrated in Figure 2-1.

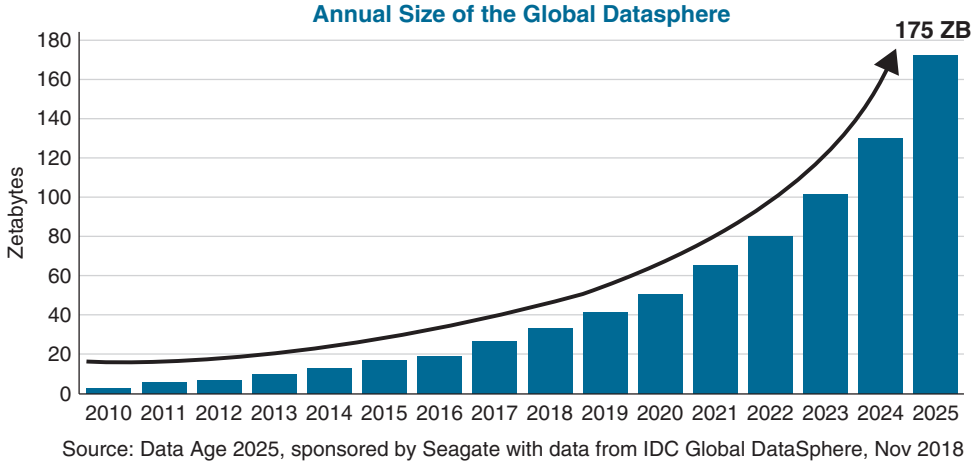


Figure 2-1 *Annual Size of the Global Datasphere Between 2010 and 2025*

2002—E-Government Act: Government Entering the Internet Age

The US E-Government Act of 2002 intended to modernize the federal government tools, communication, practices, and services to take advantage of the innovation brought by the Internet Age.

Its main provisions relate to

- **Public Information:** Agencies must provide a website that includes all "information about the agency," following the Disclosure principle.
- **Electronic Submission:** Agencies must accept a wide area of document submissions electronically (enabling the start of digital government services).
- **Electronic Dockets:** Agencies must have an Internet-accessible rule-making docket that includes all public comments and other materials created by agency rule or practice, providing transparency on each agency policy and practice.
- **Privacy Impact Assessment:** Agencies must conduct a "privacy impact assessment before collecting information" using any form of information technology, with a special focus on collecting "privacy information in an identifiable form."

The US E-Government Act of 2002 was important because it created a structure for the federal government to deliver better services to citizens using information technology and the Internet, a practice that opened the door for state-delivered safe electronic services.

2003—California State Data Breach Notification Law

With the advances in information technology, an increased amount of personal user data was being collected and stored by both businesses and government entities. The general lack of data encryption and large amount of stored data provided a significant attack surface for potential malicious actors who were interested in extracting and using this data in an unauthorized manner.

In 2003, California was the first US state to implement a law addressing the need to inform users in case state residents' unencrypted personal information was exposed in a security breach.

The California example was quickly followed by other US states, and at the time of writing this book, each US state had a similar legal requirement enacted in law.

2010—Fighting Against Identity Theft with Red Flags Protection

With the rise of electronic communication, the issue of electronic identity theft has become a major challenge, with an estimated 9 million plus Americans having a form of electronic identity stolen in 2013 and the numbers continuing to grow.

The US Federal Trade Commission (FTC) and National Credit Union Administration (NCUA) created the Red Flags Rule to force financial and credit institutions to ensure protection and early warning against identity theft.

The Red Flags Rule describes how to develop, implement, and administer an identity theft prevention program, following four major guidelines:

- A program must include reasonable policies and procedures to identify the red flags of identity theft that may occur in day-to-day operations (e.g., identifying an ID that does not look genuine is a potential red flag).
- A program must be designed to detect the red flags that have been identified (e.g., have clear procedures to determine how to identify an ID to be false).
- A program must spell out appropriate actions to be taken after identifying a red flag (e.g., alert authorities, block the user presenting the false ID).
- A program must detail how to keep it current to reflect new threats (how to make sure to continue to add mechanisms to detect false IDs created using new and more sophisticated methodologies).

The Red Flags Rule was one of the earliest attempts at mandating businesses to implement protection against identity theft attempts, recognizing this emerging threat.

2012—Right to Erase Personal Data: EU Right to Be Forgotten

With social network utilization exploding, the amount of data collected and stored by each platform increased exponentially and generated the following issue: What happens with the data when a user decides to leave the platform?

Under the previous EU Directive (the EU Data Protection Directive), the interpretation left room for platform owners to either "freeze" or "retire" data, meaning the data continued to exist even if it was not exposed anymore to the public.

Under the new EU "Right to be forgotten" regulation, the data has to be explicitly "erased without delay" at user request under certain conditions:

- The data is no longer necessary related to the original collection purpose.
- The data subject withdraws consent.
- Processing of data was unlawful.
- The data controller has a legal obligation to erase the data (e.g., following a court ruling).
- The data was collected in relation to the offering of online services.

The implementation of this regulation caused a significant shift in user perception of their rights to control how personal data is managed, especially by social platforms. On only its first day of compliance (May 30, 2014), Google received 12,000 requests to have personal details removed from its search engine, with Facebook, YouTube, and Google constantly being in the top 10 sites where data erasure was requested.

2018—Comprehensive Privacy Protection with GDPR

The EU's General Data Protection Regulation (GDPR) is a law designed to comprehensively cover data protection, privacy, and management for any systems (within or outside the EU) that handle EU citizens' data.

The GDPR supersedes the previous EU legislation covering data management and privacy (EU Data Protection, EU Right to Be Forgotten, etc.) and was created to bring the EU legislation in line with the challenges and opportunities provided by the Internet and big data technology revolution.

Under the GDPR, personal data is defined very broadly as "any information relating to an identified or identifiable natural person."

The GDPR (re)defines the general data protection principles, evolving the original OECD concepts:

- **Data Processing:** This process must be lawful, transparent, and fair (an extension of the original Disclosure principle).
- **Purpose Limitation:** This principle is in line with the original OECD Purpose principle.
- **Data Minimization:** This principle is an addition to purpose limitation, where only necessary data collection should be performed.
- **Accuracy:** This principle is in line with the original OECD Accuracy principle.
- **Storage Limitation:** This principle imposes an additional time constraint on the original Purpose limitation principle (e.g., by default, data should have an expiration date, after which it must be erased).
- **Integrity and Confidentiality:** This principle expands the original Security principle by adding an explicit requirement to the data collector to ensure appropriate security, integrity, and confidentiality by using the technology available (e.g., encryption).
- **Accountability:** This principle is in line with the original OECD Accountability principle.

In addition, the GDPR includes additional policies aimed at strengthening data security and end-user data ownership rights:

- **Data security by design:** Organizations must, "by design and by default," consider data protection.
- **User consent:**
  - Consent must be "freely given, specific, informed and unambiguous," can be withdrawn, and must be kept as evidence.
  - Children under age 13 need parental permission to provide consent (like the COPPA requirements).

- **People privacy rights (covering and expanding previous directives such as "right to be forgotten"):**
  - The right to be informed
  - The right of access
  - The right to rectification
  - The right to erasure
  - The right to restrict processing
  - The right to portability of data
  - The right to object
  - Rights in relation to automated decision-making and profiling

To ensure GDPR compliance, the fines for violating the GDPR are very high. There are two levels of fines, which can be as high €20 million or 4 percent of global revenue (whichever is higher). The GDPR also allows data subjects to ask compensation for damages.

The GDPR is considered to this date the most strict and comprehensive user data privacy and data confidentiality regulation in the world; it acts as a golden standard for similar legislation in non-EU countries.

As of 2026, the EU Commission is proposing targeted simplifications to GDPR, particularly for SMEs and small mid-cap companies (fewer than 750 people) that do not have the staff to implement the entire set of GDPR requirements, focusing on reducing administrative burdens such as record-keeping requirements while maintaining the core GDPR principles. These reforms are still under negotiation and not yet fully in force.

2020—Ensuring Consumer Privacy with CCPA

Following the GDPR enactment into law, California was the first US state to enact a similar broad privacy law. The California Consumer Privacy Act (CCPA) approved in 2020 aims to provide California state consumers more control over the personal information businesses collect about them. Its provisions are built following similar principles to GDPR, including

- The right to know what information is collected and how it is used (the Disclosure principle).
- The right to delete personal information collected (similar to the Right to erasure/Right to be forgotten). This right was further expanded, allowing consumers to correct inaccurate personal information.
- The right to opt out of sharing information (similar to the Right to object/Restrict processing). The CCPA focuses on allowing an opt-out after information is collected, whereas the GDPR mandates obtaining consent up front.
- The right to limit the use of sensitive personal information (similar to the Right to restrict processing offered by the GDPR).

Whereas the GDPR has a broad applicability, covering any potential personal information being collected from an EU citizen and strictly requires up-front consent, in most cases the CCPA does not require up-front consent, enabling consumers to opt out later. Also, it has a narrower interpretation of what constitutes personal information (mainly data identifying a consumer or household) and applies only to legal residents of the state of California.

The CCPA was the first comprehensive state privacy law introduced in the US, opening the door for other states to follow. At the time of writing this book, more than 20 states have enacted similar laws, with more states to follow, highlighting the importance of data privacy as a foundational element of trust in doing online business.

Start of the AI Era

At the start of the AI era, regulations that build on top of data privacy (what data you are allowed to obtain and how you are supposed to manage it) toward "what insights and actions you derive from the data" reflect the emergence of the enhanced capabilities of AI algorithms. While the legislation acting as baseline of AI regulation mainly addresses the use of personal data, there is also regulation of nonpersonal data (data that cannot identify a person) through laws such as the EU Data Act, which provides guidelines on handling of Internet of Things (IoT) data.

Apart from privacy concerns, the AI legislation also addresses the public need for the AI to be ethical. Ethical means that AI makes decisions and recommends actions that are transparent, trustworthy, and beneficial to society. This legislation addresses the need to guardrail and control AI from growing to become a threat that is outside human control, which is exemplified in articles such as "AI and the Paperclip Problem," written by Joshua Gans.

This issue will be described in detail in the rest of the chapter.

2024—Regulating AI Usage in EU with EU AI Act

Following the growth of global concerns on the increasing power and perceived lack of transparency related to the decisions and actions that are taken by leveraging AI algorithms, the EU has created the world's first comprehensive regulation on AI usage through the EU AI Act.

The AI Act classifies AI according to its risk and imposes regulatory constraints accordingly:

- **Prohibited AI Systems:** Any system that deploys subliminal, manipulative, or deceptive techniques to distort behavior or decision-making. In this category also fall systems that use biometric categorization (based on sensitive individual attributes such as sex, religion, political opinion) or perform real-time biometric identification (outside specific law-mandated situations) and therefore affecting personal rights and freedoms.

- **High-Risk AI Systems:** Systems used in areas such as critical infrastructure, education, employment, law enforcement, and access to essential services, where they significantly affect individuals, rights or safety. These systems must undergo rigorous risk management, data governance, and compliance assessments before deployment. Also, they must be documented in detail (technical documentation and record keeping) and allow for human oversight.

Note AI systems intended to be used as safety components in the management and operation of critical digital infrastructure communication systems, including networking, are considered critical. This means communication service providers and IT networks and surrounding management systems fall in this category if they are providing a critical connectivity service either at the national level or supporting other critical enterprises (e.g., utilities).

- **General-Purpose AI:** Any AI model that displays significant generality and is capable of competently performing a wide range of distinctive tasks as well as integrating with other AI models. These systems must provide comprehensive technical documentation and provide a detailed summary about the content they used for training before being allowed for use.
- **Limited Risk:** AI systems manipulating images, sound, or videos where there are transparency obligations of informing the users they are interacting with a product of an AI system.
- **Minimal Risk:** (e.g., video games or AI spam filters) Where most AI applications are expected to exist. These systems are unregulated; however, a voluntary code of conduct is suggested for the application owners.

As a notable exception, AI systems used for military, national security, or pure scientific research are exempted from the AI Act.

As in the GDPR case, penalties for noncompliance are high and can result in fines up to €35 million or 7 percent of annual company turnover.

The applicability of the EU AI Act (just like in GDPR cases) goes beyond the EU physical border and applies to any AI system or product that processes data or provides services to EU citizens.

2025—US State Laws Addressing AI Risks

Following the EU AI Act example, in the absence of a federal comprehensive approach, various US states have started to draft similar AI legislation. The first notable examples are as follows:

- **Colorado AI Act (CAIA):** This act, in a similar manner to the EU AI Act, focuses on high-risk AI systems that are a substantial factor in influencing user "consequential decisions," which in this context are defined as education, employment, financial

services, housing, health, or legal services. The act also protects against "algorithmic discrimination," which like the EU AI Act relies on "protected user characteristics" such as "religion, sex, or political views." High-risk system documentation and consumer awareness (disclosure) are mandated as part of the law. The CAIA is the first US state law that aims to comprehensively protect against the risk of deploying and using AI systems.

- **California AB2013: Generative AI Law:** This law regulates "generative AI," defined as AI that can generate derived synthetic content such as text, images, video, and audio. The law focuses on the quality and origin of the datasets used by the AI models and requires public disclosure of
  - Sources and owners of the datasets
  - How the data has been used in enhancing the models
  - Types and number of data points, protected, or personal information being used
  - Usage of synthetic data generation
  - Time period when the data has been collected, if the data is still being collected
- **Utah AI Policy Act (UAIPA):** This act requires full disclosure when a person interacts with an AI system for a set of regulated services (e.g., accountants, financial advisors, physicians). The disclosure must be done either verbally (at the start of the conversation) or electronically (at the start of the written interaction). This represents the first US state law that mandates transparency obligations during interactions with GenAI systems.

At the time of writing the book, more than two-thirds of US states have in place or are about to establish AI-related regulation on topics ranging from transparency to privacy and risk management. While an EU comprehensive approach is still missing at the US federal level, this state-level regulation continues to show the legislative need to address the challenges brought by deployment and usage of performant AI systems.

Overview of EU Digital Regulations

EU sees the innovation brought by the digital technologies as a fundamental force enabling its modernization and competitive advantage in the 21st century.

To formalize its vision, in 2021 the EU produced the "Digital Decade" manifest, which outlines the long-term strategy for the digital transformation of the European Union until 2030. The "Digital Decade" seeks to define a set of rights and principles for Europeans that will "empower businesses and people in a human-centered, sustainable, and more prosperous digital future." Ultimately, the EU sees data and digital domains and technologies as a truly European project domain, and as a new and emerging technology, a pan-EU unifying force, transcending geographical boundaries.

Supporting this strategy, the EU has also released EU Declaration of Digital Rights and EU Data Strategy manifestos, which add additional depth on the principles and approach envisioned for implementing the targets set for the "Digital Decade" vision.

EU Declaration of Digital Rights: Key Principles Supporting an Effective Digital Strategy Implementation

While generally high-level, the principles outlined in the EU Declaration of Digital Rights highlight an important blueprint to follow when embarking on a digital transformation journey for any organization:

1. Respect the law: "Democratic oversight of our Digital Environments"—This principle is especially aimed at ensuring the "privately-held digital companies have understood (and acted upon) the huge power and influence they can exert over society" are fully aware and following the EU legislation regulating digital technologies.
2. Technology should be understood and used by everyone: "Solidarity and Inclusion—Technology for Everyone"—This principle aims to ensure the use of technology is truly inclusive, including for those who have less advanced technical skills and who are unaware of their digital rights.
3. Access to high-speed affordable digital access—a basic human right. In this very important statement, the EU is putting high-speed Internet connectivity at the same level with basic critical services such as water or electricity by encouraging "policy-makers and businesses to design, promote and set up digital infrastructures and solutions in such a way as to ensure everyone, everywhere (at least in the EU) may access affordable high-speed digital services"
4. The right to access education, training, and lifelong learning and to be able to acquire all basic and advanced digital skills.
5. The right to have access and be able to make informed choices:
    a. Everyone should have access to a trustworthy, diverse, and multilingual online environment.
    b. Transparency and trust in data management: Everyone has the right to freedom of expression, without fear of censorship or intimidation. Everyone should have the means to know who owns or controls the media services they are using.
    c. Privacy and individual control over data: Everyone has the right to the protection of their personal data online (how said data is used and with whom it is shared). Everyone has the right to the confidentiality of their communications. Everyone should be able to determine their digital legacy.
    d. Making informed choices about artificial intelligence: Everyone should benefit from the advantages of artificial intelligence. This involves making their own, informed choices in the digital environment and being protected against any risk or harm to their health, safety. and fundamental rights.
    e. Protecting children: Children have the right to be protected from all crimes committed via or facilitated by digital technologies, have access to age-appropriate materials, and be empowered to make safe and age-appropriate decisions.

Here are the main takeaways: The principles of data privacy, the ability to make informed choices, transparency, and trust are paramount, combined with universal high-speed digital access, lawfulness, and technology being affordable, understood, and available for everyone to learn; these are the main pillars that constitute the EU approach to digital transformation.

EU Data Strategy

To facilitate and accelerate digital innovation, the EU aims to create a single market for data following the FAIR principles: Findable, Accessible, Interoperable, and Reusable.

In this strategy

1. Data can flow within the EU and across sectors, for the benefit of all.
2. European rules, in particular privacy and data protection, as well as competition law, are fully respected.
3. The rules for access and use of data are fair, practical, and clear.

Having clear principles and a unified approach to data across EU states lays the foundation for the consequent application of digital policies and provides the necessary environment for digital innovation.

As mentioned previously, this approach is universally valid for both public sector and private enterprises and will be studied in detail across other chapters in this book.

The Wave of EU Digital Regulation

Having outlined the vision, strategy, founding principles, and data strategy, EU has implemented a wide range of policies and laws addressing the digital common market, to shape the ongoing cross-EU digital transformation. The sequence of relevant EU digital regulation acts between 2016 and 2024 is presented in Figure 2-2.

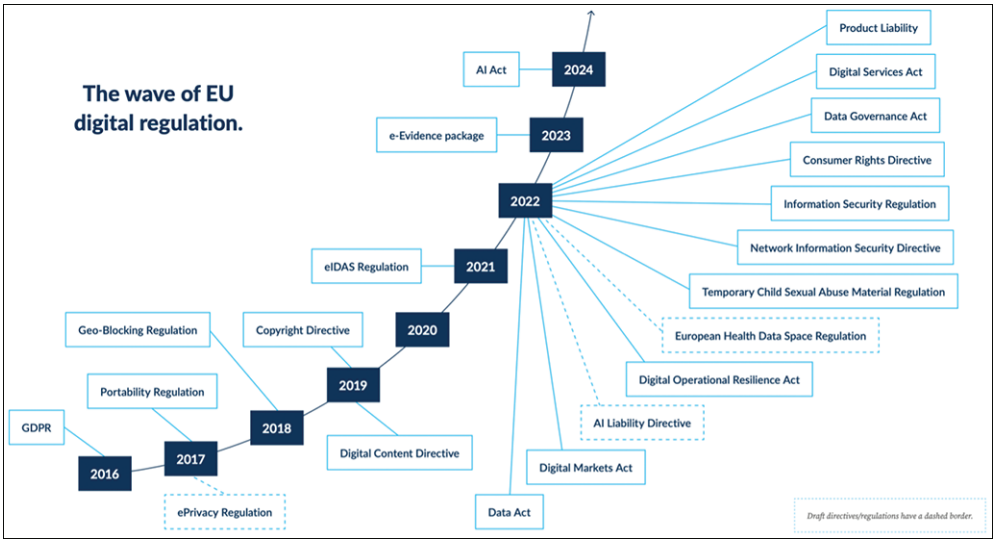


Figure 2-2 *Overview of Key EU Digital Legislation Between 2016 and 2024*

After the initial implementation of the personal data protection, privacy, and management act (or GDPR), which was in essence a horizontal law applicable across sectors and industry, the EU focused on vertical sector-specific regulation in subsequent years:

- **Content Portability and Geo-Blocking Regulation:** Content portability refers to the right of EU citizens to access paid-for online content services in other EU member states, in the same way as they would in their home country, as part of the EU digital single market. Subscribers can download and watch the same films and TV series anywhere in the EU. This relevant regulation sustains a common EU data and services framework.
- **The Directive on Copyright in the Digital Single Market:** This regulation provides a framework for fairer remuneration of content authors, especially when their content is published online, while promoting a transparent relationship with the platforms hosting their online content.

  An important provision comes in Article 17, which states that in the absence of explicit content from the original author, online providers must remove the content from their platforms. This ensures the webscalers' and other content providers' compliance with the "right to consent" GDPR provisions.
- **Data Governance Act (DGA):** This act promotes trust in voluntary sharing data mechanisms. One of its main purposes is to stimulate EU data-driven innovation and reduce the market power of large tech companies that typically sell large datasets. This law is especially important in helping smaller start-ups to get access to datasets for training their analytics algorithms and large language models (LLMs).

  DGA relies on three major pillars, also represented in Figure 2-3:

  - **Pillar 1: Use of public data for private purposes:** This pillar promotes the reuse of public sector data beyond its original purpose by non-public sector bodies and contains rules regarding how certain categories of protected data, normally outside public access, could be used.
  - **Pillar 2: Data available through intermediation services:** This pillar promotes voluntary data sharing practices between companies and facilitates the exchange of substantial amounts of data and the use of data. Among mechanisms promoting data sharing is the introduction of data intermediation services—for example, open data marketplaces—which need to follow a rigorous registration and compliance process.
  - **Pillar 3: Data altruism—voluntarily donating data:** This donation can be done by "data-altruistic organizations," which should operate on a nonprofit basis and be legally independent from any for-profit entities.

The Data Governance Act as a means to achieve the objectives of the "European Data Strategy"

Mechanisms to increase data availability	Strengthening trust in data exchange	Overcoming technical obstacles to data re-use
1st Pillar: **Re-use** of certain categories of protected data held by public sector bodies 2nd Pillar: **Data intermediation services as a** key role in the data economy 3rd Pillar: **Data altruism** – increasing the availability of data by voluntarily donating data	• Avoidance of **conflicts of interest** • **Independance** of pricing • **Structural separtion** • **Transparancy, fairness & compliance** • **Non-discriminatory access**	• Requirements - **neutrrality** and **interoperability** • Appropriate **level of protection**

Figure 2-3 *Summary of the Data Governance Act's Three Pillars*

- **Digital Markets Act (DMA):** This act aims to ensure more competition, thus preventing large digital platforms (most of the time non-EU based) from leveraging their powerful market power and facilitating conditions for new companies to enter the EU digital market. This regulation specifically targets the largest companies (named gatekeepers) operating digital content in the EU that are selected based on the volume and number of users and turnover and capitalization. When DMA appeared, six companies (the largest social media and data hyperscalers) were identified as "core platforms" and had to comply with the regulation. The list of obligations included prohibition to combine data collected from different services of the same company (e.g., Facebook and WhatsApp), rules against self-promotion, bundling, and rules aimed to ensure data portability and interoperability (aligned with "control over own data" principles).
- **Digital Services Act (DSA):** This act addresses illegal content, transparent advertising, and disinformation. Like the DMA, it applies to online platforms and intermediaries that process or contain significant end-user data and through their apps have the power to influence end-user decisions. Key requirements include disclosing to regulators how their algorithms work, providing users with explanations for content moderation decisions, and implementing stricter controls on targeted advertising, which is aligned with data processing explainability/transparency rules.
- **e-Evidence Regulation:** This regulation applies to service providers offering services in the EU. If the service provider is located outside the EU but offers electronic services to EU citizens, it must appoint legal EU representation to implement this regulation. The regulation allows judicial authorities in one member state to request electronic evidence directly from a service provider operating in another member state. Such a request (aka a production order) can be created by a judicial authority or a person who is suspected or accused in criminal proceedings. This regulation significantly simplifies the procedure to lawfully obtain user data from service providers that have EU customers.

- **EU Data Act:** This act aims to enable better access to IoT data and to strengthen the competitive fairness in user-manufacturer contractual relationships over data by providing a comprehensive set of rules and a framework for data sharing.

  It states that connected products must be "designed and manufactured…in such a manner that product data is by default, easily, securely, free of charge, in a comprehensive, structured, commonly used and machine-readable format and, where relevant and technically feasible, directly accessible to the user." This means the IoT devices subject to this act must already be manufactured in a way that respects the data requirements.

  If the data recipient wants to receive the data or share it with third parties, the data holder may charge a nondiscriminating fee and cannot impose unfair contractual terms.

  If a technical provision option to share IoT data does not yet exist, companies should introduce one in good time.

  Although the EU Data Act refers only to non-PII IoT data, if PII data coexists, GDPR rules must also be observed.

  Companies could invoke a "trade secrets" exception in sharing the data; however, this exception must be evaluated on a case-by-case basis, and the existence of such a trade secret must be proven.

  An overview of the IoT data flow, as proposed by the EU Data Act, is presented in Figure 2-4; this figure highlights the relationship between the user sharing IOT data with the manufacturer, which is the data holder sitting in between and managing end-user data sharing with commercial entities and the government.

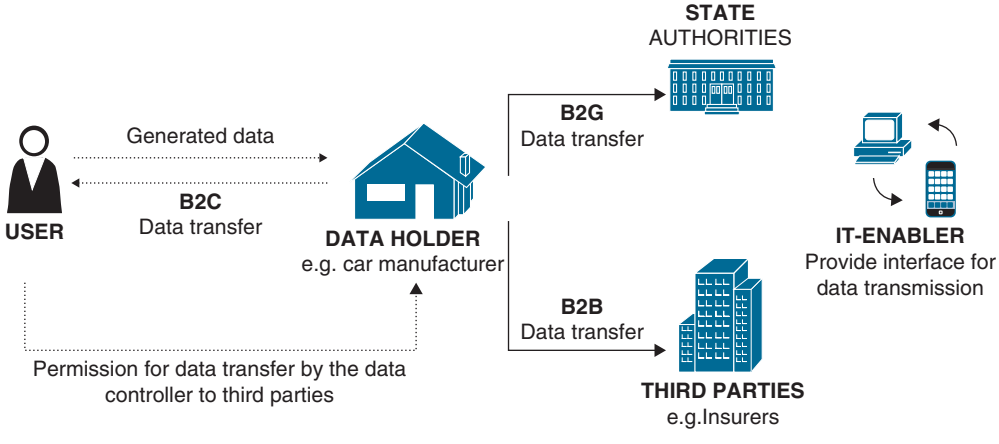


Figure 2-4 *Summary of the Data Sharing Flow Under the EU Data Act*

Note When nonpersonal data is processed together with personal data, the GDPR must also be applied. Protection of company trade secrets can be used as an exception to restrict compliance on the data sharing policy, especially toward third parties. Trade secrets have their own legislation they must follow (the EU Trade Secrets Directive) and can have primacy if they conflict with the EU Data Act. Every exception does, however, need to be reviewed and proven on a case-by-case basis.

The EU Data Act, by providing clarity regarding access to and use of data, complements the Data Governance Act, which increases trust in voluntary data-sharing mechanisms.

- **NIS2: EU Cybersecurity Regulation:** This directive mandates that each EU member state adopt a national cybersecurity strategy, including policies for supply chain security, vulnerability management, cybersecurity education, and awareness. As an important regulation, member states must also establish and regularly update a list of operators of essential services, ensuring these entities comply with the directive's requirements.

  The rules apply to communications systems and applications serving medium and large providers of critical infrastructure, such as energy, transport, finance, and water management. Additionally, digital infrastructure and service providers are also considered part of the critical services sectors. They will have to take appropriate cybersecurity risk-management measures and notify relevant national authorities of significant incidents, especially if they cause disruption or damage to the provided critical services.

  The directive also outlines audit guidelines, peer review, supervision, and enforcement guidelines, including accountability of the top management for noncompliance, thus bringing cybersecurity to the attention of the boardroom.

- **Digital Operational Resilience for the Financial Sector (DORA):** This regulation sets out requirements covering key areas of information and communications technology (ICT): risk management (including third party), incident reporting, and digital operational resilience testing.

  It sets guidelines to define the contractual and oversight relation between financial entities and third-party service providers, including these providers' participation in security audits and penetration testing.

  It also sets rules on supervisory authorities' cooperation, including supervision and enforcement.

  As such, DORA provides a sector-specific framework for the financial sector that complements and, where applicable, takes precedence over the NIS2 Directive.

- **EU AI Act**: As already mentioned in the "Short History of Data Privacy," the EU AI Act is comprehensive uniform and horizontal EU legislation aimed at promoting AI-driven innovation while protecting against its potential harmful effects. This act supports the EU's vision that people and companies are willing to take greater risks and speed up innovation as long as they trust the technology being used.

  The target AI is defined in a very broad manner as a "machine-based system that, for explicit or implicit objectives, infers from the input it receives, how to generate outputs such as predictions, content, recommendations, or decisions that can influence physical or virtual environments" to cover the entire range of AI algorithms from very simple such as regression up to complex implementations (deep learning and generative AI).

  The legislation makes a clear distinction between different actors using or producing AI systems and imposes separate obligations especially for developers (people or entities producing and putting AI systems into service) and deployers (people or entities that are using the output of an AI system for professional activities).

  The act explicitly prohibits systems that directly contradict human and digital fundamental rights such as the "right to free and informed choice and expression."

  The AI systems that significantly impact people's safety or individual rights are considered high risk and need to be properly audited, secured, documented, and governed. As a deployer, additional requirements are imposed, especially around data compliance with GDPR and intended use, incident reporting, and the requirement for end users to provide informed consent.

  A special focus is given to general-purpose, foundational, and generative AI systems (GenAI; such as OpenAI/ChatGPT models) that require an additional level of model evaluation, adversarial testing, and source data transparency. This is due to both their widespread use and impact as well as their potential to be used by downstream application users for specific business applications.

  AI systems that have limited access to end-user PII data and could mainly deceive users by providing fake images or answers or impersonating a real person are categorized as limited risk with the main obligation to inform end users that they are dealing with AI-generated outcomes.

  All other AI systems that have minimal access to PII data and can have a potential low impact on human decision-making are deemed as minimal risk with the recommendation that developers and deployers create and use them under a "code of conduct" that promotes transparency and responsible utilization.

  An overview of the risk levels and recommended actions as outlined by the EU AI Act is presented in Figure 2-5.

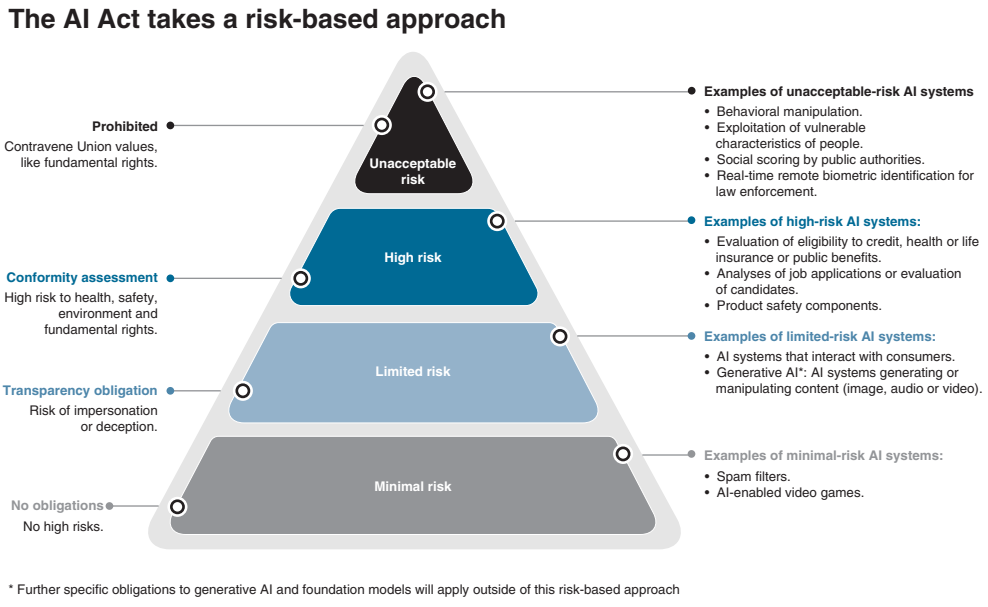


Figure 2-5 *Summary of EU AI Act Risk Levels*

The EU AI Act is relevant for large digital infrastructure providers in their role as strategic utility providers, either national (communication service providers) or with an industry scope (IT infrastructure covering critical domains such as medical, finance, public government services, justice, electricity, gas, water).

Under this act's rules, applying AI to generate insights, decisions, and actions for these network communication infrastructures falls within the high-risk category. This means that a very strict governance of AI implementation needs to be observed, including proper vetting of upstream AI algorithms being used, cybersecurity measures to protect AI systems, logging and auditing of AI activity, and not least ensuring explainability and transparency of AI decisions, especially if GenAI/LLM algorithms are being used; they are described in detail in Chapter 5, "Generative AI in the Context of Autonomous Networks."

The State of AI Regulation in the Rest of the World

As noted in the previous section, the EU regulation around AI takes a comprehensive and horizontal approach that provides a unified mandatory policy enforcement and associated fines across all industries in relation to AI system usage and compliance.

It also imposes significant fines for noncompliance (up to €35 million or 7 percent of the organization's total worldwide annual turnover for the preceding financial year) and covers any AI system or commercial offer that is used by EU customers, regardless of the geographical location of the company providing the services.

While the EU AI Act covering AI systems together with the EU Data Act, which covers IoT data, and the GDPR, which covers personal data, are considered the golden standards for AI regulation and its corresponding data usage, in the rest of the world, AI regulation is typically less unified, detailed, or restrictive. Most governments take a more lenient wait-and-see approach as they try to balance regulation and innovation in this emerging domain.

USA AI Regulation

The US legislation on AI combines federal and state laws that usually prefer to regulate topics on a per-industry level.

At the federal level, there is currently no legislation or policy regulating the use of AI because US authorities prefer to remove any barriers preventing AI innovation at a national level.

There are, however, frameworks and guidelines that intend to shape the development and use of AI, such as the White House Blueprint for an AI Bill of Rights, where five principles focused on AI safety, effectiveness, privacy, transparency, and control are outlined:

- AI systems should be developed with direct input from people and communities that will directly use them in order to enable a positive and lasting impact.
- AI systems should have embedded safeguards to protect against discrimination and bias.
- AI systems should ensure explicit consent is requested when they collect and use end-user data.
- End users should be notified when dealing with AI systems, and AI decisions should be explainable.
- End users should be able to provide feedback that changes the AI system's outcome and chooses a human decision as an alternative.

Another important framework is provided by the US National Institute of Standards and Technology (NIST), which has released the Risk Management Framework for AI (AI RMF), which articulates a methodology to develop trustworthy AI systems.

The AI RMF defines a set of characteristics that AI systems need to have to be considered trustworthy and guidance for obtaining them:

- **Validity and Reliability:** These characteristics are obtained through ongoing testing and conformance monitoring to ensure that AI systems perform as intended across different conditions and over time.
- **Safety:** An AI system should not endanger the person engaging it, achieved via safety-by-default design.
- **Security and Resilience:** Systems should be capable of defending and recovering from an attack.

- **Accountability and Transparency:** Systems should ensure confidence in the decision-making process.
- **Explainability:** They should provide information that will help the user understand the purpose and impact of the AI system,
- **Privacy:** They should protect human identity by anonymizing private data used in interactions and ensure its confidentiality.
- **Avoid Bias:** This characteristic should be provided through proper selection, processing, and usage of datasets.

NIST has also published an implementation playbook that provides guidance on ensuring compliance with its stated principles, also shown in Figure 2-6:

- **Govern:** This principle addresses policies, procedures, and practices an organization needs to employ to map, measure, and manage AI risks.
- **Map:** This principle addresses collecting context about the functioning, purpose, and usage of the AI system to better understand its potential impact and associated risks.
- **Measure:** This principle focuses on testing AI systems before deployment and continuously during their operation to ensure compliance. The framework puts emphasis on the TEVV (test, evaluate, verify, and validate) methodology to ensure AI system compliance throughout the lifecycle.
- **Manage:** This principle ensures resources are available to address the identified risks and measures based on the guidance provided by the govern function. It also includes having clear contingency measures in place to respond, recover, and communicate when an incident occurs.

The NIST AI RMF is one of the first practical guides on defining, measuring, and addressing AI risks; many private companies are using it as a reference to define their AI governance frameworks.

Certain federal institutions have also created rules for sector-specific use of AI:

- The Federal Aviation Administration Reauthorization Act provides guidance on AI usage and oversight within the aviation industry.
- The National Defense Authorization Act provides guidance on AI usage and supervision within the defense sector.

At a national level, many multinational companies, especially within the tech and cloud computing sector, have put voluntary company guidelines in place on using AI in an open, transparent, secure, and safe manner as a matter of self-regulation in the absence of country-wide legislation. This topic will be described later in this chapter.

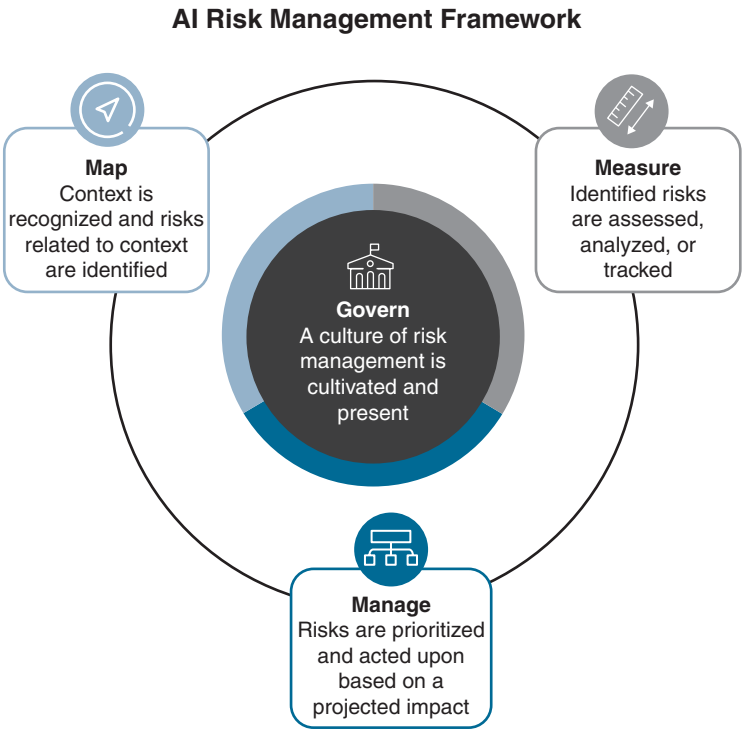


Figure 2-6 *Summary of NIST AI Risk Management Framework*

Similarly, in the absence of a general and horizontal federal law, multiple US states have created their own AI regulations with more than 40 US states having enacted planning to enforce AI laws on specific topics.

At this moment, it is increasingly difficult to obtain a clear, common set of AI rules at a national level in the US, and except for a few specific federal agencies' regulations on the use of AI for their areas of focus, companies operating in the US must do thorough work toward ensuring compliance with the various nonuniform state regulations.

Latin America: Brazil AI Regulation

Brazil is currently in the final stages of adopting its first AI regulation: the Brazil AI Act. The proposed act plans to define a comprehensive regulatory framework for AI development, deployment, and use. It also aims to ensure that the AI systems provide protection of fundamental rights, are secure and reliable, while remaining compatible with existing laws protecting human dignity, democratic values, and scientific-technological development.

The Brazil AI Act plans to follow a prescriptive risk-averse approach like the EU AI Act and split AI systems into three risk tiers with different focus and regulation intensity:

- Excessive risk (prohibited)
- High risk (heavily regulated)
- Other systems (basic requirements)

Similar to the EU AI Act, the Brazil AI Act plans to impose significant penalties—up to R$50 million per violation or 2 percent of company revenue in Brazil. Among its provisions, the policies mandate algorithmic impact assessments for high-risk systems, establish civil liability frameworks, and promote safe innovation through regulatory compliant sandboxes.

As notable policies, the law includes mandatory human oversight throughout the AI system's lifecycle, rights to explanation and contestation of AI decisions, and strict requirements for public sector AI use.

China AI Regulation

In 2017, China issued the New Generation Artificial Intelligence Development Plan, which planned in detail the development of AI from different angles, including scientific research, application, and policy support.

China has also implemented specific AI regulation, such as

- Administrative provisions on recommendation algorithms in Internet-based information services regulate the use of AI algorithms such that algorithm recommendation service providers shall not take advantage of algorithm recommendation services to engage in activities prohibited by laws and administrative regulations.
- Administrative provisions on deep synthesis in Internet-based information services impose obligations on the providers and users of so-called deep synthesis technology (deep learning, machine learning, and other algorithmic processing systems) to ensure compliance with local laws; prevent the use of deepfake services to produce, copy, publish, or disseminate information; and ensure transparency, ethical use, security, and human control.
- Interim measures for the management of generative artificial intelligence services (GenAI measures) set the rules guiding generative AI implementations, such as content moderation, responsible data selection, and usage for training purposes; tagging of generated content as being AI originated; user data privacy protection; and security recommendations.
- Scientific and technological ethics review regulation requires entities engaging in scientific research activities in life sciences, medicine, or AI to establish an ethics committee.

Contrary to the EU AI Act, China's AI regulations do not impose penalties, referring instead to the penalties imposed by the existing laws covering cybersecurity, data privacy, and civil and criminal code.

Although, in principle, the penalties imposed by the existing legal framework could be quite stiff, the AI regulations are very vague (as opposed to the EU AI Act provisions) and open for interpretation (for example, there is no clear definition of AI and generative AI). This makes compliance, implementation, and enforcement very difficult.

Japan AI Regulation

Like most of the world, Japan has taken a more gradual approach to AI regulation and has promoted a wait-and-see approach and nonbinding guidance, such as the AI Guidelines for Business Version 1.0, which outlines "unified guiding principles in AI governance in Japan to promote the safe and secure use of AI." Its nonbinding recommendations require a community approach and voluntary efforts. As in China's case, penalties for related violations would be governed by the current existing laws, such as the civil code, Product Liability Act, and the penal code.

Recently, Japan proposed a more formal AI law under the name "bill for the Act concerning the Promotion of Research and Development and the Utilization of Technologies Related to Artificial Intelligence."

Under its provisions, the Japanese government is tasked with

- Gathering information on trends in research and development and utilization of AI-related technologies
- Analyzing situations and taking appropriate measures when the rights and privacy of citizens have been impacted by improper development or use of AI technologies
- Conducting surveys and studies promoting the research and development of AI technologies

Unlike the "EU AI Act," the proposed law is rather vague in its policy definitions because its main intent is to guide rather than police the development and use of AI to support fast innovation.

There are no clear penalties for breach of the law; however, the Japanese government is being encouraged to publish a "name-and-shame" list of offenders and thus make it harder for those specific companies to do business in the future.

South Korea AI Regulation

Recently, the South Korea National Assembly passed the Act on the Development of AI and Establishment of Trust (AI Basic Act) to regulate AI systems within the country. The law seeks to bring opportunities for businesses using AI, providing responsible oversight with a clear set of ethics guidelines and reducing uncertainty on how to proceed on high-impact AI systems.

The law promotes a specific industrial growth policy, being aimed at enhancing South Korea's competitive advantage on AI with other international players. At the same time, the law bears multiple similarities with the EU AI Act, especially by advocating a risk-based approach and categorizing certain use cases as high-risk AI.

At the same time, the Basic AI Act has a different emphasis from the EU AI Act. Whereas the Basic AI Act focuses on integrating industry growth and regulatory governance with a focus on post-market oversight, the EU AI Act focuses mainly on governance regulation, risk mitigation, and compliance before an AI system is brought to the market.

The Basic AI Act also focuses only on developers of AI, balances regulation with innovation and growth incentives, imposes only moderate fines for noncompliance, and considers only foreign AI systems that have South Korean users above a certain threshold.

The South Korea Basic AI Act is considered a more industry and innovation-friendly law compared with its EU counterpart; it highlights the wait-and-see approach that most regulators are currently taking given the infancy of the technology.

Australia AI Regulation

In a similar approach to most of the world, Australia has adopted a set of guidelines on using AI that is nonbinding. Any violations are governed by its existing civil and penal laws that focus on data and consumer privacy, such as the Online Safety Act 2021, the Privacy Act 1988, and Australian consumer law.

Australia has published a number of guidelines focusing on AI, such as "AI Ethics Principles"; they are meant to act as guidance for AI development and implementation.

These principles follow the general consensus outlined by the EU AI Act, highlighting areas such as

- "AI should be beneficial and not harmful to the environment it operates into."
- "AI should respect human rights."
- "AI should be fair, transparent, explainable and provide privacy and security."
- "Human creators of AI systems should be accountable for their system outcomes."
- "AI systems should provide a consumer feed-back loop to contest and change its behavior."

While the Australian government admits its current legislation does not fully address the risks of AI, the legislation does outline the government's desire to stimulate an unrestrictive development and low-risk AI systems. To strike a balance between innovation and security, it also outlines that it is focusing on designing a legislation that "safeguards citizens, ensures transparency, testing and accountability" for high-risk AI systems, thus nodding more toward EU AI Act future legislation that is aimed at regulating high-risk AI.

State of AI Regulation Conclusions

AI is a new and fast-evolving technology that has widespread impact. It has gotten the attention of government regulators worldwide. As in every emerging technology, when regulating, governments need to strike a delicate balance between promoting innovation by being less prescriptive and promoting security and privacy.

This need for balance has split the types of regulations into three major types:

- **Security and privacy first approach, resulting in prescriptive AI laws with clear fines for noncompliance:** Here, we find the EU AI Act and the Brazil AI Act.
- **Laws striking a fine balance between supporting security and innovation, usually having a moderate level of prescriptive policies and no clear fines, leaving the punishment to be applied based on existing civil and penal codes:** Here, we find the South Korea, Japan, and China AI Acts.
- **Laws that adopt an innovation-first approach and provide a set of guidelines that should be voluntarily followed with no clear enforcement mechanism:** Here, we find countries like Australia.

Due to its size and high amount of AI innovation, the US is a special case. At the federal level, the US takes an innovation-first approach with only a set of nonbinding AI guidelines being provided while specific federal agencies provide a more binding sectorial approach (e.g., aviation, defense). At the same time, the vast majority of US states are taking matters into their own hands and regulating AI at the state level (more than 40 states have AI laws or are in the process of establishing them). This approach creates a very difficult AI legislative environment for companies doing AI-related business in the US.

Ultimately, each country will have to make a similar choice when deciding on the right approach to create its own AI legislative framework; however, the principles of privacy, transparency, security, human rights compliance, and explainability will remain relevant in relation to AI, regardless of whether they are enacted in a binding law or a nonbinding set of guidelines.

According to the UN Trade and Development Commission, in 2023 two-thirds of developed countries had a national AI strategy compared to only 6 percent of the developing nations, highlighting the early stage of AI regulatory adoption worldwide, especially among the developing countries.

Given the power of international trade, there is an expectation that in the future economic blocs with higher trade integration to adopt similar AI frameworks will facilitate doing business with each other (e.g., EU and US); however, this topic will become relevant once the AI technology becomes more established and mature.

Because of the current highly fragmented regulatory state of AI regulations worldwide, designing and implementing autonomous network systems, especially when dealing with a large amount of data, automation, and AI models driving autonomy, will need to follow a comprehensive but flexible approach to address the three tiers just described.

Different regions impose different rules on data protection, transparency, safety, and accountability. For example, in the EU, AI-driven decisions must include human oversight, especially for high-risk use cases. In less-regulated jurisdictions, systems can implement automated decisions and data processing with less regulatory oversight, but they still bear the legal responsibility if a wrong decision leads to data exposure. Therefore, most systems implementations tend to align and follow the most restrictive

legislation (in this case, the EU) because this will act as a quality mark and provide them protection in other less restrictive legislation as well.

How Companies React to Existing AI and Data Regulation

In the previous sections, we discussed the legislative framework covering data privacy and AI development and implementation. Both topics are related to emerging technologies that are linked with each other and that have experienced an exponential development. You need to have a good data strategy to derive insights, including through AI. At the same time, given its novelty, there is a global lack of alignment between countries and economic blocs regarding the intensity and punitive measures of regulating the two topics. Whereas some promote a risk-averse security-first approach by being extremely prescriptive with large associated penalties for noncompliance, others take an innovation-first philosophy, publishing nonmandatory guidelines with no associated fines. Sometimes, even at a country level, there is a misalignment between central or federal government legislation and its regional or state-level policies.

This misalignment puts most of the large, especially multinational, companies in a difficult position. What is forbidden in one place of the world can be perfectly acceptable in another and vice versa.

One of the major strategies of the private sector to address this issue was to create company-specific AI and data management policies that take a common-denominator least-path-of-resistance approach. Companies use this strategy to gain consumer trust, avoid litigation, reduce business risk, and show leadership in an uncertain and fast-changing legislative environment. In this section, we will go over some of the code of business conduct and internal policies in these companies as they relate to data and AI, showing their focus on following the most important data and AI principles highlighted in previous sections.

To be thorough, we have selected the examples of a global equipment vendor (Cisco), a global communication service provider (Deutsche Telekom), and a global enterprise specialized in managing investment in technology (Softbank).

The Cisco Approach to Data Privacy and Responsible AI

If we look at the current state of business, where more than 84 percent classify AI as being important and very important, and target business use cases to deploy AI being: IT infrastructure (84 percent), cybersecurity (83 percent), and customer experience (75 percent), Cisco is one of the key companies providing AI infrastructure, AI-driven insights, AI security, and AI use cases while leveraging its extensive IP transport know-how acquired by running the majority of Internet traffic on its equipment.

With great market power comes great responsibility, and Cisco has taken a proactive approach toward establishing a set of policies that protect customer data and ensure a responsible AI approach within the company.

Safeguarding Data Privacy: Cisco's Global Privacy Policy

Cisco's global privacy policy aims to safeguard the data privacy of Cisco's customers', partners', and employees' personal data. The following principles establish Cisco's practices for collecting and processing personal data:

- **Fairness:** The data is processed in a fair, legitimate, and transparent manner.
- **Purpose Limitation:** The data is processed in a manner that is compatible with the explicit purpose for which it is collected.
- **Proportionality:** The data is processed in a relevant manner that is not excessive for the declared scope.
- **Data Integrity:** The personal data will be kept accurate and up to date to enable accurate decision-making.
- **Data Retention and Disposal:** The data will be kept for the duration required to accomplish the stated purpose, after which it will be destroyed or anonymized to support business insights, operational improvement, or legal and compliance obligations.
- **Data Security:** Cisco will ensure reasonable physical, technical, and organizational measures to safeguard personal data and also ensure the same principles apply to the third parties processing data on Cisco's behalf.
- **Individual Rights:** Cisco will ensure compliance with applicable personal data protection and privacy laws in the countries where it operates.
- **Accountability:** Cisco will implement various tools, processes, and controls to demonstrate its compliance in processing personal data with existing privacy policies and laws.

The policy also states that when Cisco processes personal data as a processor on behalf of a customer, under the terms of a contractual agreement, Cisco not only will comply but also collaborate and facilitate the customer's compliance as well as ensure privacy by design and privacy by default.

Where there is a conflict between the various countries' applicable laws and this global privacy policy, Cisco's legal team will analyze and make a decision to harmonize the two views in a manner that is compliant with local regulations.

The current Cisco global privacy policy reflects most of the guidance and requirements present in the more prescriptive data privacy laws existing at this moment, including the GDPR.

To facilitate and allow personal data transfers between different countries and economic regions, Cisco has also been certified under various bilateral frameworks and agreements that allow for private data transfer, such as the EU–U.S. Data Privacy Framework, the UK Extension to the EU–U.S. Data Privacy Framework, and the Swiss–U.S. Data Privacy Framework.

Building on top of the data privacy principles, Cisco has also defined an ambitious set of guidelines defining the principles and methods it intends to use to ensure it builds and uses AI systems that are trustworthy, responsible, and compliant with AI regulations across the world:

As the main levers for ensuring a successful AI framework, Cisco focuses on five areas:

- An oversight committee composed of various AI experts and senior executives ensures that the policies and best practices used are compliant with the defined policy and up to date with a changing industry, legislation, and customer demand.
- Enforcing guardrail controls on the data being used for AI training and AI model reasoning to prevent bias and promote fairness, transparency, and security.
- Putting in place a robust incident management framework where any deviations of AI model output that impact human decisions or rights, such as bias and discrimination, are promptly analyzed and reported to the oversight committee for corrective action.
- Ensuring the framework and policy stays up to date and follows industry trends by participating in various industry forums where AI controls and policies are discussed.
- Continuing to promote responsible AI innovation by collaborating with external entities involved in research and development of responsible AI.

A summary of the principles is shown in Figure 2-7:

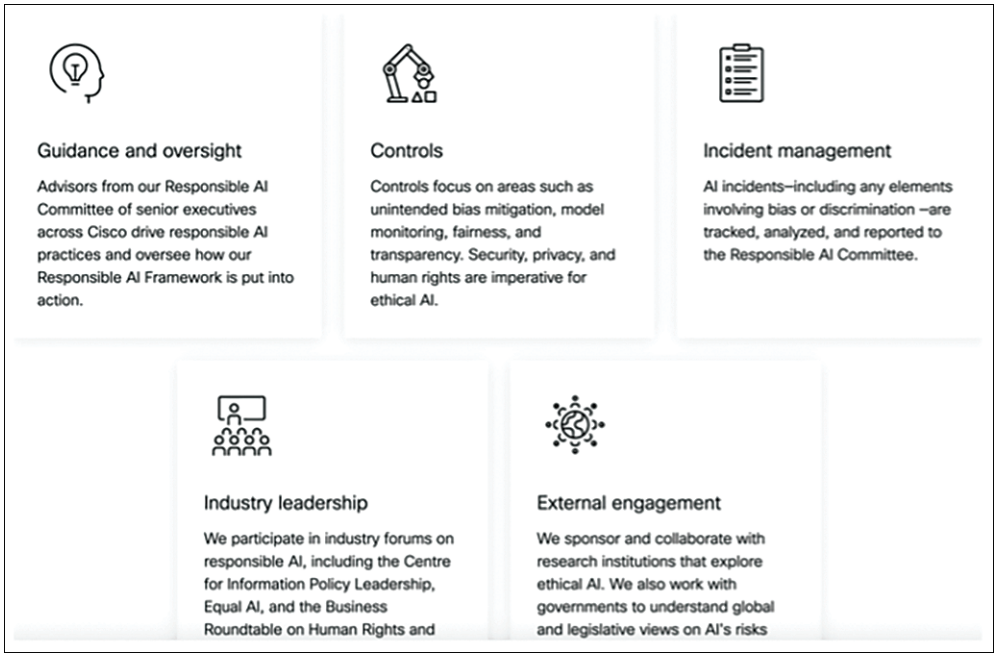


Figure 2-7 *Summary of Cisco Principles for Responsible Artificial Intelligence*

The responsible AI guiding principles follow the key global best practices recommendations with a clear focus on

- **Transparency:** The framework informs users when and how AI is employed in Cisco's technologies, the purpose and expected outcome of the AI being used, the model and data being used for training details, security, privacy, and human rights controls applied.
- **Fairness:** The framework requires documentation of AI use cases, impact assessments, and appropriate oversight by a group of cross-functional leaders. This also covers allowing users to provide actionable feedback on how the AI model can be improved.
- **Privacy:** Cisco has built privacy engineering practices into and leverages Cisco Secure Development Lifecycle (CSDL), which is used for software development, including AI-based products. CSDL is a company-specific software development methodology that allows you to design, build, and operate software with security, including data privacy being built in.
- **Security:** Next to privacy, CSDL also allows for the creation of secure development lifecycle controls, which provide protection against security threats and include testing the resilience of AI systems against cyber attacks.
- **Reliability:** As part of the mandatory AI Impact Assessment process, Cisco reviews its AI-based solutions to determine whether adequate controls are embedded in their lifecycle to maintain consistency of purpose and intent when operating in varying conditions and use cases, effectively protecting against drift or misbehavior in diverse environments.

To ensure global coverage, Cisco's guidelines are built based on NIST AI RMF recommendations, which are broad enough to address existing AI regulations. At the same time, Cisco leaves the door open to further adjust their implementation to ensure compliance when specific incompatibilities are present with local laws.

Cisco is a signatory of the EU AI Pact (an initiative launched by the EU commission to help companies prepare for the AI Act provisions) and follows the guidelines of the EU AI Act, being expected to be fully compliant when its main provisions come into force.

With its strict and comprehensive policies covering data privacy and responsible AI, Cisco is ensuring legal and commercial compliance with a diverse set of global requirements, thus ensuring its business resilience.

Deutsche Telekom Approach to Data Privacy and Responsible AI

Deutsche Telekom (DT) is the largest telecommunications service provider in Europe and one of the largest in the world. DT started its AI journey early, first by providing Copilot-like chatbots to assist operations, introducing AI into its cybersecurity and threat detection platforms, and then embedding AI into its critical network operations such as guiding and optimizing its fiber network rollout.

Introducing AI internally has been not only a technical upgrade but a fundamental culture shift. To address the privacy expectations of both its customers and employees, supporting this culture transformation, DT has realized it has to elevate the focus on fair use of AI and privacy data protection at the corporate board level and embed it into its fundamental company policies and practices.

To put this plan into practice, DT was one of the first companies in the world to introduce "AI guidelines on digital ethics" in 2018 (naming many of the globally accepted principles, five years before NIST AI guidelines!), highlighting the need for a framework that ensures AI is

- **Responsible:** Clearly defining who is accountable/responsible for which AI system outcomes
- **Careful:** Ensuring AI follows human-made rules
- **Supporting:** Having a customer-first outcome as a target
- **Transparent:** Clarifying when communication with AI takes place and how customer data is being used
- **Secure:** Protecting customer data
- **Reliable:** Ensuring the AI systems are compliant with the guidelines and outcomes being set
- **Trustworthy:** Ensuring human control to eliminate AI system misbehavior
- **Cooperative:** Ensuring AI systems work in a symbiotic collaboration with humans
- **Illustrative:** Providing, by design, the necessary documentation and skills for humans to work with AI systems

The AI guidelines have been elevated at the same fundamental corporate level as the code of conduct, guaranteeing its fundamental impact on the corporate culture.

DT has taken an additional step toward ensuring the AI guidelines become a fundamental part of its supplier code of conduct to ensure the entire supply chain alignment and conformance with its guidelines, a move that ensured early compliance with NIS2 specifications:

- "Data processing and AI algorithms, as well as the use of the data have to be documented in a comprehensible manner, to be disclosed to DTAG if required."—Illustrative principle
- "The development and usage of Artificial Intelligence is carried out according to the European fundamental values."—Reliable principle

- "The use of AI systems shall be non-discriminatory and transparent toward users."—Transparent principle
- "The Supplier has clearly defined for his AI solutions who is responsible for which system and which AI function."—Responsible principle

DT has highlighted the data privacy topic by naming a global data privacy officer who coordinates and addresses all topics related to data privacy within the company. This role reports directly to the board of DT, thus highlighting the empowerment and significance ensuring data privacy has within DT.

To address the changing needs of data privacy protection within the AI context, DT also defined the concept of "privacy by strategy" which enhances the previous "privacy by design" approach. Its main purpose is to ensure data protection is incorporated in any AI-related initiative much earlier than the design stage (the current approach). This approach advocates that the data protection organization must be up to date with upcoming digital trends in the market and current best practices to assess the impact of data privacy and make a recommendation to the decision-makers at the strategic conversation stage. The ultimate goal should be a clear assessment of whether a specific AI strategy and functionality are feasible to implement from a data privacy perspective. This approach is a good way to address the risk categorization introduced by the EU AI Act (e.g., where certain prohibited use cases are stopped from the start and others must be highly reviewed for their potential high-risk impact). Having a "privacy by strategy" approach will also allow DT to estimate the correct budget, operational implementation practices, and risks, especially when the decision is made to implement a system that can have a significant impact on end users' privacy and decision-making.

DT has ensured its principles are accepted and aligned with the workers' representatives (workers' council) and are adopted as part of the AI governance process within the company, serving as a framework for its technical, steering, and regulatory authority decision.

To ensure compliance with the EU's evolving legislation, DT is also a signatory party in the EU Commission AI Pact and continues to collaborate with German and EU industry associations and regulatory bodies on the topic of AI.

The entire set of processes governing DT's approach to digital responsibility, which includes digital ethics, data privacy, laws, and regulations, is defined in the "House of Digital Responsibility" approach, which ensures DT has a human-centered digital ethics process at the core of its technology evolution toward the future (see Figure 2-8).

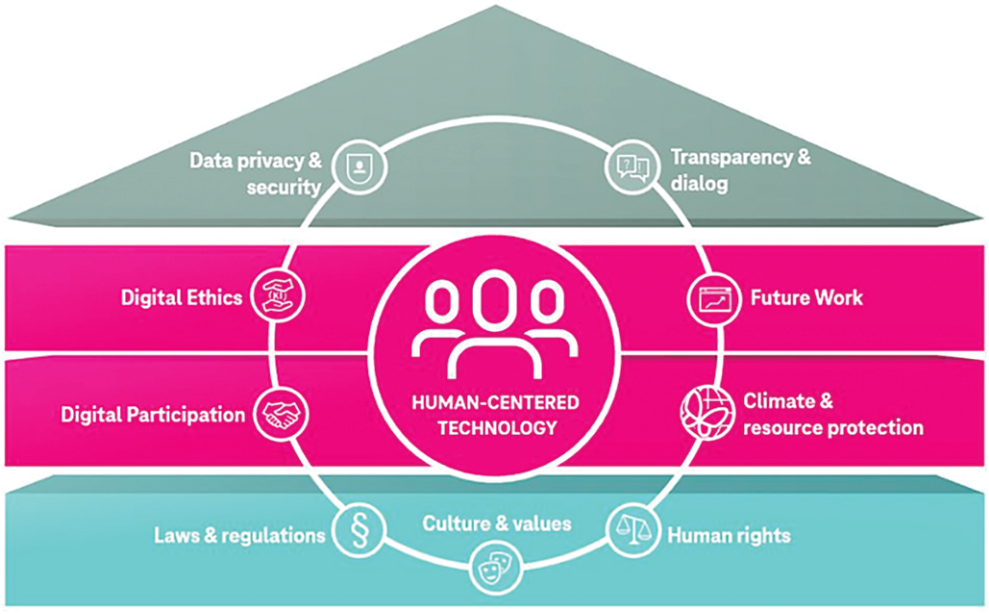


Figure 2-8 *DT House of Digital Responsibility*

Softbank Approach to Data Privacy and Responsible AI

Softbank is a Japanese multinational investment management company that has a primary focus to invest in technology companies that offer services across the world in industries such as automation and Internet connectivity.

The group focus is to stay at the forefront of the Information Revolution, which is largely defined by the outcomes derived from data insights. AI is naturally at the core of the Information Revolution, and as an aspiring leader, Softbank has seen the strategic benefit of being an early adopter of the responsible AI framework.

Softbank has established an AI Governance Working Group under the Group Risk and Compliance Committee (GRCC) to discuss an approach to AI governance.

The Governance group is composed of executives leading key divisions of the group (CEO, legal, information systems). Its purpose is sharing best practices and making decisions around the implementation and adoption of AI within the Softbank group divisions.

The AI Governance group, being only one level below directly reporting to the company board of directors, shows the strategic importance of its activity and oversight (see Figure 2-9).

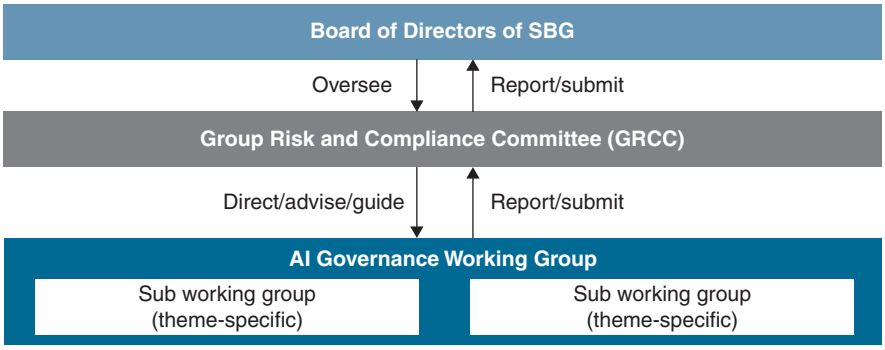


Figure 2-9 *Softbank AI Governance Working Group Reporting Structure*

Just like Deutsche Telekom, Softbank has defined its guidelines on generative AI and has included them as part of the group code of conduct to emphasize their fundamental importance.

In 2022, the company also established an AI Ethics Policy and set specific rules and checklists for AI-related planning, application, development, and outsourcing.

The AI Ethics Policy principles are in line with the other company's philosophy:

- **Human-Centered Principles:** The AI outcomes should be centers to solve real social problems.
- **Respect for Fairness:** Respect all people regardless of race, nationality, age, religion, or gender. Ensure social minorities are represented in the AI training data to avoid bias.
- **Pursuit of Transparency and Accountability:** Create an environment in which people can understand the results of AI-based judgments and verify the basis of those judgments when necessary.
- **Safety Assurance:** Strive to design and develop AI that does not threaten the life, liberty, dignity, or property of individuals.
- **Privacy Protection and Security Assurance:** Appropriately manage various types of data, including personal data, in accordance with relevant laws and regulations to prevent unauthorized access and protect privacy.
- **Education for AI Talent and Literacy:** Actively engage in education and training to enhance the skills and literacy of both the developers and users of AI.

Softbank is also a participant in the G7 Hiroshima AI Process (HAIP), which is a voluntary reporting framework to encourage transparency and accountability among organizations developing advanced AI systems.

Under its reporting, it highlights using AI risk levels grouped in four categories—Prohibited, High, Mid, and Low—considering the social and business impacts.

Softbank evaluates AI risks based on two axes: first, categories delineated by the EU AI Act and high-risk sectors defined by Japan legislation, and second, an internal risk scoring based on scale of users and revenue.

While Softbank, as a large investor in cutting-edge technical businesses, has naturally kept up to date with the international legislation covering data privacy and responsible AI, most of the enterprises are still early in the process to follow suit. With a large high-tech multinational leading the internal regulations on responsible AI due to their global exposure, communication service providers are a notable exception among national large enterprises in regard to early responsible AI policy definition given their public and trust profile.

Regulation Impact on Data and AI Projects and Mitigation Measures

AI is one of the major technology trends followed and deployed by companies worldwide. At this moment, more than 78 percent of global companies declare they are using AI for at least one line of business (up from 55 percent in 2023). Most AI use cases cover IT, marketing, sales, and service operations, with IT having the most accelerated adoption.

Similarly, generative AI has seen a jump to 71 percent adoption within businesses compared with 65 percent one year earlier with areas like marketing, sales, product development, software engineering, and service operations deploying the bulk of the use cases.

From a business perspective, according to a UN Trade and Development study, the AI market is predicted to raise from 189 billion USD in 2023 to 4.8 trillion USD by 2033, a 25-fold increase. This puts AI adoption, development, and implementation at the top of companies' strategy, being one of the main topics discussed in boardrooms.

Within the context of accelerated AI adoption, one of the main guardrails controlling widespread AI adoption is data privacy and AI fair and transparent use of private data. A Cisco 2025 Data Privacy Benchmark Study has shown that 95 percent of customers will not engage in a commercial activity if they feel their data is not properly protected. At the same time, 99 percent consider that external privacy certifications are a key vendor selection criterion, with 50 percent being aware of their country privacy regulations (see Figure 2-10).

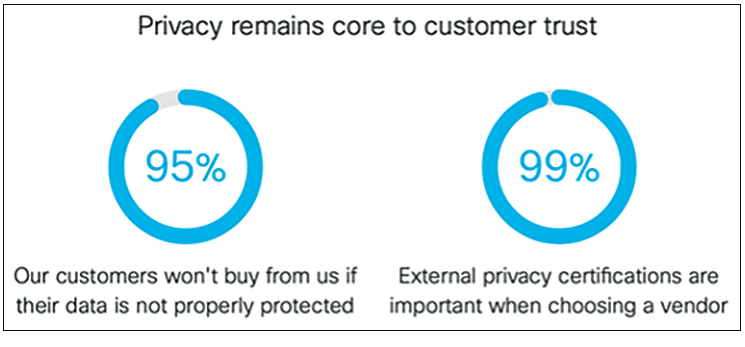


Figure 2-10 *How Important Is Data Privacy To Customer Trust?*

Solid data privacy also requires similarly solid data usage into AI tools deriving insights from it. With 64 percent being concerned about sharing sensitive information by mistake through public AI tools, this results in an overwhelming expectation that private data shared with companies should be used in an ethical manner, including for AI application (97 percent) and that a strong adoption of privacy laws would help make customers more comfortable with their private data being used by AI applications (see Figure 2-11).

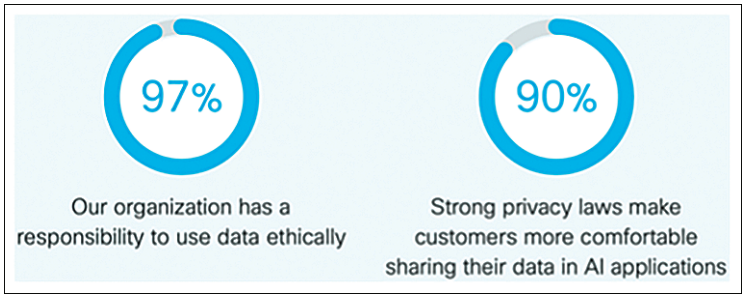


Figure 2-11 *How Important Are Ethical Business Practices and Privacy Laws to Customer Trust?*

These considerations put the formal organizational implementation of data privacy and AI fair and transparent usage at the forefront of companies' focus on adopting AI for commercial usage, with 96 percent of organizations reporting that the returns on privacy investments are an ROI-positive investment.

The organizational-structured implementation of data privacy and AI transparent and fair usage of data is mainly achieved through three mechanisms:

- Rigorous policy definitions (described in the previous section)
- An organizational setup where data and AI governance are handled as a strategic company focus
- Proper security measures

Organizational Setup

Establishing a proper and robust company governance to ensure data privacy and AI fair usage is a key topic for company executives. A recent study from the International Association of Privacy Professionals (IAPP) has highlighted that more than 69 percent of companies' chief privacy officers have taken additional responsibility in defining and running AI and data governance and ethics in their organization, highlighting the priority focus this has in organizations worldwide. This result has been documented as part of a study run by the IAPP (see Figure 2-12).

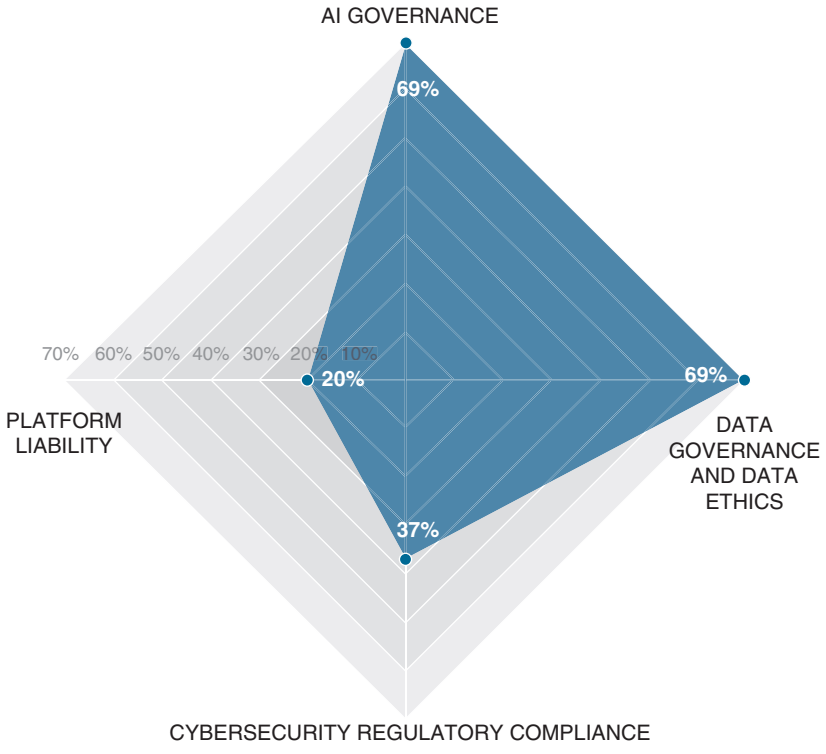


Figure 2-12 *The Expanding Remit for CPOs*

The challenge: Many companies suffer from role and responsibility overlap between different functions dealing with

- **Privacy:** Covered by the privacy office and chief privacy officer
- **Cybersecurity:** Covered by the information security office and chief information security officer
- **Technology Vision and Roadmap:** Covered by the engineering office and chief technology officer
- **Regulatory and Legal Compliance:** Covered by the legal and compliance office and chief legal/compliance officer
- **Data Usage and Manipulation:** Covered by the data management office and chief data officer
- **Risk Management:** Covered by the risk management office and chief risk officer
- **Procurement:** Led by the procurement office and chief procurement officer

This challenge causes misalignment and misjudgment, especially with the additional complexity imposed by AI transparency and data privacy, which requires a global, unitary, and cross-functional approach.

Answers to simple questions become complex topics that keep boards in endless debates. For example: Where does a company govern data privacy/AI transparency? Is it under the Privacy function (as it deals with privacy), Data function (as it deals with data), Regulatory function (as there are regulatory guidelines to be implemented), Procurement function (as most companies source data and AI processing software externally), or CTO function (as this is a new market trend)? Adding on top of that the speed at which data privacy and responsible AI legislation are changing, this puts the company in an impossible position given the rigidity and lack of alignment of its internal siloed structures.

The solution: At this moment, most companies are looking to unify data and AI governance under a single digital governance committee which has representatives covering each major function that deals with data and AI (e.g., Privacy, Cybersecurity, Technology, Legal, Data, Procurement, AI).

This specialized committee (illustrated in Figure 2-13) ensures compliance with existing and new legislation around data privacy and AI ethics, estimates and reviews the privacy and ethics risks by adopting novel technologies (new AI algorithms, new AI use cases, new data sources being ingested), ensures cybersecurity is implemented by design on existing and new AI and data systems being deployed, ensures data is handled in an ethical manner, and makes sure the company suppliers provide sufficient evidence that their software and handling of data follow the company privacy and ethical AI policies.

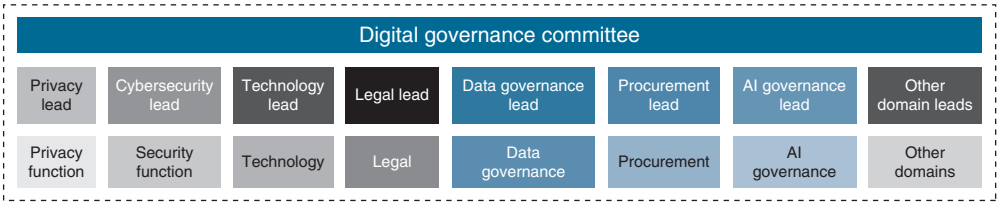


Figure 2-13 *The Digital Governance Committee*

While this is a desired state and implementation is still ongoing, the initial steps are already visible in large multinational companies that are at the forefront of digital transformation. To handle the complexities and dynamism of ensuring compliance with responsible AI and data privacy legislation, the future lies with having a multifunctional specialized team providing unitary advice to the board and business operations.

Security Control Measures

While trying to ensure the data privacy and AI transparency and fairness are maintained, the companies split this task into two major topics:

- Ensuring they follow and maintain compliance with data privacy regulations by deploying the appropriate control and protection measures

- Ensuring they control the AI systems to provide transparency and fairness by deploying appropriate AI defense mechanisms

To comply with various data privacy laws, there is a key difference between personal data (as defined by the GDPR) and personally identifiable information (as defined by the US Office of Privacy).

While PII usually covers only data that can be used directly or in combination with other data to identify a person, the GDPR takes a broader approach and defines information that can be used indirectly to identify a person as private data. This means certain types of data, especially coming from the IT department (e.g., IP, cookies, device ID), are personal data under the GDPR but not PII and will require different treatment according to the legislation being followed.

An overview of the main differences between personal data, PII, and nonpersonal data is presented in Figure 2-14.

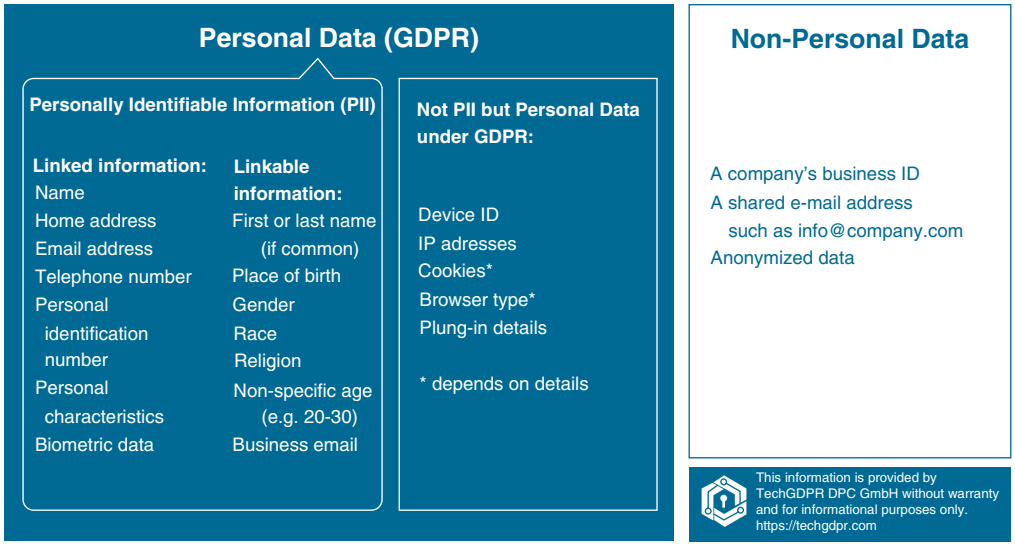


Figure 2-14 *Differences Between Personal Data, PII, and Nonpersonal Data*

The data collected and used by companies can be largely split into two major categories, assuming the more strict definition of GDPR personal data is used:

1. **Personal data:** This definition covers any information that is related to an identified or identifiable natural person. This includes all data that is or can be assigned to a person in any way. For example, phone number, financial information, Social Security number, number plate, customer number, and address are all personal data. This data falls under most countries' privacy laws (e.g., GDPR) and must be highly protected.

As mentioned previously, from a GDPR perspective, IP addresses, traffic flows, access logs, error logs, cookies, and security audit logs can be considered personal data if there is a technical way through which they can be linked to a person.

2. **Nonpersonal data:** This definition covers all the other types of data that do not include or cannot be used to identify a person directly or indirectly. This type of data usually refers to products, services, or financial data produced or used by companies (as opposed to individual users). From an IT perspective, this data usually addresses operational data linked to devices (platform telemetry, performance metrics, and alarms) rather than customers. Data collected from IoT sensors also falls in this category, and the rules allowing for it to be legally shared are provided by the EU Data Act.

The challenge: Given personal data is critical for a company; this is the prize most sought after in data breach attacks. A data breach can happen under three major attack vectors:

- **Loss or Theft of Storage Devices:** This type of attack can be mitigated through proper data encryption at rest, such as the Advanced Encryption Standard (AES).
- **Insider Attack:** This type of data breach is caused by an employee leaking information to an unauthorized third party. It can be mitigated by implementing access control measures, such as the principle of least privilege, enforcing strong passwords, and utilizing multifactor authentication (MFA). Using a strong user monitoring system and implementing data loss prevention systems are also meaningful measures to help address this threat.
- **Targeted Attack:** In this type of attack, an individual or group of attackers target specific individuals or organizations to obtain confidential information. As main methods, they use phishing (social engineering), malware, or vulnerability exploits. Rigorous user training, cybersecurity monitoring, antimalware protection, software lifecycle, and firewalls are some of the main measures used to address these attacks.

Organizations can ensure their AI models behave in an ethical and transparent manner by building a process that ensures nonbiased data is used for training, performing causal and what-if analysis to understand the model behavior and ensure model robustness for a multitude of data inputs and contexts, applying fairness constraints during model training, performing rigorous model error assessment, building model explainability features in the model, by default, by ranking feature importance in providing certain outputs, and understanding the dependence of outcomes on certain input parameters.

Specifically, for LLMs, fairness can be achieved by implementing a set of security guardrails that ensure content is checked against inappropriate, offensive language or sensitive topics. Once the answer has passed the security guardrails, additional filters can be applied to ensure the answer is relevant (on-topic) and fact-checked.

LLM security guardrails can also be centralized and adjusted to specific policies through the deployment of an AI firewall; this device sits in between, inspects, analyzes, and polices users' prompts and LLM responses. This setup ensures models and data are

behaving according to the fairness principles provided by the relevant legislation and internal company policies.

The challenge: Given an AI system's ethical output in interacting with its users is a key element reflecting the company's image, this is a preferred target in attacks against AI systems, making their security paramount.

Figure 2-15 shows a typical AI LLM system setup covering both production and development, the major potential attack vectors and mitigation measures.

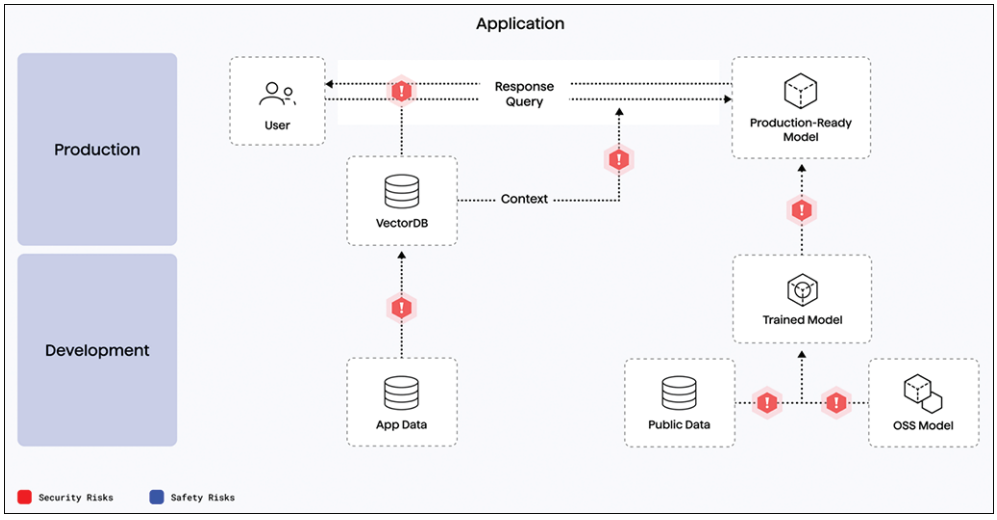


Figure 2-15 *GenAI Development and Production Typical Security Risks*

The attack vectors are as follows:

- **During development:**
  - **Training Data Poisoning:** This type of attack involves the deliberate manipulation of training data to compromise the integrity of an AI model (e.g., inserting bias in the data). This attack can be mitigated with both proper model guardrails and data sanitization techniques.
  - **Model Backdoor:** This type of attack involves inserting backdoors (triggering predetermined actions or answers based on specific inputs) into the model. This attack can be mitigated with guardrails and rigorous model verification (large training dataset, what-if analysis) prior to its deployment.
- **During production:**
  - **Direct or Indirect Prompt Injection:** This type of attack involves altering the model output and behavior by providing instructions that alter the initial request either directly or indirectly (via documents or images). This attack can be mitigated by model guardrails and the deployment of an AI firewall.

- **Sensitive Information Disclosure and Data Exfiltration:** This type of attack involves confidential or sensitive data being exposed through vulnerabilities in an AI application. This attack can be mitigated by using the AI firewall to control the data being exposed but also by employing appropriate data sanitization techniques, preventing sensitive data from being exposed to the LLM.
- **Privacy Attacks:** This type of attack involves model extraction where a functionally equivalent model is created by probing production model and data membership inference, which is used to determine which specific datasets were used to train the model. This attack can be mitigated by using data scrubbing techniques to ensure a robust dataset and by deploying specific policies on the AI firewall to prevent extraction of model/training dataset details.

Impact of Regulations on Autonomous Network Implementations

Autonomous networks, the focus of this book, are essentially a framework to allow independent self-management of network operations with minimal human intervention by leveraging a set of AI algorithms, automation use cases, and large network data collection and processing. This will be described in detail in subsequent chapters of this book.

Implementing autonomous networks depends on collecting large amounts of network observability data, including network telemetry to drive analytics, observability, AI/ML training, and inference and creating triggers for executing automated actions.

Modern privacy regulations treat network data and telemetry as a regulated governance of personal or personal-related data because it contains elements (IPs, device IDs, etc., as highlighted in Figure 2-14) that can lead to a specific individual identification. This adds specific requirements and constraints that the data collection implementation needs to account for, specifically:

- **Proportionality and Necessity:** Data must be collected for a specific reason or use case and not "just in case." Sampling and anonymization of sensitive fields are performed whenever necessary.
- **Purpose Limitation and Transparency:** Data must be used for clear purposes that are visible to users and regulators: Metadata tags must be present and highlight the purpose of each dataset being collected. There needs to be a clear separation between data collected for different purposes; for example, analytics that help understand user behavior must be processed separately from network operations data.
- **Retention and Deletion Practice:** Regulations dictate that data cannot be stored indefinitely for historical analysis and AI model training. This means that companies need to implement a tiered storage architecture that holds different types of data for different retention periods; they also need strict policies while ensuring automated data expiration and deletion. Data removal must also be propagated across all systems, including backups.

- **Sovereignty and Data Transfer Regulations:** Data must obey geographic and legal authority boundaries, which, in practice, means implementing regional data collectors, storage, and processing within specific administrative boundaries and ensuring strict control of cross-regional data transfers.
- **Accountability and Auditability:** Operators must prove lawful handling of data by enabling immutable access logs and deletion verification, ensuring that data queries are logged, evaluated, and traceable, and that overall, the data processing systems can generate valid compliance reports in line with applicable data privacy laws.

These principles apply broadly to the entire dataset and telemetry processed by autonomous systems; however, the enforcement can vary from being more relaxed when working with low-risk data (e.g., device-specific CPU utilization, packet loss, etc.), stricter when working with medium-risk data (e.g., data that contains personal elements like IP addresses or device IDs), and being the most strict when managing sensitive data such as content payloads, visited URLs, and browsing history.

Looking at the typical architecture principles and guidelines that have been created to address the data regulation requirements set by frameworks such as the EU's GDPR, US's CPRA and NIST, China's PIPL, and Japan's APPI, we see the following patterns:

- Data collection minimized through a data collection gateway that enforces feature whitelists (ensuring only specific data fields and key performance indicators [KPIs] are forward for processing), anonymization of sensitive data fields (e.g. IP, MSISDN, Device ID), and timestamp coursing.
- Purpose-segmented data ingestion pipelines, operating data through distinctly separated data pipelines; for example, operations data stream, AI/ML training data stream, and data used for incident forensic investigation.
- Automated policy-driven retention engineering, ensuring that a specific data retention limit is set for each data type (e.g., 7 days for raw telemetry, 120 days for aggregated metric).
- Implementing decision ledgers to track data to decision traceability and ensure datapoints such as policy ID, feature type, data fields, model being used, and confidence score are properly tracked.
- Destination-aware data routing engines, ensuring network data and telemetry includes metadata disclosing where the data is sent and for which reason, allowing the data routing engines to block unauthorized transfers.
- Controlling data use for secondary unintended purposes by, for example, implementing a security data tag that prevents the data with this tag being reused for user behavior analytics pipelines, to train marketing prediction models or in assessing user work productivity.
- Implementing geography-specific data collection, storage and processing, ensuring the data collected within a specific jurisdiction is processed and used within that specific region by implementing regional data lakes, inference, feature engineering, and policy engines.

- Implementing a metadata-first telemetry design by collecting and processing non-content, descriptive metadata about network activity instead of inspecting or storing the actual payload or user content as part of the autonomous network data usage.

Automated action execution and autonomous decision-making are also constrained by the data privacy laws discussed previously, with the additional constraints introduced by AI acts that exist in various countries: EU AI Act, NIST AI RMF, Japan AI governance and business guidelines, and Brazil AI Act, resulting in the following architecture implementation principles and guidelines:

- Policy-constrained automation execution demands that automated actions must be executed only after passing jurisdiction-aware policy validation. This means that every closed-loop action is gated by a code-executable policy that ensures automated actions dynamically adapt behavior per region (e.g., in the EU region, a medium-risk routing policy change that relies on pseudonymized data such as IP address hashing requires human oversight).
- The bounded execution radius principle states that automated actions and autonomous decisions must be technically constrained in scope, time, and impact, ensuring they are applied to specific network segments, rigorously tested (e.g., via canary deployments), and have well-defined auto-rollback triggers and procedures.
- Automation/autonomous decision traceability requires specific datapoints such as input features, model version, policy decision, and model outcome to be recorded in an audit-grade fashion for every automated action and autonomous decision being performed.
- Automation frameworks must align orchestration logic with the decision-making AI/ML model governance; for example, automation cannot execute with unapproved model versions; similarly, models classified as high risk within the EU require human oversight as an alternative.
- Every defined automated action must be reversible; to achieve this automation, controllers have to define and store a feasible rollback procedure for every executed automation action.
- Traceability of automated decision-making requires audit-grade evidence of the decision origin for every automated action: AI, automated scripted rule, or human oversight.
- Differentiated risk classification needs to be implemented per automation use case. The automation use case must carry a formal risk class that controls how autonomously it can operate: high risk (slice deletion) implies limited autonomy and human oversight; medium-risk (traffic engineering reroute of an L3VPN) implies conditional autonomy, which is a combination of human oversight and automated decision-making for specific subtasks; and low risk (e.g., QoS configuration change) means being fully autonomous.

AI and data privacy standards are equally relevant in governing the implementation and lifecycle of traditional AI/ML and LLMs for autonomous decision-making and data-driven insights generation, driving specific architecture principles:

- AI models need to be treated as controlled "decision components," not free-form intelligence, having an explicit intended purpose, defined action-space, and risk class that governs what it's permitted to influence. For example, an LLM-based agent may be allowed to define a network change plan but is not permitted to automatically generate configuration and deploy it to network devices; this operation is possible only by leveraging a software-defined networking (SDN) controller. The agent can select and trigger a predefined use case existing on the SDN after rigorous policy checks.
- Autonomy must be defined in a multitier manner with each tier being aligned to the perceived model risk and action impact. Typical tiers are recommend (the model makes a recommendation and stops), review and approve (the model recommends a specific action that is reviewed by a human operator), and execute (the model executes the action directly) with nondeterministic LLMs usually sitting in the first two tiers while the more deterministic AI/ML models are in the third one. For example, a high-confidence ML output can automatically tune interface queue parameters while the LLM output to instantiate a new network slice must be treated as a recommendation going through human approval first.
- Policies must exist and be placed in the information flow between the model output and corresponding automated actions to ensure regulatory requirements such as data usage, jurisdiction, safety, and change management rules are being followed.
- The separation of LLM public "knowledge" from privacy-regulated data must be ensured using retrieval-augmented generation (RAG) techniques to provide separation of topics, such as public data, internal runbooks, and logs containing private data.
- An impact assessment must be done before implementation for new AI-driven closed-loop use cases, covering privacy impact assessment, AI risk assessment, and documented mitigations (addressing topics such as oversight, action rollback, and autonomy limits).
- An algorithmic "explanation interface" for autonomous network operators and users must be provided by design: explaining why the action was taken, what the algorithmic chain of thought was, and what data influenced it. For example, if an autonomous decision is taken to reroute certain traffic, the system must show the operator the decision drivers (SLA was falling below a specific threshold), what additional data was taken into account (service status, number of users, configuration parameters), decision steps (why certain routing protocol changes were decided), and an auditable way to track the set of actions and results being achieved.

Data, AI, and system operation sovereignty is a particularly important topic, especially when dealing with systems and enterprises spanning across multiple jurisdictions. To comply with the existing and emerging legislation, organizations must design and operate for sovereignty across four distinct layers:

- Data sovereignty is architecturally best addressed through implementing geo-fenced storage, legislation-compliant data management policies that are adapted for the legislation they must adhere to, encryption that is following the national/regional cryptographic guidance, while ensuring the managed data never leaves the specific "sovereign" territory it must reside on, except through formally approved and compliant transfer mechanisms.
- Technology sovereignty is achieved by structuring and using AI models, critical software elements, and cryptographic instruments, including their licensing, deployment, and right-to-use within the given geographical jurisdictions. Additional usage of open-source, open standards and transparent supply-chain attestations can be leveraged to ensure the technology complies with the target legislation.
- Infrastructure sovereignty is ensured through physical deployment that provides compliance with the host jurisdiction of the entire infrastructure stack of network and data center assets. Hardware attestation and chain-of-custody governance are also key elements in providing the required transparency on the hardware being deployed and used.
- Operational sovereignty ensures that all the key operational tasks such as service administration, incident response, access governance, and AI model lifecycle control are being executed from legally approved jurisdictions, including the localization of the teams operating them. This ensures legal compliance and strategic independence of critical network infrastructure operations, including data, application, and infrastructure management; lawful interception; logging; and security access.

When deployed together, these four layers ensure a truly sovereign network infrastructure and operations that ensure strategic autonomy and regulatory compliance while preserving control and operational resilience.

Examples of Regulations Impacting Autonomous Network Implementations

While public data highlighting the direct regulatory impact on autonomous network implementation is scarce, regulation has a real and persistent impact in how autonomous networks are architected and deployed.

In the US, there are public documented cases of fines given by the Federal Communications Commission to major US carriers to punish unlawful sharing of real-time location data (CPNI) to third parties. This has a direct implication on any AI/ML model and use case that uses fine-grained location to drive internal network operations use cases, such as optimization, slicing, congestion mitigation, and fraud controls. The FCC fines forced US operators to use coarser-grained aggregated

mobility data and clustered KPIs for their network operations use cases. At the same time, US operators kept features like subscriber-level location out of closed loops and implemented stricter vendor data sharing controls to prevent future FCC fines.

In Sweden, the national data protection authority (IMY) penalized Tele2 over the use of the Google Analytics platform, which involved transfer of personal European user data to the US without adequate protection.

Because many "autonomous" stacks centralize observability + AIOps into global SaaS platforms, if telemetry includes identifiers (even pseudonymous) and flows to tooling outside the EU without compliant safeguards, it can and will trigger significant fines. The obvious architecture solution is implementation of a sovereign AIOps stack managing personal data on national or EU territory and jurisdiction.

Challenges and Opportunities Created by Regulations

The evolving AI and data privacy regulations worldwide are shaping the way autonomous networks are being implemented from the perspective of data collection and processing, adoption of AI algorithms, agents, and automation frameworks, resulting in a set of key challenges and opportunities when implementing closed-loop use cases.

As challenges, it is worth mentioning

- Driving insights, complex observability events, and AI decisions requires that a large amount of data be collected, while at the same time data privacy regulations require data collection to be minimized and justified by purpose. Historical data is key in understanding trends and training AI algorithms, which conflicts with the regulations demanding clear limits in data retention policies.
- The goal of autonomous networks is a setup where decisions are taken and actions are implemented automatically by systems with minimal human supervision. This contradicts the most recent regulations (such as the EU AI Act), which impose clear and accountable human oversight, especially when dealing with high-risk use cases and systems.
- Regulators increasingly insist on ensuring the decisions taken by closed-loop systems are explainable and transparent. While this is easier when statically defined policies are defined and is an emerging feature in modern LLM algorithms via the Chain-of-Thought approach, it remains a challenge when "black-box" algorithms such as neural networks, deep learning, and deep reinforcement learning are used, limiting their usage within the autonomous network framework.
- Implementing autonomous network systems significantly expands the complexity and attack surface within the management systems domain. Modern cybersecurity regulations such as NIS2 increase the security management overhead of implementing and operating such systems, with a special focus on model supply chain, data pipelines, control loop monitoring, management, and vendor coordination.

- Most of the modern AI and observability stacks rely on the significant amount of data that is available centrally for training and processing. At the same time, the most advanced algorithms and systems are offered by hyperscalers in a SaaS cloud-first fashion. The current AI and data privacy regulations promote data sovereignty in an increased fashion, which means the data should be processed and used within the jurisdiction where it is collected. At the same time, enterprises that are considered critical for the national infrastructure—for example, service providers—have even stricter restrictions of not sending data outside their own enterprise domain. This either dramatically reduces the data that is exposed to cloud-based AI algorithms and systems or reduces the autonomous network efficiency by forcing AI algorithms and observability systems to be deployed solely on the enterprise premises.

At the same time, regulations also bring a set of benefits for the enterprises implementing autonomous networks:

- Adopting "privacy by design" and following the AI and data privacy regulations in the target autonomous architecture from the start can provide a competitive advantage. Ensuring the data is anonymized, processed at the edge, and purposely collected can reduce the surface attacks and position the corresponding enterprises at the top of the list to win deals with customers who appreciate their privacy by design approach. Following regulatory standards can also become a successful procurement policy because the requirements can be cascaded to the bidding infrastructure vendors to ensure consistency.
- Compliance with the regional regulatory standards allows enterprises to scale faster into different markets and promote reusability because a one-size-fits all data management approach that fits the most restrictive standard can be easily implemented in less restrictive jurisdictions and can be used as a unique selling point.
- By adopting key regulatory requirements that promote portability and interoperability, enterprises can reduce vendor lock-in and associated high costs. Modularization is also promoted for complex elements such as AI algorithms, data collection and processing, automation framework, policy engines, SDN controllers, and observability solutions, making vendor replacement easier and reducing acquisition and operations costs.
- Ensuring full compliance with regulations increases internal and customer trust in closed-loop systems, stimulating autonomous network adoption and usage. Strong data management, policy control, and AI explainability can convince internal and external stakeholders to allow system-based autonomous decision-making and action-taking, which enables an evolution toward higher levels of network autonomy.

Summary and Future Considerations

Data privacy and ethical AI are two of the most important directions in government regulations of the digital economy. Data privacy laws are a result of a long history of individual privacy protections and are therefore more mature across the world following

common nondisclosure and protective guidance. Ethical AI is a novel legislative framework following the recent acceleration in technological development that has different implementations around the world (security over innovation in Europe and Brazil, innovation over security in the rest of the world). Except for a few countries or blocs (e.g., EU AI Act), most of the AI legislation across the world is still being developed and will see major changes as the technology becomes more mature.

Companies enact the legislation by establishing internal guidelines on how to deal with data privacy and ethical AI across multiple markets, with technology companies and service providers leading the way. With AI development in full swing, having a proper policy to deal with it will become a business resilience measure for the entire business landscape.

Apart from proper internal policies, companies are beginning to change their internal governance to ensure data privacy and ethical AI are thoroughly followed across organizations, with multifunctional teams being expected to advise directly to the board of directors soon. Securing data and AI systems is as important as having the right policies and governance. Although data protection is an established practice, AI systems security is an area that will see tremendous development in the future to ensure model, dataset, and behavioral compliance.

AI deployment and development in the next 12–24 months will continue to evolve toward implementing multi-agent and multimodal systems. Agentic AI architectures that execute autonomous tasks by accessing data and triggering actions over a standardized tool framework will become the mainstream technology that companies will use to implement autonomous networks. This will drive the need for regulation-compliant traceability and explainability supported by verbose action logging, which allows for a deterministic replay of decision and actions. Guardrails will continue to evolve to facilitate regulatory or company-mandated limitations to agent and LLM behavior and interactions. Multimodal LLM evolution will require increased input and output content validation and an increased scrutiny on the training data origin and accuracy.

References

- "A Brief History of Data Privacy, and What Lies Ahead," Skyflow: https://www.skyflow.com/post/a-brief-history-of-data-privacy-and-what-lies-ahead
- "History of Privacy Timeline," University of Michigan: https://safecomputing.umich.edu/protect-privacy/history-of-privacy-timeline
- "Overview of the Privacy Act: 2020 Edition," US Department of Justice: https://www.justice.gov/opcl/overview-privacy-act-1974-2020-edition/introduction
- "Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data," European Union: https://eur-lex.europa.eu/legal-content/en/ALL/?uri=CELEX:31995L0046

- "47 U.S. Code § 227—Restrictions on Use of Telephone Equipment," Cornell University: https://www.law.cornell.edu/uscode/text/47/227
- "Summary of the HIPAA Privacy Rule," US Department of Health and Human Services: https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html
- "Lessons Learned from the Family Educational Rights and Privacy Act," Education Data Futures: https://educationdatafutures.5rightsfoundation.com/essays/the-value-of-better-regulation/lessons-learned-from-the-famuly-educational-rights-privacy-act/
- "What Is GDPR?," GDPR.EU: https://gdpr.eu/what-is-gdpr/
- "What is COPPA?," TechTarget Network: https://www.techtarget.com/searchcio/definition/COPPA-Childrens-Online-Privacy-Protection-Act
- "Gramm-Leach-Bliley Act (GLBA): An Overview," UserCentrics: https://usercentrics.com/knowledge-hub/glba-compliance/
- "Patriot Act," Epic: https://epic.org/issues/surveillance-oversight/patriot-act/
- "Fighting Identity Theft with the Red Flags Rule: A How-To Guide for Business," US Federal Trade Commission: https://www.ftc.gov/business-guidance/resources/fighting-identity-theft-red-flags-rule-how-guide-business
- "Right to be Forgotten," Simmons + Simmons: https://www.simmons-simmons.com/en/features/european-data-protection-regulation/ck0zgbxsam5od0b78dm8ri6uk/european-data-protection-regulation-right-to-be-forgotten
- "California Consumer Privacy Act (CCPA)," State of California Department of Justice: https://oag.ca.gov/privacy/ccpa
- "European Declaration on Digital Rights and Principles for the Digital Decade," European Union: https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:52022DC0028
- "The Digital Services Act," European Commission: https://digital-strategy.ec.europa.eu/en/policies/digital-services-act
- "EU Breaks Down Digital Borders: New e-Evidence Rules Facilitate Cross-Border Investigations," White & Case: https://www.whitecase.com/insight-alert/eu-breaks-down-digital-borders-new-e-evidence-rules-facilitate-cross-border
- "The History of the General Data Protection Regulation," European Data Protection Supervisor: https://www.edps.europa.eu/data-protection/data-protection/legislation/history-general-data-protection-regulation_en
- "Global AI Governance Law and Policy: EU," IAPP: https://iapp.org/resources/article/global-ai-governance-eu/
- "AI and the Paperclip Problem," CEPR: https://cepr.org/voxeu/columns/ai-and-paperclip-problem

- "New EU Copyright Rules That Will Benefit Creators, Businesses and Consumers Start to Apply," European Commission: https://www.ecopnet.com/post/new-eu-copyright-rules-that-will-benefit-creators-businesses-and-consumers-start-to-apply
- "The EU Data Act Explained: Rights, Obligations and Challenges for Data Holders and Manufacturers," Diconium: https://diconium.com/en/blog/eu-data-act
- "Data Act," European Commission: https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/european-data-strategy_en
- "NIS2 Directive: Securing Network and Information Systems," European Commission: https://digital-strategy.ec.europa.eu/en/policies/nis2-directive
- "The Evolution of AI Regulation in Asia: A Comparative Analysis," Elastic: https://www.elastic.co/blog/artificial-intelligence-regulation-asia-comparative-analysis
- "Global AI Regulations Tracker: Europe, Americas & Asia-Pacific Overview," Legalnodes:https://legalnodes.com/article/global-ai-regulations-tracker
- "AI Watch: Global Regulatory Tracker—United States," White & Case: https://www.whitecase.com/insight-our-thinking/ai-watch-global-regulatory-tracker-united-states
- "A State-by-State guide to AI Laws in the U.S.," Techtarget Network: https://www.techtarget.com/searchenterpriseai/feature/A-state-by-state-guide-to-AI-laws-in-the-US
- "The White House's 'AI Bill of Rights' Outlines Five Principles to Make Artificial Intelligence Safer, More Transparent and Less Discriminatory," The Conversation: https://theconversation.com/the-white-houses-ai-bill-of-rights-outlines-five-principles-to-make-artificial-intelligence-safer-more-transparent-and-less-discriminatory-192003
- "NIST AI RMF," centraleyes: https://www.centraleyes.com/what-is-nist-ai-rmf/
- "NIST AI RMF Playbook," NIST: https://airc.nist.gov/airmf-resources/playbook/
- "Brazil's AI Act: A New Era of AI Regulation," GDPR Local: https://gdprlocal.com/brazils-ai-act-a-new-era-of-ai-regulation/
- "Korea's New AI Law: Not a Progeny of Brussels," European Center for International Political Economy: https://ecipe.org/insights/koreas-new-ai-law-not-a-progeny-of-brussels/
- *Cisco Live!* Keith Griffin and Harvey Jang: https://www.ciscolive.com/c/dam/r/ciscolive/emea/docs/2024/pdf/ITLGEN-1414.pdf
- "Global Privacy Policy," Cisco: https://www.cisco.com/c/en/us/about/trust-center/global-privacy-policy.html
- "Supplier Code of Conduct," T-Mobile: https://www.telekom.com/resource/blob/498830/7d0f6ce8b7f56c0664c262c3efdf3fe1/dl-240320-supplier-code-of-conduct-data.pdf

- "The House of Digital Responsibility," T-Mobile: https://www.telekom.com/en/company/digital-responsibility/details/house-of-digital-responsibility-1008292
- "Responsible AI," SoftBank Group: https://group.softbank/en/sustainability/mission/materiality/responsibleai
- "SoftBank AI Ethics Policy," SoftBank Group: https://www.softbank.jp/en/corp/aboutus/governance/ai-ethics/
- "SoftBank Corp.: G7 Hiroshima AI Process (HAIP) Transparency Report," OECD.AI: https://transparency.oecd.ai/reports/1216b7b3-eb25-473f-b294-69ba3a6dd910
- "The State of AI in 2025: Agents, Innovation, and Transformation," McKinsey & Company: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
- "Inclusive Artificial Intelligence for Development," UN Trade & Development: https://unctad.org/publication/technology-and-innovation-report-2025
- "Cisco 2026 Data and Privacy Benchmark Study, A Shifting Paradigm: Governance in the Age of AI," Cisco: https://www.cisco.com/c/en/us/about/trust-center/data-privacy-benchmark-study.html#~key-findings
- "Organizational Digital Governance Report 2024," iapp: https://iapp.org/media/pdf/resource_center/organizational_digital_governance_report.pdf
- "FCC Fines US Wireless Carriers over Illegal Location Data Sharing," Reuters: https://www.reuters.com/business/media-telecom/fcc-fines-us-wireless-carriers-nearly-200-million-over-illegal-location-data-2024-04-29/?utm_source=chatgpt.com
- "Japan's Inaugural AI Regulations: A Pro-Innovation Approach," CliffordChance: https://www.cliffordchance.com/insights/resources/blogs/talking-tech/en/articles/2025/03/japans-inaugural-ai-regulations-a-pro-innovation-approach.html
- "What Is the Difference Between Personally Identifiable Information (PII) and Personal Data?," TechGDPR: https://techgdpr.com/blog/difference-between-pii-and-personal-data/
- "Data Protection Digest 4-18 Oct 2025: Transparency the GDPR's 2026 Enforcement Goal, and the Experian Case as a Model NOT to Follow," TechGDPR: https://techgdpr.com/blog/data-protection-digest-20102025-transparency-the-gdprs-2026-enforcement-goal-and-the-experian-case-as-a-model-not-to-follow/?utm_source=chatgpt.com

Index

Numerics

A2A (Agent to Agent) protocol, 261

2G, 403

3GPP (3rd Generation Partnership Program), 29, 383–384

802.11n, 40

A

Aanestad, Margunn, "Augmenting the Algorithm: Emerging Human-in-the-Loop Work Configurations", 502

acceptance test, 320

accuracy, 207–208

Actor-Critic algorithms, 194–195

AdaBoost, 178

advanced performance analysis, 356–357

agentic data poisoning attack, 121–122

AGNTCY Identity, 261

AI, 2. *See also* generative AI; machine learning; models

1970s AI winter, 169

advances in, 221

autonomous networks, 171–172, 223

-based testing, 328–329

challenges and limitations in autonomous networks, 217

lack of data quality and label scarcity, 217–218

lack of interpretability, 218

limited causal reasoning and long-term planning, 219

robustness and failure amplification, 218

scalability and computational constraints, 218–219

company policies, 88

critical systems, 71

data correlation, 9

data privacy, 96–97

deep learning

neural networks, 170

overhead, 170–171

perceptron model, 170

developer usage ratios for AI methods, 466–467

dynamicism, 318

ethical considerations, 214–215

expert systems of the 1980s, 169

fairness, 101
hallucinations, 44, 224
inference, 44
machine learning, 169–170
micro-tasks, 222
neuro-symbolic, generative AI comparison, 223–224
origins, 168
regulation, 70, 80–81, 86–88
Australia, 86
CAIA (Colorado AI Act), 71–72
California AB2013, 72
challenges and opportunities created by, 108–109
China, 84–85
EU AI Act, 70–71
Japan, 85
Latin America, 83–84
South Korea, 85–86
United States, 81–83, 87
Utah AI Policy Act, 72
second winter of the 1980s and early 1990s, 169
testing, 316–317
transparency, 101
value of for organizations, 221–222
AI Act (EU), 79–80
AI agents, 257
alert triage agent, 438
change management, 442–445
frameworks, 257–262
goal-based operations, 255–257
infrastructure
Docker container, 263
framework selection, 262–263
Kubernetes, 263–264
microservices, 262
network automation, 254–255
observability, 421
RCA (root cause analysis) agent, 439
ReACT methodology, 258–261
resolution execution agent, 440
resolution planning agent, 439–440
validation and learning agent, 440–442
AI Ethics Principles, 86
AI governance
AI Access, 423
AI Observability, 423
AI Security, 423
CSP (communication service provider), 422–425
AI RMF (Risk Management Framework), 81–83
AIOps, 44, 355
A/B testing, 159–160
CSP (communication service provider), 425–426
agentic services, 427–428
operational data fabric, 426–427
service observability, 426
GPU activity, 157–159
model comparison, 159–160
security observability, 162–165
task completion rate/quality, 159–160
use cases, 355–358
AireOS wireless controller, 250
algorithms, 168, 185. *See also* machine learning
Actor-Critic, 194–195
clustering, 187–188
DBScan, 189
K-means, 188
linear regression, 174–175
logistic regression, 174–175
model-based, 195–196
policy-based, 193–194
REINFORCE, 194
semi-supervised, 172
value-based, 192–193
anomaly detection, 358

Ansible, 23–25, 327
Apache Flink, 297
APIs, 25–26, 55
network platforms, 249–250
OpenTracing project, 136
architectural planning for network observability, 154–156
architecture
CSP AIOps, 425–428
GameX, 446
mobile network, 403–404
NFV (network function virtualization), 353
SDN (software-defined networking), 353–354
transformer, 171, 224–225
Arize Phoenix, 163–164
artificial neural networks. *See* neural networks
Ask Sage, 278
association rule learning, 190–191
attention mechanisms, 224, 227–228, 229–231. *See also* embeddings
audit trails, autonomous networks, 147–148
Australia, AI regulation, 86
autoencoders, 171–172, 200–201
auto-instrumentation, 153
automation, workflow, 418–419, 426. *See also* network automation
automation tools
Apache Flink, 297
telemetry broker, 291–294
autonomous networks, 411–412. *See also* AI agents; CSP (communication service provider), autonomous networks
AI, capabilities, 223
AI agents, 257
frameworks, 257–262
ReACT methodology, 258–261
AI/ML algorithms, 171–172
audit trails, 147–148
business benefits
Level 3, 450–451
Level 4, 450–452
Levels 0–2, 449–450
challenges and limitations of AI, 217
lack of data quality and label scarcity, 217–218
lack of interpretability, 218
limited causal reasoning and long-term planning, 219
robustness and failure amplification, 218
scalability and computational constraints, 218–219
challenges in achieving autonomy, 39–40
chaos engineering, 331–334
closed-loop operations, 4–5
CSP (communication service provider)
CI/CD pipeline, 418–420
common messaging bus technologies, 416–417
dependencies and obstacles, 452–454
guiding principles, 411–412
Kafka, 416–417
objectives, 408–409
scalable agentic platform, 420–422
security, 421–422
workflow automation and service catalog, 418–419
CSP AIOps, 425–426
agentic services, 427–428
operational data fabric, 426–427
service observability, 426
data management, 34–39
data sources, 5–7
decentralized security, 467
decentralized identity, 468–472
smart contracts, 472–474
Web3, 468

decision-making, 9–11
end-to-end evolution roadmap, 484–486
application hosting and edge computing, 488
Kubernetes, 489
network access control, 486
enterprise domain, 459–460
alignment with core business applications, 489–494
capacity planning, 463
configuration management, 463
fault management, 461
integration with existing systems, 495–499
knowledge management, 461
observability, performance, and assurance, 463
planning and deployment, 494–495
security management, 463
technical challenges, 499–505
uptime, 460
execution process, 11
generative AI, techniques, 222
hosted architectures for inference and training, 43–45
human factor, 50–52
impact of regulations, 103–108
KPIs, 557–560
Level 3+ use-case deployment, 428
capacity planning and forecasting, 433–435
predictive maintenance, 431–432
RCA (root cause analysis), 432–433
SIA (service impact analysis), 429–431
Level 4+ use-case deployment, 435–436
agentic incident resolution framework for CSP networks, 437–438
agentic incident response, alert triage agent, 438
agentic incident response, RCA (root cause analysis) agent, 439
agentic incident response, resolution execution agent, 440
agentic incident response, resolution planning agent, 439–440
agentic incident response, validation and learning agent, 440–442
levels of autonomy, 277–278, 410–411, 461
liability aspects, 51–52
MLOps, 11–12
network automation, 3
network domains, 334–337
NOC/SOC operations, 544–546
observability. *See also* observability
AI operations (GPU activity), 157–159
AI operations (model comparison), 159–160
AI operations (security observability), 162–165
AI operations (task completion rate/quality), 159–160
microscope, 414
network automation with application performance, 156–157
telescope, 413–414, 415–416
OTel (Open Telemetry), 44–45
regulation, challenges and opportunities created by, 108–109
rule-based approach, 12–13
skills required for AN-based operations
data fluency, 563–565
data science fluency, 565–566, 567–569
programming languages, 560–562
vibe coding, 569–571

SSL (semi-supervised learning), 186
systems integration, 56–58
autonomous operations, 30–32
Autonomous Operations Maturity Model, 560
autonomy, network, 1–2

B

A/B testing, 159–160, 294–296
baggage, 142–143, 151–152
bagging, 176–177
bandwidth, federated learning, 265–267
barriers to network autonomy, 39–40
Basic AI Act (South Korea), 85–86
batch processing, 298–299
benchmarking, 139
Bezos, Jeff, 54–55
big data, 170, 550
"black box anxiety", 554–555
Boole, George, 168
boosting, 177–178
AdaBoost, 178
GBMs (gradient boosting machines), 179–180
bootstrap aggregation. *See* bagging
bottlenecks, IP networks, 16–17
Boyd, John, 5, 360
Brazil AI Act, 83–84
Broda, Eric, *Implementing Data Mesh*, 34–35
building containers, 305–308
multi-architecture builds, 310
multistage building, 308–309
tags and labels, 310–312
WebAssembly, 312–313
BUs (business units), 504–505
business continuity, 18
business drivers, network autonomy, 4–8

C

CA (certificate authority) certificates, 127
CAIA (Colorado AI Act), 71–72
California AB2013, 72
California state data breach notification law, 66
canonical logging, 146
capacity planning and forecasting, 433–435, 463
case law, *US vs US District Court 407 US 297*, 60
CCPA (California Consumer Protection Act), 69–70
CD (continuous delivery), 27. *See also* DevOps
CD (continuous deployment), 304–305
centralized execution, network automation, 21
change management, 147
agentic, 442–445
autonomous, 391–394
ITSM (IT service management), 127
chaos engineering, 331–334
chatbots, Ask Sage, 278. *See also* AI agents
Children's Online Privacy Protection Act. *See* COPPA (Children's Online Privacy Protection Act)
China, AI regulation, 84–85
chunking, 242–244
CI (continuous integration), 27, 302–304. *See also* DevOps
CI/CD pipeline, 287–289, 418–420
Cilium Hubble communication flows, containers, 162
Cisco
AGNTCY Identity, 261
global privacy policy, 89–91
project execution, 46–49
responsible AI principles, 90–91

Cisco devices, TACACS+, 124
Cisco Virtual Kubelet, 314–315, 489
Cisco XDR, 55–56, 251
CLAMS framework, 529
closed loops, 345–346, 348–349
autonomous networks, 4–5
challenges and limitations
closed loop stability through decision automation, 396
data quality and real-time observability, 395
integration with legacy and multivendor environments, 396–397
operational trust, governance, and security, 397
policy and intent translation, 395–396
future evolution, 397–398
IBN (intent-based networking), 359
negative feedback, 349
network automation, 3GPP (3rd Generation Partnership Program), 383–384
network operations, 360
business layer, 366
change management, 391–394
data collection and telemetry, 371–373
data processing and analytics, 373–374
domain and cross-domain layer, 367–369
dynamic service assurance and closed-loop optimization, 389–391
end-user layer, 366–367
ETSI ZSM, 377–379
ONAP, 381–383
ONF (Open Networking Foundation), 379
OODA loop, 360–365
orchestration and integration with SDN, 376–377
policy and decision engine, 374–376
predictive network maintenance and self-healing, 386–389
resource operations layer, 365
service layer, 366
TM Forum, 379–381
use cases, 384–386
in networking, 350
positive feedback, 349–350
cloud and hybrid cloud networking, 342
clustering, 187–188. *See also* Kubernetes
DBScan, 189
K-means, 188
silhouette score, 210–211
WCSS (Within-Cluster Sum of Squares), 211–212
CMDB (configuration management database), 5, 127
CNCF (Cloud Native Computing Foundation), 2, 33–34
project proposal process, 285–286
Project Sveltos, 160
CNNs (convolutional neural networks), 171–172, 225
COBIT, 517–519
collaboration tools, as data source, 127–128
"collect once, use many time" philosophy, 416–417
Collier, Mark, 466
confidential computing, 478–479
confidentiality, federated learning, 267–268
configuration management, 463
configurations, standards, 43
confusion matrix, 208
containers
building, 305–308

multi-architecture builds, 310
multistage building, 308–309
tags and labels, 310–312
WebAssembly, 312–313
Cilium Hubble communication flows, 162
context propagation, 150–151
control systems, 346–347. ***See also*** **closed loops**
closed loops, 345–346, 348–349
negative feedback, 349
in networking, 350
positive feedback, 349–350
digital controllers, 346
feedback control, 346
feedback-based regulators, 345
military applications, 346
MIMO (multiple inputs/multiple outputs), 346
open-loop systems, 347–348
PID (proportional-integral-derivative) controllers, 346
SISO (single inputs/single outputs), 346
convergence test, 324
Conway's law, 460–461
COPPA (Children's Online Privacy Protection Act), 64
core dumps, as data source, 128
CoT (Chain-of-Thought) prompting, 240–241
co-training, 185
counters, 138–139. ***See also*** **metrics**
CPO (chief privacy officer), 98
CPU, profiling, 144–145. ***See also*** **GPUs**
CRISP-DM (Cross Industry Standard Process for Data Mining), 35–37
Crispin, Lisa, ***Agile Testing: A Practical Guide for Testers and Agile Teams*****, 318**
CRISP-ML[Q] (Cross Industry Standard Process for Machine Learning and Quality Assurance), 37
critical AI systems, 71
CSP (communication service provider), 116. ***See also*** **mobile data networks; organizations; telecommunications**
access technologies, 405
AI governance, 422–425
autonomous networks
CI/CD pipeline, 418–420
common messaging bus technologies, 416–417
dependencies and obstacles, 452–454
guiding principles, 411–412
Kafka, 416–417
Level 0–2, business benefits, 449–450
Level 3, business benefits, 450–451
Level 4, business benefits, 450–452
levels of autonomy, 410–411
objectives, 408–409
observability, 413–416
scalable agentic platform, 420–422
security, 421–422
workflow automation and service catalog, 418–419
CapEx investment, 406–407
customer expectation, 405–406
customer services, 401–402
E2E (end-to-end) service management, 405
IP transport network, 404–405
mobile data networks
4G, 404
architecture, 403–404
criticality, 14–16
CS (circuit-switched domain), 403–404
PS (packet-switched) domain, 403–404
standards, 29
UTRAN (UMTS Terrestrial Radio Access Network), 403–404
network operations, 509–510

NOC/SOC operations, 540
autonomous networks, 544–546
“dark” NOC, 515–516
integrated operations center, 513–515
NOC engineer, 546–547
service assurance, 512–513
siloed setup, 544
squad-based model, 547
three-tier setup, 541–544
OpEx (operational expenses), 407
reference AIOps architecture, 425–426
agentic services, 427–428
operational data fabric, 426–427
service observability, 426
regulation, 401–402
revenue growth, 406–407
use-case deployment for Level 3+, 428
CSP (content service provider), 29–30. *See also* **SPs (service providers)**
culture change, 553
addressing “black box anxiety”, 554–555
addressing the engineer identity crisis, 555–556
democratizing data, 554
eliminating the “bespoke” culture, 553–554
Cybersecurity Regulation (EU), 78

D

Dapper, 150
“Dapper, a Large-Scale Distributed Systems Tracing Infrastructure”, 135–136
“dark” NOC, 515–516
Dartmouth Conference, 168
data
correlation, 9
transformation, 9
Data Act (EU), 77–78
data bus, 425–426
data centralization, 354–355
Data Governance Act (EU), 75
data mesh, 149–150
data mining, 35–37
data model-based automation, 25–26
data pipeline
network automation, 289–290
security, 480–482
data preprocessing, 202
data cleaning, 203
data splitting, 205–206
data transformation, 203–204
feature engineering, 204
handling imbalanced data, 204–205
data privacy, 59–60
AI adoption and, 96–97
digital governance committee, 99
data processing
batch processing, 298–299
data validation and deduplication, 296
extraction, 290–291
stream processing, 297–298
streaming data, 299–301
telemetry broker, 291–294
data science, 277, 290–291
data sources, 116–120
autonomous networks, 5–7
CA certificates, 127
CMDB (configuration management database), 127
collaboration tools, 127–128
core dumps, 128
DHCP (Dynamic Host Configuration Protocol), 126
DNS (Domain Name System), 126–127
end-user QoE, 134–135
enterprise orchestrators, 128–131
granular data on devices, 128
ITSM (IT service management), 127

MDT (model-driven telemetry), 122, 131–133
NetFlow/IPFIX, 125–126
OTel (Open Telemetry), 133
packet captures, 136
PKI certificates, 127
protocols, 135–136
RADIUS, 125
SNMP (Simple Network Management Protocol), 122
syslog, 120
agentic data poisoning attack, 121–122
formatting of data, 120–121, 122
time series evaluation of exceptions and anomalies, 120
TLS mappings, 122
TACACS+, 124
on Cisco platforms, 124
encryption, 124
troubleshooting, 8
verbosity, 155
data visualization, 9–10
DataOps, 12, 420, 532
DBScan, 189
decentralized security, 467
decentralized identity, 468–472
smart contracts, 472–474
Web3, 468
decision gates, 11
decision trees, 176
decision-making, 9–10
automation, 367–369, 396
expert systems, 169
human-only, 10
machine learning, ethical considerations, 214–215
semi-supervised, 10, 51–52
supervised, 10
declarative-based provisioning, 23
deep learning, 169–170, 196–197, 225, 226. ***See also*** **generative AI; supervised deep learning**
AlexNet, 170
CNNs (convolutional neural networks), 171–172, 225
overhead, 170–171
performance metrics, 212
training loss, 213
validation loss, 212
deep reinforcement learning, 171–172
Dehghani, Zhamak, ***Data Mesh*****, 149**
democratizing data, 554
deployment lifecycle, 287
Deutsche Telekom
"AI guidelines on digital ethics", 92
data privacy officer, 93
"privacy by strategy", 93
devices, granular data, 128
DevOps, 12, 529, 530–531
CI/CD pipeline, 287–289
CLAMS framework, 529
infinity loop, 27
main processes, 529–530
"shift left" philosophy, 530
DGA (Data Governance Act), 75
DHCP (Dynamic Host Configuration Protocol), 126
digital controllers, 346
"Digital Decade", 72
digital governance committee, 99
dimensionality reduction, 189–190
diminishing returns, 18–19
distributed execution, network automation, 21
DMA (Digital Markets Act), 76
DNS (Domain Name System), 126–127
Docker container
AI agent, 263
building, 305–308
Dockerfile, 307–308

domestic intelligence surveillance, due process, 60
DORA (Digital Operational Resilience for the Financial Sector), 78
DPU (data processing unit), 234–235
DQNs (Deep Q-Networks), 201–202
DSA (Digital Services Act), 76
due process, domestic intelligence surveillance, 60
dynamicism, 318

E

edge computing and IoT connectivity, 340
e-Evidence Regulation, 76
E-Government Act, 66
ELT (Extract, Load Transform), 290–291
embeddings, 228
pretrained word vectors, 228–229, 230
RAG (retrieval-augmented generation), 242–243
encryption
syslog, 122
TACACS+, 124
end-user QoE, as data source, 134–135
ensemble models, 176
bagging, 176–177
boosting, 177–178
AdaBoost, 178
GBMs (gradient boosting machines), 179–180
stacking, 180–182
enterprise networks, 32. *See also* organizations
associations, 32–34
automation, 3
autonomous, 459–460
alignment with core business applications, 489–494
capacity planning, 463
configuration management, 463
end-to-end evolution roadmap, 484–489
fault management, 461
high availability, 463
integration with existing systems, 495–499
knowledge management, 461
observability, performance, and assurance, 463
planning and deployment, 494–495
security management, 463
technical challenges, 499–505
uptime, 460
multivendor approach, 28
network operations, 510
NOC/SOC operations
"dark" NOC, 515–516
integrated operations center, 513–515
service assurance, 512–513
observability, 146–147
orchestrators, as data source, 128–131
QoS (quality of service), 459
security, 467
data pipeline, 480–482
decentralized identity, 468–472
federation, 482
model weights, 474–475
smart contracts, 472–474
Web3, 468
shared data center networks, 475–478
software development, 460–461
use cases, 464–465
ethical considerations, AI, 214–215
eTOM, 524–527
ETSI (European Telecommunications Standards Institute), 29
autonomous networks, definition, 4
ZSM (Zero-touch network and Service Management), 377–379

ZSM reference architecture, 116–117
EU (European Union)
AI Act, 70–71, 79–80
content portability and geo-blocking regulation, 75
Convention of Human Rights, 63
Cybersecurity Regulation, 78
Data Act, 77–78
Data Governance Act, 75
Data Protection Directive, 63
Declaration of Digital Rights
FAIR principles, 74
key principles, 73
"Digital Decade", 72
Directive on Copyright in the Digital Single Market, 75
DMA (Digital Markets Act), 76
DORA (Digital Operational Resilience for the Financial Sector), 78
DSA (Digital Services Act), 76
e-Evidence Regulation, 76
GDPR (General Data Protection Regulation), 68, 69
end-user data ownership rights, 68–69
general data protection principles, 68
"Right to be forgotten" regulation, 67
event correlation, 356–357
execution process, 11
expert systems, 169

F

F1-score, 208
FAIR principles, 74
fairness, AI, 101
Family Education Rights and Privacy Act, 61
Fastly, 488
fault management, 341, 461
FCAPS, 527–529
feature engineering, 204
federated learning, 264–265
bandwidth, 265–267
confidentiality and privacy, 267–268
federation, 482
feedback control, 346
feedback-based regulators, 345
FERPA. *See* Family Education Rights and Privacy Act
few-shot/multi-shot prompting, 238
fine-tuning, 246–248
agentic frameworks and tools, 249
embedded function calling, 248–249
versus pretraining, 248
templates, 247
first AI winter, 169
first-order logic, 168
formatting, syslog data, 120–121
foundation models, 233–234
frameworks, 284, 539
agent, 257–258
LangChain, 258–262
selecting, 262–263
AI RMF (Risk Management Framework), 81–83
CLAMS, 529
COBIT, 517–519
DevOps, 529–532. *See also* DevOps
eTOM, 524–527
FCAPS, 527–529
IG1218F Autonomous Networks Framework v1.2.0, 278
ITIL, 522–524
open-source, 284
platform engineering, 532–533
SRE (site reliability engineering), 534–538
unified automation, 550

Frege, Gottlob, 168
functional and nonfunctional testing, 317–318

G

G7 Hiroshima AI Process (HAIP), 95
GameX, 445–446
architecture, 446
flow, 448–449
main elements, 446–448
TM Forum Open APIs, 446
Gartner, 18, 355
gauges, 139–140
GBMs (gradient boosting machines), 179–180
GDPR (General Data Protection Regulation), 68, 69
end-user data ownership rights, 68–69
general data protection principles, 68
generative AI, 221–222. ***See also*** **generative AI**
autonomous networks, techniques, 222
federated learning, 264–265
bandwidth, 265–267
confidentiality and privacy, 267–268
GPUs (graphics processing units), 171
hallucinations, 224
multimodality, 235–236
neuro-symbolic AI comparison, 223–224
prompting, 236–237
Chain-of-Thought, 240–241
few-shot/multi-shot, 238
one-shot, 238
zero-shot, 237–238
RAG (retrieval-augmented generation), 241–243
chunking, 242–244
embeddings, 242–243
graph-based approach, 245–246
limitations, 245
security risks, 244
regulation, 72
security risks, 102–103
tokenization, 238–239
transformer architecture, 171, 224–225
attention mechanisms, 225, 227–228, 229–231
capabilities, 225–226
embeddings, 228–229
GitOps, 301–302
CD (continuous deployment), 304–305
CI (continuous integration), 302–304
"glass cockpit" effect, 556–557
global privacy policy, Cisco, 89–91
GNNs (graph neural networks), 199–200
goal-based automation, 255–257
Google Research
"Attention Is All You Need", 224–225
"Dapper, a Large-Scale Distributed Systems Tracing Infrastructure", 135–136
GPT (Generative Pretrained Transformer), releases, 231–233
GPUs (graphics processing units), 171
inference, 268–270, 274–276
transformer models/architecture, 234
Gramm-Leach-Bliley Act, 64
graph databases, 125–126, 245–246
graph-based training, 185
Gregory, Janet, ***Agile Testing: A Practical Guide for Testers and Agile Teams*****, 318**
Grønsund, Tor, "Augmenting the Algorithm: Emerging Human-in-the-Loop Work Configurations", 502

H

hallucinations, 44, 224
high availability, enterprise networks, 463

HIPAA (Health Insurance Portability and Accountability Act), 63–64
horizontal service observability, 426
hotspots, monitoring, 414–415
HPC (high-performance computing), 43
Huang, Jensen, 48–49
human-only decision-making, 10
Hunt, Andrew, *The Pragmatic Programmer*, 240

I

IaC (Infrastructure as Code), 12
monitoring, 157
tools, 332
IBN (intent-based networking), 359
identity theft, Red Flags Rule, 66–67
IDPs (internal developer platforms), 532–533
IETF (Internet Engineering Task Force), 32, 284
IG1218F Autonomous Networks Framework v1.2.0, 278
imperative-based provisioning, 23
industry frameworks, 27–28
inference, 44
infrastructure
back-end network, 272
compute resources, 272
front-end fabric, 270–271
GPUs, 268–270, 274–276
security, 272–274
storage clusters, 271
infrastructure
agent
Docker container, 263
framework selection, 262–263
Kubernetes, 263–264
microservices, 262
inference
back-end network, 272
compute resources, 272
front-end fabric, 270–271
GPUs, 268–270, 274–276
security, 272–274
storage clusters, 271
telecom, 402–403
insights, 363–365
instrumentation, 141–142
auto-, 153
OTel (Open Telemetry), 153
integrated operations center, 513–515
integration test, 319
intent-based automation, 281
IoT (Internet of Things), feedback control, 346
IP networks, 13. *See also* data sources; mobile data networks
bottlenecks, 16–17
business continuity, 17–18
criticality, 12–16
diminishing returns, 18–19
hotspots, monitoring, 414–415
operational issues, 19–20
technical debt, 16–17
IPFIX, 125–126
ISO standards, 519–522
IT system integration, 53–58
ITIL, 522–524

J-K

Japan, AI regulation, 85
Kafka, 416–417, 425–426
KEIs (key efficiency indicators), 558–559
K-means, 188
knowledge management, 461
KPIs, autonomous networks, 557–560. *See also* metrics
KT (Kepner Traego), 52
Kubeflow, 34

Kubernetes, 43
AI agent, 263–264
Cisco Virtual Kubelet, 314–315, 489
Project Sveltos, 160, 294–296

L

Langfuse, 330
LangSmith, 331
Langwatch.ai, 330
Latin America, AI regulation, 83–84
LECs (local exchange companies), 402–403
legislation, privacy. ***See also*** **EU (European Union); legislation**
California state data breach notification law, 66
CCPA (California Consumer Protection Act), 69–70
COPPA (Children's Online Privacy Protection Act), 64
E-Government Act, 66
Family Education Rights and Privacy Act, 61
Gramm-Leach-Bliley Act, 64
HIPAA (Health Insurance Portability and Accountability Act), 63–64
Privacy Act, 61
USA PATRIOT Act, 65
Level 3+ deployment, 428
capacity planning and forecasting, 433–435
predictive maintenance, 431–432
RCA (root cause analysis), 432–433
SIA (service impact analysis), 429–431
Level 4+ deployment, 435–436
agentic incident resolution framework for CSP networks, 437–438
alert triage agent, 438
RCA (root cause analysis) agent, 439
resolution execution agent, 440
resolution planning agent, 439–440
validation and learning agent, 440–442
levels of autonomy, 410–411
CSP domain
Level 3, 450–451
Level 4, 450–452
Levels 0–2, 449–450
enterprise domain, 461
liability, autonomous networks, 51–52
linear regression, 174–175
Linux Foundation, 2, 33, 284, 466–467
LLMOps, 532
LLMs (large language models), 224, 368. ***See also*** **generative AI**
GPT releases, 231–233
security guardrails, 101–102
security risks, 102–103
logging, 140–141
canonical, 146
disparate formats within a single-vendor ecosystem, 141
logistic regression, 174–175
Lopez de Padro, Marcos, ***Advances in Financial Machine Systems*****, 118**
LSTM (Long Short-Term Memory) networks, 171–172, 198–199, 225, 227

M

machine learning, 168, 169–170. ***See also*** **MLOps (machine learning operations); RL (reinforcement learning); SSL (semi-supervised learning); supervised learning; unsupervised learning**
commercial applications, 170
data preprocessing, 202
data cleaning, 203
data splitting, 205–206
data transformation, 203–204

feature engineering, 204
handling imbalanced data, 204–205
deep learning, 196–197
performance metrics, 212–213
ensemble models, 176
bagging, 176–177
boosting, 177–180
stacking, 180–182
ethical considerations, 214–215
linear regression, 174–175
logistic regression, 174–175
neural networks, 197–198
AlexNet, 170
convolutional, 171–172
deep learning, 169–170
graph, 199–200
perceptron model, 170
training, 198
open source, 170
RL (reinforcement learning), 191–192, 367
Actor-Critic algorithms, 194–195
DQNs, 201–202
model-based algorithms, 195–196
policy-based algorithms, 193–194
value-based algorithms, 192–193
SSL (semi-supervised learning), 184–186
supervised deep learning, 198–199
supervised learning, 173–174, 367
SVC (support vector classifier), 183–184
SVM (support vector machine), 182–183
SVR (support vector regression), 183–184
unsupervised deep learning, 200–201
unsupervised learning, 186–187, 210, 367
association rule learning, 190–191
clustering, 187–189
dimensionality reduction, 189–190
management consulting, 47–49
Maniyka, James, 222
Marick, Brian, 318
maturity evaluation, 286–287
McKinsey and Company, 49–50
MCP (Model Context Protocol), 261
MDT (model-driven telemetry), 5, 122, 131–133
metrics. ***See also*** **performance metrics**
benchmarking, 139
cluster performance
silhouette score, 210–211
WCSS (Within-Cluster Sum of Squares), 211–212
engineering, 49–50
gauges, 139–140
model performance, 207
accuracy, 207–208
confusion matrix, 208
F1-score, 208
recall, 208
ROC (Receiver Operating characteristic) curve, 208–209
observability, 138–140
one-hot encoded information, 138–139
QoE, 134–135
Six Sigma, 148–149
SRE, 536–537
SRO (site reliability objective), 50–51
time series, 139–140
microscope observability, 414
microservices, 262
micro-tasks, 222
MIMO (multiple inputs/multiple outputs), 346
MLOps (machine learning operations), 11–12, 215–217, 420, 531–532
MLPs (multilayer perceptrons), 226

mobile data networks
2G, 403
4G, 404
architecture, 403–404
criticality, 14–16
CS (circuit-switched domain), 403–404
PS (packet-switched) domain, 403–404
standards, 29
UTRAN (UMTS Terrestrial Radio Access Network), 403–404
MOD (Model Openness Framework), 37–39
model-based algorithms, 195–196
model-driven telemetry, 44
ModelOps, 12
models. *See also* LLMs (large language models); supervised learning; transformer models/architecture
ensemble, 176
bagging, 176–177
boosting, 177–180
stacking, 180–182
evaluation, validation, and monitoring, 206–207
fine-tuning, 246–248
agentic frameworks and tools, 249
embedded function calling, 248–249
versus pretraining, 248
templates, 247
foundation, 233–234
graph-based, 171–172
linear regression, 174–175
logistic regression, 174–175
perceptron, 170
performance evaluation, 329–331
performance metrics, 207
accuracy, 207–208
confusion matrix, 208
F1-score, 208
recall, 208
regression metrics, 209–210
ROC (Receiver Operating characteristic) curve, 208–209
temporal, 171–172
training, 173
validation, 173
weights, security, 474–475
monitoring
digital exerience monitoring, 134–135
IaC (Infrastructure as Code), 157
models, 207
network hotspots, 414–415
SNMP (Simple Network Management Protocol), 115–116
syslog, 115–116
multimodality, generative AI, 235–236
multivendor approach, 28, 42–43
Muzzolini, Valerio, 46
MYCIN, 169

N

NDT (network digital twin), agentic change management, 442–445
negative feedback, 349
Neo4J, 125–126
Netflow, 125–126
NetOps
characteristics, 351
siloed operations, 351–352
network automation, 2–3, 281, 368–369
agentic, 254–255
AI model evaluation, 329–331
AI-based testing, 328–329
Ansible, 23–25
APIs, 25–26
CD (continuous delivery), 27
centralized versus distributed execution, 21–24
CI (continuous integration), 27
CI/CD pipeline, 287–289

data model-based, 25–26
data pipeline, 289–290
decision-making, 367–369, 396
declarative-based provisioning, 23
domain vs. cross-domain architecture, 334–337
functional and nonfunctional testing, 317–318
goal-based, 255–257
imperative-based provisioning, 23
network test automation, 315–316
open-source tools, 41
programmatic, 25–26
rule-based, 255
scripted, 21–22, 254
service catalog, 418–419
site reliability engineering, 19–20
telemetry broker, 291–294
Terraform, 23, 25
use cases
cloud and hybrid cloud networking, 342
edge computing and IoT connectivity, 340
fault management and self-healing, 341
network operations and maintenance, 342
NOC/SOC operations, 339–340
performance optimization, 341
predictive maintenance, 337–338
security and compliance, 342
service provisioning and lifecycle management, 340
trend analysis, 338
vendor-specific orchestrators, 41

network autonomy, 1–2, 3. ***See also*** **autonomous networks**
business drivers, 4–8
service providers, 30

network domains, 334–337

network operations, 509. ***See also*** **NOC/SOC operations**
CSP domain, 509–510
enterprise domain, 510
NOC team, 509–510
service provider, 342

network platforms
integration of multimodal input sources, 250–252
limited API interaction, 249–250
object detection, 252

networks, closed loops, 350

neural networks, 197–198, 225. ***See also*** **CNNs (convolutional neural networks); RNNs (recurrent neural networks)**
AlexNet, 170
convolutional, 171–172
graph, 199–200
MLPs (multilayer perceptrons), 226
optimized weights and biases, 226–227
perceptron model, 170
recurrent, 224, 225
training, 198

neuro-symbolic AI, generative AI comparison, 223–224

NFV (network function virtualization), 353

nightlys, 287–288

NIST, AI RMF (Risk Management Framework), 81–83

NLP (natural language processing), 2, 224, 238–239

NOC/SOC operations, 339–340
birth of standards frameworks, 512
CSP domain, 540
autonomous networks, 544–546
siloed setup, 544
squad-based model, 547
three-tier setup, 541–544
“dark” NOC, 515–516
digital shift and start of decentralized compute, 511–512

integrated operations center, 513–515
IT enterprise, 548–553
mainframe/analog era, 511
service assurance, 512–513
nonpersonal data, 101
NRE (network reliability engineering), 550
NVIDIA, 48–49

O

object detection, network platforms, 252
observability, 361, 426. ***See also*** **data sources; OTel (Open Telemetry)**
AI agent, 421
architectural planning, 154–156
audit trails in autonomous networks, 147–148
data mesh, 149–150
data sources, 8
data visualization, 9–10
enterprise domain, 463
enterprise network deployment, 146–147
instrumentation, 141–142, 153
logging, 140–141
metrics, 138–140
microscope, 414
profiling, 143–145
sampling, 147
SP networks, 146–147
spans, 137–138, 150–151
standards, 136
telescope, 413–414, 415–416
traces, 137–138, 150–151
transformation of data, 9
use cases
AI operations (GPU activity), 157–159
AI operations (model comparison), 159–160
AI operations (security observability), 162–165
AI operations (task completion rate/quality), 159–160
network automation with application performance, 156–157
Observability 1.0, 135–136
Observability 2.0, 145
canonical logging, 146
IPv4 OSPF adjacency log message, 145–146
OECD (Organization for Economic Cooperation and Development), privacy recommendations, 61–62
ONAP (Open Network Automation Platform), 381–383
one-hot encoded information, **138–139**
one-shot prompting, 238
ONF (Open Networking Foundation), 379
OODA loop, 5–8, 360–365. ***See also*** **decision-making; observability**
OpenCensus, 136
OpenConfig Project, 42–43
open-loop systems, 347–348
open-source frameworks, 284
open-source tools, network automation, 41
OpenTofu, 327
OpenTracing project, 136
OpEx (operational expenses), 4, 407
orchestrators
categorization of network relevant data types, 129–131
as data source, 128–131
differences with system-direct derived telemetry, 128–129
vendor-specific, 41
organizations. ***See also*** **enterprise networks**
communication structures, 48–49
culture change, 553

addressing "black box anxiety", 554–555
addressing the engineer identity crisis, 555–556
democratizing data, 554
eliminating the "bespoke" culture, 553–554
digital governance committee, 99
management consulting, 47–49
people and processes, 45–49
systems integration, 53–58
value of AI, 221–222
OSDL (Open Source Development Labs), 33
OTel (Open Telemetry), 33, 44–45, 136
baggage, 142–143, 151–152
as data source, 133
instrumentation, 153
instrumentation libraries, 141–142
profiling, 143–145
overfitting, 176

P

packet captures, 136
PCA (principal component analysis), 189–190
perceptron model, 170
performance metrics
clustering
silhouette score, 210–211
WCSS (Within-Cluster Sum of Squares), 211–212
deep learning, 212
training loss, 213
validation loss, 212
model, 207
accuracy, 207–208
confusion matrix, 208
F1-score, 208
recall, 208
regression metrics, 209–210
ROC (Receiver Operating characteristic) curve, 208–209
performance optimization
enterprise domain, 463
service provider, 341
Perl, 21–22
Perrin, Jean-Georges, *Implementing Data Mesh*, 34–35
personal data, 100–101
PHI (protected health information), 64
PID (proportional-integral-derivative) controllers, 346
PII (personally identifiable information), 61, 100–101
Pinto, Ivo, *Network Automation Made Easy*, 281
Pinto, Joe, 52
PKI (public key infrastructure) certificates, 127
platform engineering, 532–533
Playwrite, 327
Podman, 306
policies, privacy. *See also* legislation; global privacy policy, Cisco; privacy
Cisco, 89–91
Deutsche Telekom, 91–94
Softbank, 94–96
policy-based algorithms, 193–194
positive feedback, 349–350
PPO (proximal policy optimization), 194–195
predictive analytics, 355–356
predictive maintenance, 337–338, 431–432
pre-production lab, 420
pre-standard adoption, 40
pretrained word vectors, 228–229, 230
privacy, 59–60
Cisco global privacy policy, 89–91
Deutsche Telekom policy, 93
EU Data Protection Directive, 63

EU GDPR, 68, 69
end-user data ownership rights, 68–69
general data protection principles, 68
federated learning, 267–268
identity theft, Red Flags Rule, 66–67
legislation
California state data breach notification law, 66
CCPA (California Consumer Protection Act), 69–70
COPPA (Children's Online Privacy Protection Act), 64
E-Government Act, 66
Family Education Rights and Privacy Act, 61
Gramm-Leach-Bliley Act, 64
HIPAA (Health Insurance Portability and Accountability Act), 63–64
Privacy Act, 61
TCPA (Telephone Consumer Protection Act), 62
USA PATRIOT Act, 65
OECD recommendations, 61–62
regulation, challenges and opportunities created by, 108–109
"right to", 60, 63
right to be forgotten, 67
Softbank policy, 94–96
Privacy Act, 61
Privacy Rule, HIPAA, 63
profiling, 143–145
programmatic automation, 25–26
Project Sveltos, 160, 294–296
prompting, 236–237
Chain-of-Thought, 240–241
few-shot/multi-shot, 238
one-shot, 238
zero-shot, 237–238
protocols, as data source, 135–136
PyATS, 325–326
Python, 22

Q

Q-learning, 192–193
QoS (quality of service), 459
quantum technology, 2

R

RADIUS, 125
RAG (retrieval-augmented generation), 241–243
chunking, 242–244
embeddings, 242–243
graph-based approach, 245–246
limitations, 245
security risks, 244
random forest, 176–177
RCA (root cause analysis), 432–433
ReACT methology, 258–261
reasoning, 168, 241
recall, 208
Red Flags Rule, 66–67
Redhat Ansible. ***See*** **Ansible**
regression metrics, model performance, 209–210
regulations. ***See also*** **policies**
AI, 70, 80–81, 86–88
Australia, 86
CAIA (Colorado AI Act), 71–72
California AB2013, 72
China, 84–85
EU AI Act, 70–71, 79–80
Japan, 85
Latin America, 83–84
South Korea, 85–86
United States, 81–83, 87
Utah AI Policy Act, 72

autonomous network impacts, 103–108
challenges and opportunities created by, 108–109
content portability and geo-blocking, 75
CSP (communication service provider), 401–402
REINFORCE, 194
responsible AI principles, Cisco, 90–91
RFC 5425, 122
RFPs (requests for proposals), 40–41
right to be forgotten, 67
right to privacy, 60, 63
RL (reinforcement learning), 171, 191–192, 367
Actor-Critic algorithms, 194–195
DQNs (Deep Q-Networks), 201–202
model-based algorithms, 195–196
policy-based algorithms, 193–194
value-based algorithms, 192–193
RNNs (recurrent neural networks), 224, 226
LSTM (Long Short-Term Memory) networks, 171–172, 198–199, 225, 227
sequential processing bottleneck, 225
Robot Framework, 325–326, 327
ROC (Receiver Operating characteristic) curve, 208–209
rubberducking, 240
rule-based approach, autonomous networks, 12–13
Russell, Stuart, *Human Compatible*, 221

S

sampling, 147
scale/property test, 320
scripts, 12–13
limitations, 254
network automation, 21–22
Python, 22
SDN (software-defined networking), 353–354
second AI winter, 169
security
confidential computing, 478–479
CSP autonomous network, 421–422
data pipelines, 480–482
decentralized, 467
decentralized identity, 468–472
smart contracts, 472–474
Web3, 468
enterprise domain, 463
federation, 482
inference infrastructure, 272–274
model weights, 474–475
observability, 162–165
RAG (retrieval-augmented generation), 244
service provider, 342
shared data center networks, 475–478
SPIFFE runtime environment, 482–484
Selenium, 327
self-healing, 341
self-training, 185
semi-supervised decision-making, 10, 51–52, 172
service assurance, 512–513
service catalog, 418–419
service providers
domains, 30
failure to achieve network autonomy, 39
multivendor deployment, 42–43
network autonomy, 30
systems integration, 53–58
service provisioning and lifecycle management, 340
service test, 320–321
shared data center networks, 475–478
“shift left” philosophy, 530
SIA (service impact analysis), 429–431

silhouette score, 210–211
single-vendor ecosystem, disparate logging formats, 141
SISO (single inputs/single outputs), 346
Site Reliability Engineering: How Google Runs Production Systems, **20, 49**
Six Sigma, 148–149
skills required for AN-based operations
data fluency, 563–565
data science fluency, 565–566
new skills adoption, 570–571
programming languages, 560–562
prompt and context engineering, 567–569
vibe coding, 569–571
SLA (service-level agreement), 125, 490–491
SLI (service-level indicator), 490–491
SLO (service-level objective), 490–491
Slurm, 43
smart contracts, 472–474
SNMP (Simple Network Management Protocol), 115–116, 122
SOA (service-oriented architectures), 55
Softbank
AI Ethics Policy, 95
AI governance, 94–95
AI policy, 95–96
software development. ***See also*** **DevOps**
Conway's law, 460–461
enterprise networks, 460–461
source of truth, 289–290, 291
South Korea, AI regulation, 85–86
sovereign infrastructure, 28, 104
spans, 137–138, 150–151
SPIFFE runtime environment, 482–484
Splunk ML, 159–160
SPs (service providers). ***See also*** **CSP (communication service provider)**
cloud and hybrid cloud networking, 342
edge computing and IoT connectivity, 340
fault management and self-healing, 341
Level 4 autonomy, 277–278
network automation, 2
network operations and maintenance, 342
observability, 146–147
performance optimization, 341
security and compliance, 342
service provisioning and lifecycle management, 340
SQM (service quality management), 512–513
SRE (site reliability engineering), 19–20, 126, 534–538
metrics, 536–537
SRO (site reliability objective), 50–51
SRO (site reliability objective), 50–51
SSL (semi-supervised learning), 184–186
stacking, 180–182
standards, 512
configurations, 43
IETF, 284
ISO, 519–522
mobile network, 29
observability, 136
pre-standard adoption, 40
syslog data format, 120–121
telemetry, 44
standards bodies, 284
startup organizations, 46
Steenwinkel, Franck van, 52
stream processing, 297–298, 299–301
supervised decision-making, 10
supervised deep learning, 198–199
supervised learning, 173–174, 367
SVC (support vector classifier), 183–184
SVM (support vector machine), 182–183
SVR (support vector regression), 183–184
SVC (support vector classifier), 183–184

SVM (support vector machine), 182–183
SVR (support vector regression), 183–184
syslog, 115–116, 120
agentic data poisoning attack, 121–122
encryption, 122
formatting of data, 120–121, 122
time series evaluation of exceptions and anomalies, 120
TLS mappings, 122
system test, 319
systems integration, 53–58

T

TAC (Technical Assistance Center), 17, 52
TACACS+, 124, 162
on Cisco platforms, 124
encryption, 124
tag switching, 40
TCL interpreted language, 21
TCPA (Telephone Consumer Protection Act), 62
technical debt, 16–17
technology
"Digital Decade", 72
pre-standard adoption, 40
privacy concerns, 60
telecommunications. ***See also*** **mobile data networks**
automatic switching systems, 402–403
early networks, 402–403
inter-city connectivity, 402–403
IP transport network, 404–405
LECs (local exchange companies), 402–403
mobile communications, 403
telemetry. ***See also*** **data sources; monitoring; observability; OTel (Open Telemetry)**
instrumentation, 141–142
model-driven, 44
standards, 44
telemetry broker, 291–294
Telephone Consumer Protection Act, 62
telescope observability, 413–416
templates, fine-tuning, 247
temporal models, 171–172
Terraform, 23, 25, 327
test tools, 324–327
testing, 316
acceptance test, 320
AI systems, 316–317
AI-based, 328–329
A/B, 159–160, 294–296
convergence test, 324
extended testing considerations, 321–324
functional and nonfunctional, 317–318
integration test, 319
metrics, 325
scale/property test, 320
service test, 320–321
system test, 319
unit test, 319
Thomas, Dave, ***The Pragmatic Programmer*****, 240**
time series metrics, 139–140
TM Forum, 2, 29, 222, 282, 284, 410
2024 survey, 4
approach toward maturity evaluation, 286–287
autonomous network levels, 30–31
autonomous operations levels, 31
closed-loop automation, 379–381
GameX, 446
Outlook for Level 4 Autonomy in Service Providers, 277–278
tokenization, 238–239
topology-aware service visualization, 358
TraceID, 151
traces
context propagation, 150–151
observability, 137–138

training, 173
co-, 185
ethical considerations, 214–215
graph-based, 185
neural networks, 198
reskilling vs. replacing, 555–556
self-, 185
training loss, 213
transformation of data, 9
transformer models/architecture, 171, 224–225, 233–234
attention mechanisms, 225, 227–228, 229–231
capabilities, 225–226
embeddings, 228–229
GPUs, 234
transparency, AI, 101
"Tree of Thoughts: Deliberate Problem Solving with Large Language Models", 241
Tree-of-Thought reasoning, 241
trend analysis, 338
troubleshooting, 52
data sources, 8
KT (Kepner Traego), 52
Turing, Alan, 168
Turing Test, 168

U

UDHR (Universal Declaration of Human Rights), 60
UDP (User Datagram Protocol), 115–116, 121
unified automation framework, 550
unit test, 319
United States
AI regulation, 81–83, 87
COPPA (Children's Online Privacy Protection Act), 64
Family Education Rights and Privacy Act, 61
Gramm-Leach-Bliley Act, 64
HIPAA (Health Insurance Portability and Accountability Act), 63–64
Privacy Act, 61
TCPA (Telephone Consumer Protection Act), 62
unsupervised deep learning, 200–201
unsupervised learning, 186–187, 210, 367
association rule learning, 190–191
clustering, 187–188
DBScan, 189
K-means, 188
silhouette score, 210–211
WCSS (Within-Cluster Sum of Squares), 211–212
dimensionality reduction, 189–190
uptime, enterprise networks, 460
US Department of Defense, Ask Sage chatbot, 278
USA PATRIOT Act, 65
use cases
AIOps, 355–358
autonomous network operations, 384–386
change management, 391–394
dynamic service assurance and closed-loop optimization, 389–391
predictive network maintenance and self-healing, 386–389
industrial and manufacturing sector networks, 464–465
Level 3+ deployment, 428
capacity planning and forecasting, 433–435
predictive maintenance, 431–432
RCA (root cause analysis), 432–433
SIA (service impact analysis), 429–431
Level 4+ deployment, 435–436

agentic incident resolution framework for CSP networks, 437–438
agentic incident response, alert triage agent, 438
agentic incident response, RCA (root cause analysis) agent, 439
agentic incident response, resolution execution agent, 440
agentic incident response, resolution planning agent, 439–440
agentic incident response, validation and learning agent, 440–442
network automation
cloud and hybrid cloud networking, 342
edge computing and IoT connectivity, 340
fault management and self-healing, 341
network operations and maintenance, 342
NOC/SOC operations, 339–340
performance optimization, 341
predictive maintenance, 337–338
security and compliance, 342
service provisioning and lifecycle management, 340
trend analysis, 338
observability
AI operations (GPU activity), 157–159
AI operations (model comparison), 159–160
AI operations (security observability), 162–165
AI operations (task completion rate/quality), 159–160
network automation with application performance, 157–159
Utah AI Policy Act, 72
UTRAN (UMTS Terrestrial Radio Access Network), 403–404

V

validation, model, 173, 206–207
validation loss, 212
value-based algorithms, 192–193
vendor-specific orchestrators, 41
verbosity, data sources, 155
vibe coding, 569–571
VMs (virtual machines), 16–17

W

Watergate scandal, 61
WCSS (Within-Cluster Sum of Squares), 211–212
Web3, 468
WebAssembly, 312–313
White House Blueprint for an AI Bill of Rights, 81
Wi-Fi alliance, 32
workflow automation, 418–419, 426

X-Y

XCON, 169Yang models, 131
YANG models, 153

Z

Zero Trust, 11, 46–47
zero-shot prompting, 237–238
Zymtrace, 158