

DATA GOVERNANCE AND ANALYTICS

A Philippine-Context Guide to Data Collection,
Governance, Security and Analytics



Cristeta G. Tolentino, DIT

DATA MANAGEMENT

*A Philippine-Context Guide to Data Collection, Governance,
Security and Analytics*

Cristeta G. Tolentino, DIT

Copyright Page

All rights reserved.

No part of this publication may be reproduced, distributed, stored in a retrieval system, or transmitted in any form or by any means, whether electronic, mechanical, photocopying, recording, scanning, or otherwise, without the prior written permission of the author and publisher, except for brief quotations used in reviews, scholarly works, and other uses permitted by applicable copyright laws.

DATA MANAGEMENT: A PHILIPPINE-CONTEXT GUIDE TO DATA COLLECTION, GOVERNANCE, SECURITY, AND ANALYTICS

Published by: Philippine Technical Institute Inc.
Guilig Street, Poblacion, Lingayen, Pangasinan 2401, Philippines

Mobile No.: +63 918 620 0902

Email: po@pti.edu.ph

Website: press.pti.edu.ph

Date of Publication: June 30, 2026

ISBN:

PDF Downloadable Edition: 978-621-8546-06-6

PDF Read-Only Edition: 978-621-8546-05-9

Copyright © 2026 by:

Cristeta G. Tolentino

Cover Design: PTI Press Creative Team

The views and opinions expressed in this publication are those of the author and do not necessarily reflect the views of the publisher or its affiliated institutions.

Published in the Philippines by PTI Press.

Foreword

Graduate education in data-intensive fields faces a recurring, quiet gap: students learn database theory, statistical methods, and security concepts as separate subjects, but rarely as one connected discipline that practitioners must hold together in real work. A graduate student may understand relational schema design yet still be unprepared when a real dataset arrives dirty, undocumented, and legally sensitive. *Data Management: A Philippine-Context Guide to Data Collection, Governance, Security, and Analytics* closes that gap with rare coherence for a single-author graduate text.

What distinguishes this book is its architecture. Its fourteen chapters, organized into six deliberate Parts, guide students through the disciplined sequence a data professional follows in practice: governance, collection and storage, quality assurance, processing, security and compliance, and analysis and communication. The cumulative Data Project Build threaded throughout the chapters is, in my view, the book's strongest pedagogical decision. It does not allow data management to remain abstract; it requires students to practice it, step by step, using a dataset of their own choosing.

I have had the privilege of collaborating with the author, Cristeta G. Tolentino, across several research and publishing efforts in information technology and data management education. In this book, I recognize the same qualities I associate with her scholarship: evidence-based frameworks, strong attention to Philippine regulatory and organizational realities, and genuine concern for what graduate students can actually do with what they have learned. This book does not treat Filipino data professionals as an audience for generic imported theory; it places them at the center.

I recommend this book without reservation to instructors seeking a Philippine-context foundation for data management courses and to graduate students who want more than vocabulary, who want to finish the course with a complete, defensible data management portfolio of their own.

Randy Joy Ventayen, DBA, DIT

Elmar B. Noche, MIT

Preface

This book grew out of a simple observation: the existing MDMN 202 Learning Module gave graduate students the vocabulary of data management, but not always the connective structure that lets one module's concepts build deliberately on the last. Governance means little without collected data to govern. Quality assurance means little without an understanding of where data actually comes from. Security means little without an appreciation of what makes data worth protecting in the first place. This book was written to make that connective structure explicit, chapter by chapter, so that by the final page, a student has not merely studied fourteen separate topics, but walked a single, coherent data management process from first collection to final, communicated insight.

Every chapter follows a consistent template deliberately: learning outcomes, historical grounding, key definitions, core theory, a worked table, a Philippine Context feature, a case study, best practices and emerging trends, a chapter summary, key terms, reflection and discussion questions, a Data Project Build activity step, a self-check quiz, and full references. This consistency is intentional; it allows students to know exactly what to expect from each chapter and instructors to adapt, reorder, or supplement individual sections without disrupting the book's overall architecture.

This edition has been formatted to the Philippine Technical Institute Inc. (PTI Press) house style. Two items remain outstanding before final publication: a full reference-accuracy audit against APA 7th edition standards, and confirmation against the final approved MDMN 202 syllabus of any scope adjustments the teaching faculty may require. Both are noted transparently here rather than presented as complete.

Cristeta G. Tolentino, DIT

How to Use This Book

For Students: Read each chapter in sequence, since later chapters deliberately assume familiarity with frameworks introduced earlier. Complete each chapter's Data Project Build step as you go, applying every new framework to the same evolving dataset, rather than waiting until Chapter 14 to begin. Use the Chapter Self-Check quiz to confirm understanding of key terms and concepts before moving to the next chapter, and treat the reflection and discussion questions as genuine thinking prompts, not merely questions to skim past.

For Instructors: Each Part corresponds to a natural instructional unit (Foundations, Collection and Storage, Quality and Integrity, Processing and Transformation, Security and Privacy, Analytics and Visualization), and Parts may be paced across weeks according to your course calendar, consistent with the MDMN 202 syllabus's six-module structure. The cumulative Data Project Build activity serves as the course's primary summative assessment, culminating in the Chapter 14 capstone presentation. The accompanying Appendix A worksheet packet may be distributed at the start of the term as a running companion document.

Icons and Features: Throughout this book, colored feature boxes signal specific content types, Key Terms (definitions), Philippine Context (local application), Case Study (illustrative composite scenarios, not depicting real named individuals or organizations unless explicitly cited), and Data Project Build (your cumulative assignment). Part-opening pages use a consistent color per Part, carried through into that Part's chapter headings, tables, and feature-box accents, to help orient students to which Part of the data management journey they are currently in.

A Note on Depth: Every chapter deliberately combines conceptual grounding with practical application; the Extended Application and Frequently Asked Questions subsections found in several chapters are not optional supplementary material, but a continuation of the chapter's core teaching, often addressing the specific misapplications and edge cases students most commonly encounter when first applying a given framework to their own dataset. Instructors assigning partial chapter readings should be aware that skipping these subsections omits genuine content, not mere elaboration.

Acknowledgments

This book builds directly on the MDMN 202 Data Management course syllabus and its accompanying Modules in Data Management learning module (Pangasinan State University, Graduate School), whose six modules and seven course outcomes provided the foundational structure this fourteen-chapter expansion restructures and deepens. The book's development process mapped each chapter's scope back to a defined, approved course outcome and module topic rather than to ad hoc topic selection.

The author gratefully acknowledges the colleagues, co-authors, and academic community whose ongoing collaboration in Philippine information technology, data management, and business intelligence education has shaped the perspective this book brings to a Philippine-context treatment of an inherently global subject.

Table of Contents

Book Cover	i
Title Page	ii
Copyright Page	iii
Foreword	iv
Preface	v
How to Use This Book	vi
Acknowledgments	vii
Table of Contents	viii
Part I. Foundations of Data Management	1
Chapter 1. Understanding Data Management	2
Chapter 2. Key Components and Data Governance	12
Chapter 3. The Data Lifecycle	22
Part II. Data Collection and Storage	31
Chapter 4. Data Sources, Types, and Collection Methods	32
Chapter 5. Data Storage Systems and Architectures	41
Part III. Data Quality and Integrity	50
Chapter 6. Dimensions and Challenges of Data Quality	51
Chapter 7. Data Cleansing and Validation Techniques	62
Part IV. Data Processing and Transformation	72
Chapter 8. ETL Processes and Data Transformation Techniques	73
Chapter 9. Data Processing Tools and Automation	82
Part V. Data Security and Privacy	91
Chapter 10. Data Security Threats, Encryption, and Access Control	92
Chapter 11. Data Privacy, Compliance, and Regulatory Standards	102
Part VI. Data Analytics and Communication	112
Chapter 12. Data Analytics Fundamentals and Types	113
Chapter 13. Data Visualization and Communicating Insights	123
Chapter 14. Emerging Trends and Technologies in Data Management	132
List of Abbreviations	142
List of Figures	143
List of Tables	144
Glossary	145
Appendix A: The Complete Data Project Build Worksheet Packet	147
Appendix B: Philippine Legal and Regulatory Quick Reference	148

Appendix C: Index of Key Frameworks and Course Outcome Alignment	149
Appendix D: Suggested Course Calendar	151
Suggested Further Reading	152
Notes on Sources and Academic Integrity	153
About the Author	154

PART I

FOUNDATIONS OF DATA MANAGEMENT

Part I lays the conceptual groundwork for everything that follows. Before an organization can collect, process, secure, or analyze its data, it must first understand what data management actually is, what components a mature program requires, and how data moves through its complete lifecycle from creation to disposal.

UNDERSTANDING DATA MANAGEMENT

Foundations of a Discipline Every Organization Now Depends On

“An organization does not lack data. It lacks a disciplined way of turning that data into something it can trust.”

1.0 Chapter Overview

This chapter opens your study of data management by establishing what the discipline actually is, why it emerged, and why organizations that once treated data informally can no longer afford to do so. Before you can collect data responsibly, secure it, or extract insight from it, you need a clear and accurate picture of the field as a whole and the vocabulary that describes it. This chapter also introduces the Data Project Build, the cumulative project you will carry through all fourteen chapters of this book, culminating in a complete data governance and analytics portfolio built around a real or realistic dataset of your choosing.

1.1 Learning Outcomes

- LO 1.1 Define data management and explain its importance in modern organizations.
- LO 1.2 Distinguish data management from adjacent disciplines such as data science and information technology.
- LO 1.3 Identify the organizational pressures that have made formal data management a necessity rather than a luxury.

1.2 Introduction

Every organization, from a two-person sari-sari store to a national government agency, generates data constantly: transactions, records, communications, sensor readings, forms. What separates an organization that benefits from this data from one that is quietly buried under it is not the volume of data it holds, but whether it manages that data as a deliberate discipline. Data

management is that discipline: the practices, roles, and technologies an organization uses to ensure its data is collected properly, stored securely, kept accurate, governed responsibly, and made usable for decisions.

This distinction matters because data, unlike most organizational assets, does not announce its own decay. A broken machine stops working visibly. Data that has quietly become inconsistent, duplicated, or outdated continues to look like data, and continues to be used, often for years, before its unreliability surfaces in a bad decision, a compliance failure, or a security breach. The chapters that follow teach you to prevent that decay through disciplined practice rather than discover it after the fact.

It is worth being explicit about how this book is organized, since the sequence itself teaches something. Part I, this Part, deliberately establishes vocabulary and governance before any technical practice begins. Parts II through V then follow the data lifecycle in the order in which data is actually processed: collected, stored, quality-checked, processed, and secured. Part VI closes with analytics and visualization, the stage every earlier chapter has quietly been building toward. A reader who skips ahead to a later Part without this foundation will still learn the technical content, but will miss the governance and quality discipline that determines whether that technical content can actually be trusted in practice.

1.3 Historical Background

Formal data management traces its institutional roots to the rise of database management systems in the 1970s, when organizations first needed dedicated software to reliably store and retrieve structured records, replacing ad hoc paper filing and disconnected spreadsheets. Edgar F. Codd's 1970 paper on the relational model gave the field its first rigorous theoretical foundation, and the decades that followed saw data management mature from a purely technical concern, keeping a database running, into an organizational discipline spanning governance, quality, security, and ethics.

The shift accelerated sharply from the 2000s onward for two converging reasons. First, the sheer volume and variety of data organizations generated grew far beyond what traditional relational databases were designed to handle, as unstructured data such as documents, images, and social media content became routine organizational assets alongside structured transactional records. Second, a series of high-profile data breaches and the resulting public and regulatory backlash, culminating in laws such as the European Union's General Data

Protection Regulation and the Philippines' Data Privacy Act of 2012, made data management a matter of legal compliance and organizational risk, rather than merely operational convenience.

Codd's original relational model is worth briefly dwelling on, since its core insight still shapes how this book approaches data management, even outside relational database contexts. Codd argued that data should be organized around its logical structure, independent of how it happened to be physically stored, a separation that let applications and analyses be built without needing to know the underlying storage details. That same principle, that data's meaning and organization should be managed deliberately and separately from whatever specific technology currently stores it, is precisely why this book treats data management as a discipline that outlasts any single database technology or storage platform, a theme Chapter 14 returns to directly.

1.4 Definitions

KEY TERMS

Data Management: The set of practices, roles, and technologies an organization uses to collect, store, secure, govern, and derive value from its data throughout its useful life.

Data, Raw, unprocessed facts and figures, such as a single transaction record or sensor reading, that have not yet been organized into a meaningful context.

Information, Data that has been organized, processed, or structured in a way that gives it meaning and usefulness to the person or system consuming it.

Data Governance: The framework of policies, roles, and decision rights that determines how an organization's data is managed, who is accountable for it, and how quality and compliance are enforced.

Data Lifecycle: The stages data passes through from creation or collection to eventual archival or deletion, developed fully in Chapter 3.

1.5 Core Concepts and Frameworks




Figure 1.1. From Data to Insight.

1.5.1 Data Management as a Discipline, Not a Department

A common early misconception is that data management is simply the job of the IT department, just as maintaining email servers is. In practice, effective data management is cross-functional: legal and compliance teams define what regulations require, business units define what quality and accessibility they need to do their jobs, and only the technical implementation, the databases, pipelines, and security controls, belong primarily to IT. The chapters that follow deliberately reflect this breadth, covering not only technical storage and processing but also governance, quality, security, and ethical use.

1.5.2 Data Management, Data Science, and Information Technology

These three fields are frequently conflated but serve distinct purposes. Information technology is concerned broadly with the hardware and software infrastructure an organization runs on, of which data storage is only one part. Data science is concerned with extracting predictive or exploratory insight from data, typically through statistical and machine learning methods, and depends on the data it analyzes already being well managed. Data management sits upstream of both: it is the discipline that ensures data is trustworthy, accessible, and compliant before it ever reaches a data scientist's model or an executive's dashboard.

1.5.3 Why Formal Data Management Is Now a Necessity

Three organizational pressures have converged to make informal, ad hoc data practices untenable. Regulatory pressure requires organizations to demonstrate specific data-handling practices, particularly for personal data, under laws such as the Philippines' Data Privacy Act. Competitive pressure rewards organizations that can extract reliable insight from their data faster than rivals, which is impossible without well-managed underlying data. Operational pressure comes from the sheer scale of modern data generation, which overwhelms manual, undocumented practices that may have worked adequately at a smaller scale a decade earlier.

1.6 Table 1.1, Data vs. Information vs. Insight

Stage	Description	Example
Data	Raw, unprocessed facts	“450” recorded in a sales log
Information	Data organized with context	“450 units of Product A sold in March”
Insight	Information interpreted to support a decision	“Product A sales rose 20% after the March promotion; extend the promotion”

Table 1.1. Moving from raw data to actionable insight.

1.7 Government Policy and Professional Standards

In the Philippines, the Data Privacy Act of 2012 (Republic Act No. 10173) establishes the foundational legal framework for the management and administration of personal data, administered by the National Privacy Commission (NPC). Organizations that collect, process, or store personal data, which in practice means nearly every organization of any size, are legally required to implement specific data management safeguards, as this book develops in later chapters, particularly Chapters 6, 10, and 11. Professional data management practice is further shaped by international frameworks such as DAMA International's Data Management Body of Knowledge (DMBOK), which this book draws on for its treatment of governance and quality.

1.8 Philippine Context and Case Study

PHILIPPINE CONTEXT

Philippine government agencies increasingly maintain open data portals, publishing datasets on health, education, and economic indicators for public and research use. This practice, alongside private-sector adoption of formal data management driven by Data Privacy Act compliance, has accelerated demand for graduates who understand data management as a full discipline rather than a narrow technical skill, precisely the demand this book and the MDM program are designed to meet.

This demand extends well beyond the largest Philippine enterprises and government agencies. Small and medium enterprises undergoing digital transformation, cooperatives modernizing member records, and local government units building citizen-facing digital services all increasingly require the same disciplined foundation this chapter introduces, even when they cannot yet justify a dedicated, full-time data management department. Graduates who understand the discipline well enough to apply it proportionately, at whatever scale an organization can support, will find the demand for this competency considerably broader than the largest, most visible employers alone would suggest.

CASE STUDY

Illustrative composite case.

A mid-sized Philippine cooperative bank had, for years, maintained member records across a mix of spreadsheets, a legacy loan system, and paper files, with no single source of truth for a given member's complete financial relationship with the bank. When the bank sought to launch a new digital lending product, staff discovered that basic questions, such as how many active members held both savings and loan accounts, could not be answered without days of manual reconciliation. The bank's subsequent investment in a formal data management function, a designated data steward, a consolidated core system, and documented governance policies was justified not by abstract best practice but by this very concrete, costly gap.

The reconciliation delay proved costly beyond the immediate inconvenience: the digital lending product's launch was pushed back nearly two months while staff manually cross-referenced records across all three systems, a delay that allowed a competing cooperative to launch a similar product first. Board members who had previously viewed a dedicated data management function as an unnecessary overhead cost reversed that view once the connection between the delay and the absence of that function became concrete and quantifiable, illustrating a pattern common across the case studies in this book: data management investment is easiest to justify in hindsight, after its absence has already cost something measurable.

1.9 Best Practices, Current Issues, and Emerging Trends

- **Best Practice:** Treat data management as a cross-functional responsibility from the outset, involving legal, business, and technical stakeholders rather than delegating it entirely to IT.
- **Current Issue:** Many organizations, particularly small and medium enterprises, still lack a designated individual or team responsible for data management, leaving the function fragmented across departments.
- **Emerging Trend:** The rise of AI and machine learning has raised the stakes of poor data management considerably, since flawed or biased underlying data directly produces flawed or biased automated decisions at scale.

1.10 Summary and Key Takeaways

Data management is the disciplined practice of collecting, storing, securing, governing, and deriving value from an organization's data throughout its useful life. It is distinct from, though foundational to, both information technology and data science, and has shifted from a

technical convenience to an organizational necessity under regulatory, competitive, and operational pressure.

- Common Pitfall: Assuming data management is solely IT's responsibility, rather than a cross-functional discipline requiring legal, business, and technical involvement.
- Common Pitfall: Treating regulatory compliance as the only reason to invest in data management, overlooking the competitive and operational value that well-managed data provides.
- Common Pitfall: Waiting for a visible data failure, a bad decision, or a compliance complaint before investing in basic data management discipline.

KEY TAKEAWAYS

- Data management is cross-functional, spanning legal, business, and technical stakeholders, not solely an IT responsibility.
- Data, information, and insight represent successive stages of value extraction from the same underlying facts.
- The Data Privacy Act (RA 10173) makes formal data management a legal requirement for organizations handling personal data in the Philippines.
- Poor data management compounds silently, unlike physical asset failures, and often surfaces only when a costly decision or compliance failure occurs.

1.11 Key Terms, Reflection, and Discussion

Data Management. Data. Information. Data Governance. Data Lifecycle. DAMA-DMBOK. Data Privacy Act.

- Reflection: Think of an organization you have interacted with, as an employee, customer, or student. What signs suggested its data management was strong or weak?
- Discussion: Should data management be led by IT, by business units, or shared equally? Defend your position.
- Discussion: Is the growing legal weight behind data management, such as the Data Privacy Act, primarily a burden on organizations or a necessary protection for individuals? Can it be both?
- Discussion: Table 1.1 distinguishes data, information, and insight. Can you think of a situation where more data actually made a decision worse, rather than better?

1.12 Activity and Data Project Build, Step 1

ACTIVITY: DATA MANAGEMENT MATURITY AUDIT

Step 1: Select an organization you have access to (a workplace, student organization, or family business).

Step 2: Identify what data it collects and where that data currently lives.

Step 3: Assess informally whether the organization has any documented data governance, quality, or security practices.

Deliverable: A one-page maturity summary, rating the organization Low, Medium, or High on data management discipline, with justification.

DATA PROJECT BUILD, STEP 1

Select the dataset or data domain you will carry through this entire book's Data Project Build (a public open dataset, a dataset from your workplace, or a realistic dataset you construct). Write a one-paragraph statement of what organizational question this dataset could help answer, and identify which data management challenges (collection, quality, security, governance) you expect to be most significant for it.

1.13 References

- Codd, E. F. (1970). A relational model of data for large shared data banks. Communications of the ACM, 13(6), 377 to 387.
- DAMA International. (2017). DAMA-DMBOK: Data management body of knowledge (2nd ed.). Technics Publications.
- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- Silberschatz, A., Korth, H. F., & Sudarshan, S. (2020). Database system concepts (7th ed.). McGraw-Hill.
- National Privacy Commission. Data Privacy Act primer (latest edition).

End of Chapter 1. Continue to Chapter 2: Key Components and Data Governance

KEY COMPONENTS AND DATA GOVERNANCE

The Structural Elements That Hold a Data Management Program Together

“Governance is not the paperwork around data management. It is the reason data management survives past the person who first set it up.”

2.0 Chapter Overview

Chapter 1 established what data management is and why it matters. This chapter breaks the discipline into its constituent components and introduces data governance, the framework that coordinates those components into a coherent, accountable program rather than a loose collection of independent practices. Understanding governance early matters because every later chapter, from storage to security to analytics, operates within the accountability structure this chapter establishes.

2.1 Learning Outcomes

- LO 2.1 Identify the key components of a comprehensive data management system.
- LO 2.2 Explain the role of data governance and its core principles.
- LO 2.3 Distinguish between key data governance roles and their respective accountabilities.

2.2 Introduction

An organization can own excellent database technology, a skilled technical team, and a genuine intention to manage its data well, and still fail, if no one is clearly accountable for data quality, no policy defines who may access sensitive fields, and no process exists for resolving disagreements about what a given data field actually means. These are governance failures, not technical ones, and they are strikingly common. This chapter teaches you to recognize the components a mature data management program requires and the governance structure that binds them.

It is worth noting why this chapter comes second, immediately after the foundational overview of Chapter 1 and before the more technical chapters on collection, storage, or processing. Every technical practice this book covers, from choosing a storage architecture in Chapter 5 to running a cleansing routine in Chapter 7, is more effective and more durable when it operates within a clear governance structure. Skipping ahead to technical practice without this foundation tends to produce isolated, ungoverned pockets of good practice that do not scale across an organization.

Students sometimes find governance the least immediately engaging topic in this book, less visibly technical than a storage architecture decision, less immediately consequential than a security control. That impression tends to fade once a student has actually experienced a governance failure firsthand, whether in a workplace, a group project, or their own Data Project Build later in this course. Two people, or two systems, quietly disagreeing about what a shared piece of data actually means is one of the most common and most preventable sources of organizational dysfunction that this book will describe, and this chapter is where you learn to prevent it.

2.3 Historical Background

Data governance emerged as a formalized discipline later than data management technology itself, largely in response to large organizations discovering, often painfully, that having sophisticated databases did not prevent inconsistent, duplicated, or contradictory data across different systems. Financial services and healthcare organizations were early adopters of formal governance structures, driven by regulatory requirements that demanded a demonstrable chain of accountability for data quality and access, well before governance became common practice in other sectors.

The Data Management Association (DAMA International) played a significant role in standardizing governance vocabulary and practice through its DMBOK framework, first published in the 2000s and revised since, providing organizations across industries with a shared reference model rather than having to reinvent governance structures independently. This standardization is part of why this book, and the MDM program it supports, draws on DAMA's framework for its own treatment of governance roles and principles.

It is worth noting why financial services and healthcare specifically led this early adoption, since the pattern reveals something durable about when governance investment becomes

organizationally urgent rather than merely advisable. Both sectors faced regulatory examiners who could, and did, ask pointed questions such as who is accountable for this specific data field and how do you know it is accurate, questions an organization without named stewards and documented definitions simply could not answer convincingly. Other sectors have followed this same trajectory more recently as their own regulatory environments, including the Philippines' Data Privacy Act, began asking similarly pointed questions.

2.4 Definitions

KEY TERMS

Data Governance: The framework of policies, roles, and decision rights that determines how data is managed, who is accountable for it, and how quality and compliance are enforced across an organization.

Data Steward: An individual accountable for the quality, definition, and appropriate use of a specific data domain, acting as the operational link between governance policy and day-to-day data handling.

Data Owner: The individual or role with ultimate decision-making authority over a given data domain, typically a senior business stakeholder rather than a technical role.

Data Dictionary: A centralized, documented reference defining every significant data field an organization uses, including its meaning, format, and permitted values.

Master Data: The core, authoritative data entities an organization relies on across multiple systems, such as customer, product, or employee records.

2.5 Core Concepts and Frameworks




Figure 2.1. Data Governance Roles.

2.5.1 The Key Components of a Data Management System

A comprehensive data management system rests on several interlocking components: data architecture, the overall design of how data is structured and flows through an organization; data storage and operations, the physical and logical infrastructure covered in Chapters 4 and 5; data security, covered in Chapters 10 and 11; data quality, covered in Chapters 6 and 7; and data governance itself, the accountability layer this chapter focuses on. No single component functions well in isolation from the others.

2.5.2 Core Governance Principles

Effective data governance rests on a few durable principles. Accountability requires that every significant data domain have a named owner and steward, not a diffuse, shared responsibility that in practice belongs to no one. Transparency requires that data definitions, quality standards, and access rules be documented and discoverable, not held informally in individual employees' heads. Consistency requires that the same data element mean the same thing everywhere it appears in the organization, enforced through tools such as the data dictionary and master data management practices.

2.5.3 Governance Roles and the RACI Structure

Mature governance programs typically define roles using a RACI structure, specifying who is Responsible for day-to-day data handling, who is Accountable for outcomes (typically the data owner), who must be Consulted before changes are made, and who must simply be Informed. Applied to data governance, this usually places data stewards in the Responsible role for a given domain, data owners in the Accountable role, and a central data governance committee or council in a Consulted and coordinating capacity across domains.

2.6 Table 2.1, Data Governance Roles at a Glance

Role	Typical Accountability	Example
Data Owner	Ultimate decision authority over a data domain	VP of Sales owns customer data
Data Steward	Day-to-day quality and definition management	Analyst maintaining the customer data dictionary
Data Governance Council	Cross-domain policy and dispute resolution	Committee resolving conflicting definitions of “active customer”

Role	Typical Accountability	Example
Chief Data Officer (or equivalent)	Organization-wide governance strategy and sponsorship	Executive sponsoring the governance program

Table 2.1. Core data governance roles and their typical accountabilities.

2.7 Government Policy and Philippine Context

The Data Privacy Act's requirement that organizations designate a Data Protection Officer (DPO) is, in effect, a legally mandated governance role, giving Philippine organizations a regulatory anchor point around which broader data governance structures can be built. The National Privacy Commission's guidelines on DPO responsibilities overlap substantially with the data governance principles described in this chapter, particularly regarding accountability and documented policy.

PHILIPPINE CONTEXT

Many Philippine organizations, particularly smaller enterprises, satisfy the letter of the Data Privacy Act's DPO requirement without building out the broader governance structure this chapter describes, appointing a DPO on paper without the data stewards, data dictionary, or governance council that would make the role effective in practice. Graduates who understand governance as a full structure, not a single compliance appointment, are positioned to close this common and consequential gap.

This gap is often not a matter of neglect but of unfamiliarity: many Philippine organizations that appointed a DPO in direct response to the Data Privacy Act's requirement did so as a discrete, one-time compliance task, without a clear internal reference for what a fully functioning governance structure around that role should look like. A graduate entering such an organization with this chapter's RACI framework already internalized can often make a concrete, immediate contribution simply by proposing the missing pieces, named stewards, a documented dictionary, a functioning council, rather than needing to convince leadership that governance matters at all.

2.8 Industry Practice and Case Study

Large financial institutions typically operate the most mature governance structures of any Philippine industry sector, driven by both the Data Privacy Act and Bangko Sentral ng Pilipinas regulatory requirements, often maintaining dedicated data governance offices with formally chartered councils and documented escalation procedures for data disputes.

CASE STUDY

Illustrative composite case.

A Philippine retail chain's marketing and finance departments each maintained their own definition of “active customer,” marketing counting anyone who had browsed the loyalty app in the past ninety days, finance counting only those with a completed purchase in the same window. Reports using the two definitions produced customer counts that differed by nearly 40%, and executive meetings routinely stalled over which number was correct. The company's eventual formation of a data governance council, chaired by a newly appointed data owner for customer data, resolved the conflict by establishing one authoritative definition, documented in a shared data dictionary, that both departments were required to use.

Resolving the definitional conflict required more than simply picking one department's existing definition over the other, since each definition served a genuinely different, legitimate purpose. Marketing's broader definition supported engagement campaigns, while finance's narrower definition supported revenue forecasting. The governance council's actual solution retained both definitions as separately named, documented metrics, engaged customers, and paying customers, rather than forcing a single compromise definition that would have served neither department's original purpose well, a resolution that depended entirely on the council having the authority to formalize two coexisting definitions rather than being expected to produce one.

2.9 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Name a specific data owner and steward for every significant data domain before attempting to enforce quality or security policy for that domain.

- **Current Issue:** Governance programs frequently stall when they are treated as a one-time documentation exercise rather than an ongoing accountability structure with active dispute resolution.
- **Emerging Trend:** Automated data governance tools, which programmatically enforce data definitions and lineage tracking, are increasingly supplementing purely manual governance processes.
- **Common Pitfall:** Appointing a Data Protection Officer to satisfy legal requirements without building the broader stewardship and documentation structure that makes governance actually function.
- **Common Pitfall:** Allowing multiple departments to maintain conflicting definitions of the same core data entity without a governance council empowered to resolve the conflict.
- **Common Pitfall:** Treating the data dictionary as a one-time deliverable rather than a living document maintained as data definitions evolve.

2.10 Summary and Key Takeaways

A comprehensive data management system rests on interlocking components, architecture, storage, security, quality, and governance, none of which functions well in isolation. Data governance provides the accountability structure, through named owners, stewards, and documented policy, that allows these components to operate consistently and durably across an organization rather than as isolated pockets of good practice.

It is worth being direct about the stakes of getting this chapter's material right before proceeding. Every remaining chapter in this book teaches a technical or analytical practice, collection, storage, cleansing, security, and analytics that depend on the governance foundation described here to scale beyond a single well-intentioned individual. A founder or analyst who skips ahead to technical practice without this foundation risks building good habits that quietly disappear the moment they leave the organization.

- Common Pitfall: Leaving a data domain without a named owner or steward, so that quality and access disputes have no clear path to resolution.
- Common Pitfall: Treating the data dictionary as a one-time document rather than a living reference updated as definitions evolve.
- Common Pitfall: Satisfying the Data Protection Officer's legal requirement without building the broader stewardship structure that role depends on.

KEY TAKEAWAYS

- A comprehensive data management system includes architecture, storage, security, quality, and governance as interlocking components.
- Data governance rests on accountability, transparency, and consistency, enforced through named roles like data owners and stewards.
- The RACI structure clarifies who is responsible, accountable, consulted, and informed for a given data domain.
- The Data Privacy Act's DPO requirement provides a legal anchor for Philippine governance programs, but effective governance requires more than a single compliance appointment.

2.11 Key Terms, Reflection, and Discussion

Data Governance. Data Steward. Data Owner. Data Dictionary. Master Data. RACI. Data Protection Officer.

- Reflection: For the organization or dataset you selected in Chapter 1's Data Project Build, who should be the data owner, and who should be the data steward? Justify your answer.

- Discussion: Can effective data governance exist without a dedicated governance council, relying instead on informal coordination? Under what conditions might this work, and when would it fail?
- Discussion: Should the Data Protection Officer role required by the Data Privacy Act be expanded by law to include broader data governance responsibilities, or kept narrowly focused on privacy compliance?
- Discussion: Table 2.1 assigns clear roles to the owner, steward, and council. What happens in a very small organization where the same person must play more than one of these roles?

2.12 Activity and Data Project Build, Step 2

VENTURE-STYLE ACTIVITY: GOVERNANCE ROLE MAPPING

For a real or hypothetical organization, draft a one-page governance charter identifying the data owner, data steward, and governance council membership for one significant data domain (such as customer data or student records), along with two governance policies that each role would be responsible for enforcing.

DATA PROJECT BUILD, STEP 2

For the dataset you selected in Step 1, identify who would serve as its data owner and data steward, and draft a short data dictionary entry (name, definition, format, permitted values) for three key fields in your dataset.

2.13 References

- DAMA International. (2017). DAMA-DMBOK: Data management body of knowledge (2nd ed.). Technics Publications.

- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- National Privacy Commission. Guidelines on the appointment of data protection officers.
- Connolly, T., & Begg, C. (2015). Database systems: A practical approach to design, implementation, and management (6th ed.). Pearson.

End of Chapter 2. Continue to Chapter 3: The Data Lifecycle

THE DATA LIFECYCLE

Following Data from Creation to Disposal

“Data is not permanent by default. It is permanent only where an organization has deliberately decided it should be.”

3.0 Chapter Overview

This chapter closes Part I by tracing data through its entire lifecycle, from its creation or collection to its eventual archival or deletion. Understanding the lifecycle as a whole, before Part II examines collection and storage in technical depth, gives you a map of where each later chapter's material fits within the larger journey any piece of organizational data actually takes.

3.1 Learning Outcomes

- LO 3.1 Describe the stages of the data lifecycle and the key activities within each.
- LO 3.2 Explain the purpose of data retention and disposal policies.
- LO 3.3 Apply lifecycle thinking to identify governance gaps in a real or hypothetical data domain.

3.2 Introduction

Organizations tend to invest heavily in the early stages of the data lifecycle, collection and storage, while giving far less deliberate attention to what happens as data ages: when it should be archived, when it should be deleted, and what obligations attach to data the organization has decided to keep. This imbalance is not accidental. Collection and storage produce visible, immediate value. In contrast, disposal produces no visible benefit until an audit, breach, or regulatory inquiry reveals that data that should have been deleted years earlier is still sitting, unmonitored, in a forgotten system.

This chapter treats the lifecycle as a complete loop, deliberately including the stages that organizations most often neglect. By the end of this chapter, you should be able to look at any

dataset, including the one you selected for your own Data Project Build, and identify not only how it should be collected and stored but also how long it should be retained and what should happen to it afterward.

This chapter also closes Part I with a deliberate purpose: it hands you a complete map before Part II asks you to zoom into any single stage of that map in technical depth. Keeping the full lifecycle in view, even while immersed in the collection-specific detail of Chapter 4 or the storage-specific detail of Chapter 5, prevents a common student error, treating each subsequent chapter's stage as though it existed in isolation, disconnected from the retention and disposal obligations this chapter has just established for it.

3.3 Historical Background

Lifecycle thinking in data management draws directly on records management, a much older discipline concerned with the retention and disposal of physical documents, which itself predates digital data. Government and legal recordkeeping requirements, dictating how long specific document types must be retained before disposal, provided an early template that data management adapted for the digital era. However, digital data's near-zero marginal storage cost initially made organizations far less disciplined about disposal than paper recordkeeping had required.

This changed decisively with the arrival of stringent data protection regulations, particularly the European Union's General Data Protection Regulation in 2018, which explicitly grants individuals a right to have their personal data deleted under certain conditions, the so-called right to erasure. This legal development forced organizations to treat the disposal stage of the lifecycle with the same rigor previously reserved for collection and storage, a shift the Philippines' own Data Privacy Act broadly mirrors.

The near-zero marginal cost of digital storage deserves particular attention as a historical explanation for why disposal discipline lagged so far behind collection and storage discipline. Physical recordkeeping imposed a natural, tangible cost on indefinite retention: filing cabinets filled up, and warehouses had finite space, making disposal decisions unavoidable in the long run, regardless of organizational discipline. Digital storage removed this natural forcing function almost entirely, and it took the arrival of legal requirements specifically targeting retention and disposal to reintroduce a comparable discipline, since cost alone no longer provided the pressure that physical storage had once automatically exerted.

3.4 Definitions

KEY TERMS

Data Lifecycle: The complete sequence of stages data passes through, typically collection, processing, storage, usage, archiving, and disposal.

Data Retention Policy: A documented rule specifying how long a given category of data must or may be kept before archival or deletion.

Data Archiving: Moving data that is no longer actively used to lower-cost, longer-term storage, while preserving it for potential future reference or legal requirement.

Data Disposal: The permanent, secure deletion of data once it is no longer needed, and no retention obligation requires it to be kept.

Right to Erasure: a legal right, prominent in data protection regulation, that allows individuals to request the deletion of their personal data under specified conditions.

3.5 Core Concepts and Frameworks



Figure 3.1. The Data Lifecycle.

3.5.1 The Six Stages of the Data Lifecycle

Data typically moves through six identifiable stages. Collection is the point of origin, covered fully in Chapter 4. Processing transforms raw collected data into a usable form,

covered in Chapters 8 and 9. Storage holds data for active or future use, covered in Chapter 5. Usage is the stage at which data serves its organizational purpose, whether through operational systems or analytics, as covered throughout Part VI. Archiving moves data no longer in active use to lower-cost, longer-term storage. Permanently and securely disposing of data removes it once no purpose or legal obligation requires its retention.

3.5.2 Why Retention Policy Requires Active Decisions

A retention policy answers a deceptively difficult question for each data category an organization holds: how long should this data be kept, and why? The answer typically balances three considerations that often pull in different directions. Legal and regulatory requirements may mandate a minimum retention period, such as for financial records. Business value may justify retention beyond the legal minimum if the data supports ongoing analytics or decision-making. Privacy and risk considerations argue for the shortest retention period consistent with these other needs, since data that no longer serves a purpose is pure liability, valuable to no one but costly to secure, and a growing target the longer it exists.

3.5.3 Archiving vs. Disposal: A Deliberate Choice, Not a Default

Organizations without a deliberate lifecycle policy tend to default to archiving everything indefinitely, since archiving requires no difficult decision, and modern storage costs are low. This default is a governance failure, not a neutral choice: indefinitely archived data still carries security risk, still falls under data subject access and erasure rights where applicable, and still requires the organization to know what it holds and why. A mature lifecycle policy treats archiving as a deliberate, time-bound decision, not as the absence of a decision.

3.6 Table 3.1, The Six Lifecycle Stages

Stage	Key Governing Question	Chapter Reference
Collection	Is this data being gathered lawfully and accurately?	Chapter 4
Processing	Is this data being transformed correctly and traceably?	Chapters 8 to 9
Storage	Is this data stored securely and appropriately for its use?	Chapter 5
Usage	Is this data being used for its intended, authorized purpose?	Part VI

Stage	Key Governing Question	Chapter Reference
Archiving	Should this data still be kept, and in what form?	Chapter 3
Disposal	Has the retention period ended, and can this data be safely deleted now?	Chapter 3

Table 3.1. The data lifecycle and its governing questions.

3.7 Government Policy and Philippine Context

The Data Privacy Act's principle of proportionality, requiring that personal data be retained only as long as necessary for the purpose for which it was collected, directly implements lifecycle thinking into Philippine law. The National Privacy Commission has issued guidance indicating that indefinite retention without a documented, purpose-linked justification is itself a compliance violation, independent of whether the data is ever misused.

PHILIPPINE CONTEXT

Philippine organizations frequently retain customer and employee records well beyond any legal or operational necessity, often because no retention policy exists to trigger disposal, rather than through any deliberate decision to keep the data. This is a widespread, quietly accumulating compliance risk under the Data Privacy Act, and one of the most immediately actionable governance improvements a new data professional can identify and address in a Philippine organization.

A useful first step for a graduate encountering this pattern is not to propose an immediate, large-scale disposal effort, which can feel risky and disruptive to an organization unaccustomed to it, but to begin with a modest, defined retention policy applied only to newly created records in the future, while separately auditing and addressing the existing backlog on its own timeline. This staged approach tends to secure organizational buy-in far more easily than an all-at-once disposal initiative, and it directly mirrors the phased, evidence-based approach this book has modeled throughout.

3.8 Industry Practice and Case Study

CASE STUDY

Illustrative composite case.

A Philippine university's admissions office retained complete application files, including personal and financial documents, for every applicant dating back over a decade, including thousands who were never admitted, reasoning informally that the data “might be useful someday.” A data governance review found no documented retention policy justified this practice, no legal requirement mandated it beyond a much shorter period, and the accumulated volume of unnecessary personal data represented a significant, unmanaged compliance exposure under the Data Privacy Act. The office's subsequent retention policy, specifying a defined retention period for unsuccessful applications followed by secure disposal, reduced this risk and freed storage resources previously used to maintain data with no remaining purpose.

Implementing the new policy also surfaced a secondary problem the original review had not anticipated: several of the oldest application files existed only in a discontinued document format that current staff no longer had software capable of opening, meaning the office had been paying to store personal data it could not have actually retrieved, even if a legitimate reason to do so had arisen. This discovery reinforced the office's decision to secure disposal on a defined schedule in the future, rather than defaulting again to indefinite retention, since data that has become technically inaccessible still carries full legal retention risk without offering any of the practical value retention was originally meant to preserve.

3.9 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Define a documented retention period for every significant data category at the time it is first collected, rather than retroactively once storage or compliance concerns arise.
- Current Issue: Automated enforcement of retention and disposal policies remains uncommon in smaller Philippine organizations, leaving compliance dependent on manual, easily neglected processes.

- Emerging Trend: Data lifecycle management platforms increasingly automate the archiving and disposal stages, flagging or deleting data according to policy without requiring manual review of every record.
- Common Pitfall: Defaulting to indefinite archiving because storage is cheap, without recognizing the ongoing security and compliance liability this creates.
- Common Pitfall: Writing a retention policy but never implementing the disposal mechanism that would actually enforce it.
- Common Pitfall: Treating disposal as equivalent to deletion from a single visible system, without confirming data has also been removed from backups and secondary copies.

3.10 Summary and Key Takeaways

This chapter closes Part I by giving you a complete map of the data lifecycle, six stages spanning collection through disposal, and a governing question for each. Part II now turns to the first of these stages in technical depth. Still, the lifecycle view you have built here should stay active throughout: every collection or storage decision you make in the chapters ahead has downstream consequences for how that same data will eventually need to be archived or disposed of.

The data lifecycle encompasses all stages data passes through, from creation and collection to eventual archival or deletion, and organizations that manage only the early stages well while neglecting archiving and disposal accumulate quietly, compounding risk. Deliberate retention and disposal policy, grounded in legal requirements, business value, and privacy risk, closes this gap.

- Common Pitfall: Defaulting to indefinite retention because deletion feels risky, without recognizing that unnecessary retention is itself a compliance and security liability.
- Common Pitfall: Writing a retention policy without implementing the disposal mechanism that would actually enforce it in practice.

- Common Pitfall: Treating disposal as complete once data is removed from the primary system, without confirming that backups and secondary copies are also addressed.

KEY TAKEAWAYS

- The data lifecycle has six stages: collection, processing, storage, usage, archiving, and disposal.
- Retention policy should balance legal requirements, business value, and privacy risk for each data category.
- Indefinite archiving by default is a governance failure, not a neutral or safe choice.
- The Data Privacy Act's proportionality principle makes purpose-linked retention a legal requirement in the Philippines, not merely good practice.

3.11 Key Terms, Reflection, and Discussion

Data Lifecycle. Data Retention Policy. Data Archiving. Data Disposal. Right to Erasure. Proportionality Principle.

- Reflection: For your Data Project Build dataset, what would an appropriate retention period be, and what would justify it, legal requirement, business value, or both?
- Discussion: Should individuals have an unconditional right to demand deletion of their own data, or should organizational and public-interest exceptions apply? Where should the line be drawn?
- Discussion: This chapter argues indefinite archiving is a governance failure, not a safe default. Can you think of a legitimate reason an organization might deliberately choose to retain data far longer than any legal requirement demands?

3.12 Activity and Data Project Build, Step 3

DATA PROJECT BUILD, STEP 3

Draft a lifecycle map for your Data Project Build dataset, specifying, for each of the six stages, what happens to the data and who is responsible for it. Propose a specific retention period for the dataset and justify it against legal requirements, business value, and privacy risk.

3.13 References

- DAMA International. (2017). DAMA-DMBOK: Data management body of knowledge (2nd ed.). Technics Publications.
- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- Ma, X., Fox, P., Rozell, E., West, P., & Zednik, S. (2014). Ontology dynamics in a data lifecycle: Challenges and recommendations from a geoscience perspective. *Journal of Earth Science*, 25(2), 407 to 412.
- National Privacy Commission. Guidelines on data retention and disposal.

End of Chapter 3. Continue to Chapter 4: Data Sources, Types, and Collection Methods

PART II

DATA COLLECTION AND STORAGE

Part II turns to the first substantive stages of the data lifecycle: where data actually comes from, and where it goes once collected. Sound collection and storage decisions determine how effectively subsequent stages of data management can operate.

DATA SOURCES, TYPES, AND COLLECTION METHODS

Where Data Comes from and How to Gather It Responsibly

“The quality of every downstream analysis is set, and limited, at the moment data is first collected.”

4.0 Chapter Overview

Part I gave you the vocabulary, governance structure, and lifecycle view that frame this entire book. Part II now turns to the first substantive stage of that lifecycle: collection. This chapter examines where organizational data actually comes from, the fundamental distinction between structured and unstructured data, and the collection methods available for gathering each responsibly and accurately.

4.1 Learning Outcomes

- LO 4.1 Distinguish primary and secondary data sources, and structured and unstructured data types.
- LO 4.2 Identify appropriate data collection methods and techniques for a given data need.
- LO 4.3 Evaluate the advantages and challenges associated with different collection approaches.

4.2 Introduction

Every downstream chapter in this book, quality, security, processing, and analytics, works with data that already exists by the time those chapters begin. This chapter is the exception: it addresses the moment data first enters an organization's systems, the point at which errors are cheapest to prevent and most expensive to fix later. A flawed collection instrument, an ambiguous survey question, an incorrectly configured sensor, and an undocumented manual

entry process plant a defect that every later stage of the lifecycle inherits and often cannot fully correct.

Part II, spanning this chapter and Chapter 5, is where the book's abstractions from Part I meet concrete practice for the first time. Governance and lifecycle thinking remain essential background, but this chapter asks a narrower, more immediate question: given everything Part I established about why data management matters, how do you actually bring data into an organization's systems in a way that respects those principles from the very first moment of contact?

4.3 Historical Background

Data collection as a formal discipline predates computing entirely, rooted in centuries of statistical survey methodology developed for census-taking, scientific research, and government administration. What changed with digitization was not the underlying logic of good collection practice, representative sampling, minimizing measurement error, documenting method, but the sheer scale and automation now possible: a sensor network or web application can collect more data in a day than a traditional survey team could gather in years.

This scale shift introduced a genuinely new category of data collection: passive, automated collection that occurs as a byproduct of some other activity, a customer completing a purchase, a user browsing a website, a device transmitting a sensor reading, rather than through any deliberate act of data gathering. This passive collection now represents the majority of data most organizations hold, and it raises collection challenges, particularly around consent and transparency, that traditional survey methodology never had to address.

The census-taking tradition this chapter traces its roots to is worth noting for a specific reason: census methodology developed strong norms around documenting method and limitation precisely because census results directly shaped resource allocation and political representation, making errors visible and consequential in ways that motivated methodological rigor. Modern automated data collection, precisely because it happens passively and at enormous scale, can lose sight of this same documentation discipline, producing data whose collection method and limitations are effectively invisible to the analysts who later use it, a gap this chapter's emphasis on documentation is meant to close.

4.4 Definitions

KEY TERMS

Primary Data: Data collected directly from a source or activity for a specific purpose, first-hand and unprocessed.

Secondary Data, Data originally collected for another purpose, obtained from an existing source such as government statistics or a prior study, and reused.

Structured Data, Data organized according to a predefined model or schema, typically stored in rows and columns, such as a transactional database.

Unstructured Data, Data that does not follow a predefined model or schema, such as text documents, images, video, or social media content.

Data Collection Instrument: The specific tool or mechanism used to gather data, such as a survey form, sensor, or automated logging system.

4.5 Core Concepts and Frameworks




Figure 4.1. Structured vs. Unstructured Data.

4.5.1 Primary vs. Secondary Data Sources

Primary data sources involve data collected directly from original activities, events, or experiments for a specific purpose. Its chief advantage is that it is tailored precisely to the collecting organization's own question, though gathering it well requires deliberate method and typically more resources than reusing existing data. Secondary data sources draw on data already collected by another party for a different original purpose, such as government census

data or industry reports. This offers considerable speed and cost advantages, but carries the challenge that the data may not perfectly match the current question, its collection methodology may be unclear, and it may already be outdated by the time it is reused.

4.5.2 Structured and Unstructured Data

Understanding data types is critical in data management because it determines how data can be stored, handled, and analyzed effectively. Structured data fits neatly into the row-and-column model traditional relational databases were built for, and includes examples such as transaction records, sensor readings logged at regular intervals, and form-based survey responses. Unstructured data, encompassing text files such as emails and documents, multimedia files such as images and video, and social media content, does not fit this model and requires specialized storage and analysis tools, which are further developed in Chapter 5. A growing middle category, semi-structured data, such as JSON or XML files, carries some organizational markers without conforming to a rigid schema.

4.5.3 Collection Methods: Quantitative, Qualitative, Mixed, and Automated

Quantitative data collection gathers numerical, measurable data, typically through structured surveys, sensors, or transactional logging, well-suited to statistical analysis. Qualitative data collection gathers descriptive, non-numerical data, typically through interviews, focus groups, or open-ended survey responses, well-suited to understanding context, motivation, and nuance that a purely numerical approach would miss. Mixed-methods data collection deliberately combines both approaches within a single study, capturing the statistical generalizability of quantitative methods alongside the explanatory depth of qualitative methods. Automated collection, increasingly the dominant method for organizational data, gathers data passively through system logs, sensors, or application usage tracking, without requiring active participation from the person or system being measured.

4.6 Table 4.1, Structured vs. Unstructured Data

Dimension	Structured Data	Unstructured Data
Format	Predefined schema, rows, and columns	No predefined format

Dimension	Structured Data	Unstructured Data
Examples	Transaction records, sensor logs, survey responses	Emails, documents, images, video, social media posts
Storage	Relational databases	Data lakes, document stores, object storage (Chapter 5)
Analysis Tools	SQL, traditional statistical software	Text mining, computer vision, natural language processing

Table 4.1. Comparing structured and unstructured data.

4.7 Government Policy and Professional Standards

Data collection involving personal information is directly governed by the Data Privacy Act's principle of transparency, requiring that data subjects be informed of what data is collected about them, for what purpose, and with what consent, before or at the point of collection. This obligation applies equally to active collection methods, such as surveys, and passive, automated collection. One of them is website tracking, a distinction that catches many organizations unprepared, since automated collection often occurs without any explicit consent interaction resembling a traditional survey's informed consent process.

4.8 Philippine Context and Case Study

PHILIPPINE CONTEXT

The Philippine Statistics Authority operates as the country's primary secondary data source for demographic, labor, and economic data, freely accessible and widely used across academic and private-sector research. For primary data collection, the widespread use of mobile messaging and social media platforms among Filipino respondents has made these channels increasingly viable and cost-effective survey distribution methods, particularly for reaching younger or geographically dispersed populations that traditional paper or phone surveys reach less effectively.

Founders and researchers relying on PSA data should note that its regional and provincial breakdowns can lag national figures in release timing, and that certain indicators are collected only at specific census intervals rather than continuously,

meaning a secondary-data researcher may need to supplement PSA figures with more recent primary collection for time-sensitive questions. Understanding this publication cadence in advance prevents a common early mistake, discovering partway through a project that the specific, recent regional figure a study depends on does not yet exist in published PSA form.

CASE STUDY

Illustrative composite case.

A Philippine local government unit sought to understand community health needs. Initially, it relied entirely on existing secondary data from national health statistics, which showed broad regional trends but could not answer specific questions about barriers residents faced in accessing local health services. A subsequent mixed-methods primary collection effort, a structured household survey for quantitative prevalence data combined with focus group discussions for qualitative insight into access barriers, revealed transportation cost as an under-recognized obstacle that the secondary data alone had never surfaced, directly informing a targeted mobile clinic program that the purely secondary-data approach would not have identified.

The focus group discussions specifically surfaced a nuance the household survey's structured questions had not been designed to capture: transportation cost was less about the fare itself than about the unpredictability of travel time on certain routes, which made residents reluctant to risk missing work for a clinic visit that might take an entire day rather than the expected hour. This qualitative detail directly shaped the mobile clinic program's eventual scheduling design, offering predictable, brief appointment windows rather than the open queuing system the quantitative data alone would have suggested was sufficient.

4.9 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Document the collection method, instrument, and known limitations for every dataset at the time of collection, not retroactively when questions arise later.

- **Current Issue:** Automated, passive data collection frequently outpaces organizations' consent and transparency practices, creating compliance exposure under the Data Privacy Act.
- **Emerging Trend:** Synthetic data generation, creating artificial data that preserves the statistical properties of real data without exposing actual personal information, is increasingly used to reduce collection and privacy risk during testing and development.
- **Common Pitfall:** Reusing secondary data without verifying its original collection methodology, leading to unrecognized bias or incompatibility with the current question.
- **Common Pitfall:** Deploying automated, passive data collection without a clear, documented consent and transparency mechanism.
- **Common Pitfall:** Assuming quantitative data alone tells the complete story, missing the explanatory context that qualitative methods would have revealed.

4.10 Summary and Key Takeaways

Data enters an organization through primary sources, collected directly for a specific purpose, or secondary sources, reused from existing data collected for another purpose, and takes the form of structured, unstructured, or semi-structured data, depending on whether it conforms to a predefined schema. Collection methods span quantitative, qualitative, mixed, and increasingly automated approaches, each suited to different questions and each carrying distinct advantages, challenges, and Data Privacy Act obligations.

- **Common Pitfall:** Reusing secondary data without verifying its original collection methodology, risking unrecognized bias or mismatch with the current question.
- **Common Pitfall:** Deploying automated, passive collection without a clear consent and transparency mechanism, creating unrecognized Data Privacy Act exposure.

- Common Pitfall: Relying entirely on quantitative data when qualitative context is needed to understand the actual cause behind a pattern.

KEY TAKEAWAYS

- Primary data is tailored to a specific purpose but costlier to gather; secondary data is faster and cheaper but may not perfectly fit the current question.
- Structured data fits predefined schemas; unstructured data, an increasing majority of organizational data, does not.
- Mixed-methods collection combines quantitative generalizability with qualitative explanatory depth.
- Automated, passive collection requires the same Data Privacy Act transparency obligations as active survey collection, though it is frequently overlooked.

4.11 Key Terms, Reflection, and Discussion

Primary Data. Secondary Data. Structured Data. Unstructured Data. Semi-Structured Data. Quantitative Collection. Qualitative Collection. Mixed Methods.

- Reflection: Is your Data Project Build dataset primary or secondary, structured or unstructured? What collection method would be, or was, most appropriate for it?
- Discussion: As automated, passive collection becomes the dominant method for organizational data, does traditional informed consent remain a meaningful concept? What would a genuinely transparent automated collection practice look like?

- Discussion: This chapter argues that secondary data is faster but may not fit the current question well. Can you think of a situation where secondary data would be actively misleading rather than merely imperfect for a given purpose?

4.12 Activity and Data Project Build, Step 4

DATA PROJECT BUILD, STEP 4

Classify your Data Project Build dataset by source (primary or secondary) and type (structured, unstructured, or semi-structured). Draft a one-paragraph data collection statement describing how the data was, or would be, collected, including the specific consent or transparency mechanism appropriate under the Data Privacy Act if it involves personal information.

4.13 References

- Babbie, E. (2010). The practice of social research (12th ed.). Cengage Learning.
- Creswell, J. W. (2014). Research design: Qualitative, quantitative, and mixed methods approaches (4th ed.). SAGE Publications.
- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- Philippine Statistics Authority. Open data and statistical standards portal.
- Levity AI. What is data extraction? Structured and unstructured data explained.

End of Chapter 4. Continue to Chapter 5: Data Storage Systems and Architectures

DATA STORAGE SYSTEMS AND ARCHITECTURES

Choosing Where and How Data Should Live

“The right storage architecture is invisible when it works, and catastrophic when it was never chosen deliberately.”

5.0 Chapter Overview

Chapter 4 examined the sources of data. This chapter examines where it goes once collected: the storage systems and architectures available to an organization, and the factors that should drive the choice among them. This chapter closes Part II by completing the collection-and-storage stage of the data lifecycle before Part III turns to ensuring the quality of the data these systems hold.

5.1 Learning Outcomes

- LO 5.1 Identify major types of data storage systems and their appropriate use cases.
- LO 5.2 Compare data storage architectures, including centralized and distributed approaches.
- LO 5.3 Evaluate factors relevant to choosing a data storage solution, including security considerations.

5.2 Introduction

Choosing a storage system is not a purely technical decision to be delegated entirely to IT, though it is frequently treated that way. The choice has direct consequences for data quality, since some storage systems enforce structure and validation that others do not; for security, since different systems carry different inherent risk profiles; and for cost, since storage architecture decisions often lock an organization into a particular cost structure for years. This

chapter equips you to participate meaningfully in that decision, whether or not you are the one implementing it technically.

This chapter closes Part II by completing a natural pair with Chapter 4: having addressed where data comes from, it now addresses where that data actually lives once it arrives. The two decisions are more connected than they first appear, since the collection method chosen in Chapter 4 often constrains which storage systems are practically available, and a storage decision made without reference to how the data was collected can create friction that neither chapter's material alone would predict.

5.3 Historical Background

Relational databases dominated organizational data storage from the 1970s through the early 2000s, offering strong consistency guarantees and a mature query language, SQL, well-suited to the structured, transactional data most organizations held. The explosive growth of unstructured and semi-structured data from the 2000s onward, alongside the sheer scale of data generated by web-scale applications, exposed real limitations in the relational model for certain use cases, driving the development of NoSQL databases and, later, data lakes capable of storing raw, unstructured data at massive scale without requiring a predefined schema at the point of storage.

Cloud computing's maturation over the past decade has further reshaped storage architecture decisions, shifting many organizations away from owning and maintaining physical storage infrastructure toward consuming storage as an on-demand service. This shift has lowered the barrier to entry for smaller organizations, including many in the Philippines, that previously could not justify the capital investment required by traditional on-premises storage infrastructure.

The relational-to-NoSQL transition is sometimes mischaracterized as one technology simply replacing another, but the more accurate historical reading is additive rather than substitutive. NoSQL databases did not render relational databases obsolete; they expanded the available toolkit to cover use cases relational databases were never well suited for in the first place, while relational databases have remained, and remain today, the standard choice for exactly the structured, consistency-critical use cases they were originally designed to serve. This additive pattern, new tools expanding rather than replacing the toolkit, recurs throughout this book's later treatment of processing and analytics technology as well.

5.4 Definitions

KEY TERMS

Relational Database: A database that organizes data into structured tables with predefined schemas, related through keys, and queried using SQL.

NoSQL Database: A database designed for flexible, schema-less or semi-structured data, often optimized for scale and speed over strict relational consistency.

Data Warehouse: A centralized repository optimized for storing structured, processed data specifically for analysis and reporting rather than transactional operations.

Data Lake: A storage repository that holds vast amounts of raw data, structured and unstructured, in its native format until it is needed.

Cloud Storage: Data storage maintained, operated, and billed by a third-party provider, accessed over the internet rather than on organization-owned physical infrastructure.

5.5 Core Concepts and Frameworks




Figure 5.1. Centralized vs. Distributed Storage.

5.5.1 Types of Data Storage Systems

Relational databases remain the standard choice for structured, transactional data requiring strong consistency, such as financial records or inventory systems. NoSQL databases suit applications that require flexible schemas or massive horizontal scaling, such as social media

platforms or IoT sensor networks. Data warehouses serve analytical rather than transactional needs, consolidating cleaned, structured data from multiple source systems specifically to support reporting and business intelligence. Data lakes serve as a landing zone for raw data of any type, structured or unstructured, deferring schema decisions until the data is actually needed for a specific analytical purpose.

5.5.2 Centralized vs. Distributed Architectures

Centralized architectures store data in a single location or tightly coupled system, offering simpler management and stronger consistency guarantees but creating a single point of failure and potential performance bottleneck as scale increases. Distributed architectures spread data across multiple physical or logical locations, offering greater scalability and resilience against a single point of failure, at the cost of increased complexity in maintaining consistency across the distributed copies. Most large modern organizations operate a hybrid of both, using centralized systems for core transactional data and distributed systems for analytical or high-volume data.

5.5.3 On-Premises vs. Cloud Storage

On-premises storage, owned and maintained by the organization, offers maximum control over security and compliance configurations but requires a significant upfront capital investment and ongoing technical maintenance capacity. Cloud storage shifts this burden to a third-party provider, offering lower upfront cost and rapid scalability, but requires careful vendor evaluation, particularly around where data is physically stored, since this has direct implications for which jurisdiction's data protection laws apply.

5.6 Table 5.1, Choosing a Storage System

Storage System	Best Suited For	Key Limitation
Relational Database	Structured, transactional data needing strong consistency	Less flexible for unstructured or rapidly changing schemas
NoSQL Database	Flexible-schema, high-scale applications	Weaker consistency guarantees than relational systems
Data Warehouse	Structured data for reporting and analytics	Requires cleaned, processed data; not for raw storage

Storage System	Best Suited For	Key Limitation
Data Lake	Raw data of any type, purpose deferred	Risk of becoming an unmanaged data swamp without governance

Table 5.1. Matching storage systems to data needs.

5.7 Government Policy and Data Security in Storage

Storage architecture decisions carry direct Data Privacy Act implications, particularly the requirement to implement reasonable and appropriate security measures proportionate to the sensitivity of the data stored. Cloud storage arrangements additionally raise data sovereignty questions, since the Data Privacy Act's protections are most straightforwardly enforceable when personal data of Philippine residents remains within jurisdictions with comparable protections, a consideration organizations must weigh explicitly when selecting cloud providers and specific data center regions.

5.8 Philippine Context and Industry Practice

PHILIPPINE CONTEXT

Philippine organizations are increasingly adopting cloud storage for its lower upfront costs and scalability, particularly for small and medium enterprises that could not previously justify on-premises infrastructure investment. This shift has made cloud vendor data-residency options; several major providers now offer Southeast Asia-region data centers, an increasingly practical and relevant consideration for Philippine organizations balancing cost, performance, and Data Privacy Act compliance.

Internet connectivity reliability outside major Philippine urban centers remains a practical factor that cloud storage decisions must account for, since an organization's day-to-day operations can be meaningfully disrupted by connectivity gaps in ways an on-premises system, whatever its other limitations, would not. Organizations serving provincial or rural operations sometimes adopt a hybrid approach specifically for this reason, keeping operationally critical data available locally while using cloud storage for analytical, archival, or less time-sensitive data, a pattern this chapter's centralized-

versus-distributed framework helps clarify even though it was not originally designed with connectivity reliability specifically in mind.

5.9 Case Study and Best Practices

CASE STUDY

Illustrative composite case.

A Philippine logistics company initially stored all data in a single relational database, structured shipment records alongside unstructured delivery photos and customer communication logs, never designed for unstructured content, resulting in degraded performance and awkward workarounds for storing images as encoded text fields. A redesigned architecture separating structured shipment data into a relational database and unstructured content into cloud object storage, connected through a unified access layer, resolved the performance issues and allowed each data type to be stored using tools actually designed for it.

The performance improvement was substantial and immediate; query response times for shipment lookups improved severalfold once the database was no longer burdened with encoded image data, which it was never designed to handle efficiently. The migration also required the company to rebuild a number of internal reports that had been written directly against the old single-database structure. This migration cost, while real, was ultimately smaller than the ongoing performance cost the company would have continued to absorb indefinitely had the storage architecture problem gone unaddressed, illustrating a pattern this book returns to repeatedly: a well-justified architectural correction usually costs less over time than tolerating a poor fit indefinitely.

- Best Practice: Match storage system to data type and use case explicitly, rather than forcing all organizational data into a single system chosen for convenience.

- Best Practice: Evaluate cloud vendor data-residency options against Data Privacy Act obligations before finalizing a cloud storage decision, not after data has already been migrated.

5.10 Summary and Key Takeaways

Data storage systems, relational databases, NoSQL databases, data warehouses, and data lakes each suit different data types and use cases, and architecture decisions between centralized and distributed, on-premises and cloud, carry real consequences for cost, performance, security, and compliance. This chapter closes Part II; Part III now turns to ensuring the quality of the data these systems hold.

- Common Pitfall: Forcing all organizational data into a single storage system chosen for convenience, regardless of whether it fits the data's actual type and use case.
- Common Pitfall: Selecting cloud storage without evaluating data-residency options against Data Privacy Act obligations before migration.
- Common Pitfall: Allowing a data lake to become an unmanaged data swamp by storing raw data without any governance or documentation.

KEY TAKEAWAYS

- Relational databases suit structured, transactional data; data lakes suit raw data of any type awaiting future purpose.
- Distributed architectures offer scalability and resilience at the cost of consistency complexity, relative to centralized systems.
- Cloud storage lowers upfront cost but requires careful evaluation of data-residency and Data Privacy Act implications.

- Storage architecture should be chosen deliberately per data type, not defaulted to a single system for organizational convenience.

5.11 Key Terms, Reflection, and Discussion

Relational Database. NoSQL Database. Data Warehouse. Data Lake. Cloud Storage. Data Sovereignty.

- Reflection: What storage system would best suit your Data Project Build dataset, and why?
- Discussion: Should Philippine organizations be required by law to store personal data of Philippine residents within the country, or is cross-border cloud storage with adequate safeguards sufficient?
- Discussion: Table 5.1 warns that a data lake can become an unmanaged data swamp without governance. What specific governance practice from Chapter 2 would you apply first to prevent this?

5.12 Activity and Data Project Build, Step 5

DATA PROJECT BUILD, STEP 5

Recommend a storage system and architecture (centralized or distributed, on-premises or cloud) for your Data Project Build dataset, justified against Table 5.1's criteria and any Data Privacy Act considerations relevant to your data.

5.13 References

- Silberschatz, A., Korth, H. F., & Sudarshan, S. (2020). Database system concepts (7th ed.). McGraw-Hill.
- Armbrust, M., et al. (2010). A view of cloud computing. Communications of the ACM, 53(4), 50 to 58.
- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- Apache Hadoop. Apache Hadoop documentation (latest edition).

End of Chapter 5. Continue to Chapter 6: Dimensions and Challenges of Data Quality

PART III

DATA QUALITY AND INTEGRITY

Part III asks whether collected and stored data can actually be trusted. It introduces the formal dimensions along which data quality is measured, and the cleansing and validation techniques used to repair and prevent the defects that erode it.

DIMENSIONS AND CHALLENGES OF DATA QUALITY

Measuring Whether Data Is Actually Fit to Use

“Data quality is not a property data has or lacks. It is a measure of how well data serves the purpose someone actually needs it for.”

6.0 Chapter Overview

With data now collected (Part II) and given a home in an appropriate storage system, Part III turns to a question that determines whether any of that data can actually be trusted: Is it good enough to use? This chapter introduces the formal dimensions along which data quality is measured and the common challenges that erode it, setting the stage for Chapter 7's treatment of the cleansing and validation techniques used to repair it.

6.1 Learning Outcomes

- LO 6.1 Explain the concept of data quality and its dependence on context of use.
- LO 6.2 Identify and apply the major dimensions of data quality.
- LO 6.3 Diagnose common data quality challenges and their organizational sources.

6.2 Introduction

Ask three different people in the same organization whether a given dataset is high quality, and you will often get three different answers, not because any of them is wrong, but because data quality is not a single, objective property a dataset either has or lacks. It is, in the most widely used definition, fitness for use: the degree to which data successfully serves the purpose of the person or process consuming it. A customer address dataset missing postal codes may be perfectly suited to a marketing email campaign and completely unsuitable for a shipping logistics system, using the same underlying data.

This context-dependence is precisely why data quality requires formal, measurable dimensions rather than a single vague judgment of good or bad. This chapter introduces those dimensions, giving you a vocabulary for diagnosing exactly which quality property a given dataset is lacking and for which specific use case, rather than settling for the blunt, often unhelpful verdict that the data is simply bad.

Part III, this chapter, and Chapter 7 sit at a hinge point in the book's overall structure. Everything before it, mindset, governance, collection, storage, has produced data that now exists somewhere in an organization's systems. Everything after it, processing, security, and analytics, assumes that the data is trustworthy enough to build on. This Part is where that assumption is actually tested and, where necessary, repaired, rather than hoped for.

6.3 Historical Background

Formal data quality theory was substantially developed by researchers such as Richard Wang and Diane Strong in the 1990s, who reframed data quality from a narrow technical accuracy question into a broader, multidimensional concept explicitly grounded in fitness for use. This reframing mattered because it shifted organizational attention away from purely technical error-correction toward a wider set of questions, including whether data was complete, timely, and consistent, that technical accuracy checks alone would never surface.

Later synthesis work, notably by David Loshin and by database theorists Elmasri and Navathe, consolidated a growing, sometimes inconsistent, academic vocabulary into the more standardized dimension sets, accuracy, completeness, consistency, timeliness, and related measures that this chapter and most contemporary data quality practice now use. Even so, no single set of dimensions is universally agreed upon among data quality professionals, and organizations often adapt the core set to their own specific context.

Wang and Strong's fitness-for-use reframing is worth revisiting, as it explains why data quality theory took longer to mature than data storage theory. Storage and structure are properties a database administrator can verify largely by inspecting the system itself. Fitness for use, by definition, cannot be verified by inspecting the data alone; it requires understanding the purpose the data is meant to serve, which varies by consumer and context in ways that resisted the kind of clean, universal formalization that Codd's relational model achieved for storage roughly two decades earlier.

6.4 Definitions

KEY TERMS

Data Quality: The degree to which data is fit for the person or process using it, measured along multiple distinct dimensions.

Data Quality Dimension: A specific, measurable aspect of data quality, such as accuracy or completeness, that can be assessed independently of the others.

Data Integrity: The overall accuracy, consistency, and reliability of data maintained across its lifecycle, often used to describe the combined effect of multiple quality dimensions.

Data Profiling: The systematic analysis of a dataset to understand its structure, content, and quality characteristics before use or cleansing.

6.5 Core Concepts and Frameworks



Figure 6.1. Core Data Quality Dimensions.

6.5.1 The Core Data Quality Dimensions

Several dimensions recur across most data quality frameworks. Accuracy refers to the degree of closeness of data values to real, true values. Completeness is the degree to which all required data values are present, with no missing fields where a value should exist. Consistency

is the degree to which data values agree with each other across records, files, or points in time, rather than contradicting one another. Timeliness is the degree to which data is available within an appropriate period relative to when it was created or is needed. Uniqueness is the degree to which records occur only once in a dataset, without unintended duplication. Validity is the degree to which data values comply with defined rules or formats, such as a postal code matching an expected pattern.

6.5.2 Additional Dimensions: Currency, Traceability, and Availability

Beyond the core dimensions, several additional measures matter in specific contexts. Currency is the degree to which data values remain up to date, distinct from timeliness, which concerns whether the value itself still reflects current reality, not merely how quickly it was made available. Traceability is the degree to which data lineage, its origin, and every transformation it has undergone can be established and audited. Availability is the degree to which data can be consulted or retrieved by consumers or processes that need them, and it can be compromised by poor access design even when the underlying data is accurate and complete.

6.5.3 Data Quality Is Multidimensional and Trade-Off-Prone

A critical, easily missed point is that these dimensions do not always move together, and improving one can sometimes come at the cost of another. Prioritizing timeliness, releasing data as fast as possible, can reduce the time available for accuracy validation. Prioritizing completeness, requiring every field to be filled before a record is accepted, can slow collection and reduce timeliness. Effective data quality management, developed further in Chapter 7, requires deciding which dimensions matter most for a given use case rather than assuming all dimensions can be maximized simultaneously without cost.

6.6 Table 6.1, Core Data Quality Dimensions

Dimension	Definition	Example Check
Accuracy	Closeness of data values to real values	Does the recorded address match the actual address?
Completeness	Presence of all required values	Are any required fields left blank?

Dimension	Definition	Example Check
Consistency	Agreement of values across records or files	Does a customer's status match across all systems?
Timeliness	Availability within an appropriate time period	Is this month's sales data available before the report is due?
Uniqueness	Records occur only once	Are there duplicate customer entries?
Validity	Compliance with defined rules or formats	Does the email field contain a properly formatted address?

Table 6.1. The core data quality dimensions and how each is assessed.

6.7 Common Data Quality Challenges

Several recurring organizational patterns erode data quality across most industries. Manual data entry remains one of the most persistent sources of accuracy and consistency errors, since human transcription introduces typos, inconsistent formatting, and interpretation differences that automated collection avoids. Data silos, where different departments maintain separate, disconnected copies of overlapping data, produce consistency failures as those copies inevitably drift apart over time without a synchronization mechanism. System migrations and integrations frequently introduce quality defects when data is mapped between systems with different schemas or validation rules, sometimes silently truncating or misinterpreting values.

6.8 Government Policy and Philippine Context

Data quality is not merely an operational concern in the Philippine regulatory context. The Data Privacy Act's principle of data quality explicitly requires that personal data be accurate and, where necessary, kept up to date, giving data quality direct legal weight beyond its business value. The National Privacy Commission has treated significant, uncorrected data quality failures affecting individuals, such as persistently inaccurate records used in decision-making, as a compliance concern in its own right.

PHILIPPINE CONTEXT

Philippine organizations transitioning from paper-based or fragmented legacy systems to unified digital platforms frequently discover extensive data quality issues only during the transition, including accumulated duplicate customer records from years of manual entry, inconsistent address formats, and outdated contact information. This makes data quality assessment a near-universal, immediate first step for Philippine organizations undertaking digital transformation, not an optional refinement to address later.

Philippine address data poses a specific, recurring quality challenge worth naming directly: the absence of a single standardized national addressing system, combined with informal or landmark-based addressing common in many communities, means address fields that would be considered well-formed and valid in many other national contexts routinely fail simple validity checks when applied to genuine Philippine addresses. Organizations that import overly rigid address validation logic from a foreign software vendor without adapting it to this reality often reject or mangle a meaningful share of otherwise perfectly usable, accurate customer data.

6.9 Case Study and Best Practices

CASE STUDY

Illustrative composite case.

A Philippine microfinance institution's loan approval system relied on customer income data that, during a data profiling review, was found to be complete for 98% of records but accurate for an estimated 60%, since loan officers under time pressure frequently entered rounded, best-guess income estimates rather than verified figures. The completeness dimension alone had masked a serious accuracy problem that surfaced only when the institution's data quality review explicitly measured accuracy as a separate dimension rather than treating a fully populated field as evidence of good data.

The institution's response involved more than a one-time cleansing pass on existing records, since the underlying cause was procedural, loan officers lacked time, and, in

some cases, there was no convenient method to verify income during the application process. A revised workflow adding a lightweight, mobile-friendly income verification step at the point of entry addressed the root cause directly, illustrating why the diagnosis stage of a quality investigation, tracing a defect back to why it occurred, matters as much as detecting the defect itself; a purely corrective cleansing effort without this procedural fix would have required repeating the same correction indefinitely as new, equally rushed applications continued to enter the system.

- Best Practice: Measure data quality along multiple explicit dimensions rather than relying on a single, vague quality judgment.
- Best Practice: Conduct data profiling before assuming a dataset's quality; a fully populated field can still be inaccurate.

6.10 Extended Application: Diagnosing a Quality Problem by Dimension

Consider a concrete diagnostic scenario. A university registrar's office notices that its graduation eligibility reports frequently differ from the records of individual academic advisors. Applying the dimensions from Table 6.1 systematically, rather than simply declaring the data wrong, quickly narrows the actual problem: completeness checks show no missing fields, accuracy spot-checks against original transcripts show values match source documents, but a consistency check reveals that the registrar's system and the advising system apply different rules for when a repeated course replaces a failing grade in the cumulative average calculation. The problem is not that data is inaccurate or incomplete; it is a consistency failure rooted in two systems implementing the same underlying business rule differently.

This diagnostic habit, checking each dimension independently rather than assuming a single cause, transfers directly to your own Data Project Build. When you encounter a quality problem in your dataset in the chapters ahead, resist the temptation to label it broadly as bad data and instead identify precisely which dimension, or combination of dimensions, is actually failing, since the appropriate cleansing technique in Chapter 7 differs substantially depending on the answer.

6.10.1 Frequently Asked Questions

- Can data be 100% high quality across every dimension simultaneously? Rarely in practice, and rarely necessary; the goal is fitness for the specific use case, not maximal quality on every dimension regardless of cost or need.
- Which dimension matters most? It depends entirely on use case; a real-time fraud detection system may prioritize timeliness over completeness, while a financial audit may prioritize accuracy and traceability above all else.
- Is data profiling a one-time activity? No, data quality degrades over time as systems, processes, and business rules change, so profiling should be repeated periodically rather than treated as a single initial assessment.

6.11 Summary and Key Takeaways

Data quality is fitness for use, measured along multiple distinct dimensions, accuracy, completeness, consistency, timeliness, uniqueness, and validity, among them, rather than a single blunt judgment of good or bad data. Common organizational challenges, such as manual entry, data silos, and system migrations, erode these dimensions in identifiable, diagnosable ways, setting the stage for the cleansing and validation techniques Chapter 7 develops to address them.

- Common Pitfall: Treating a fully populated field as evidence of good data, without separately verifying accuracy.
- Common Pitfall: Assessing data quality once at the start of a project and never revisiting it as systems and processes change.
- Common Pitfall: Attempting to maximize every quality dimension equally, rather than prioritizing the dimensions that matter most for the specific use case.

KEY TAKEAWAYS

- Data quality is fitness for use, dependent on context, not a single objective property.
- Core dimensions include accuracy, completeness, consistency, timeliness, uniqueness, and validity.
- Quality dimensions can trade off against each other; maximizing every dimension simultaneously is rarely feasible or necessary.
- The Data Privacy Act's data quality principle gives accuracy and currency direct legal weight for personal data in the Philippines.

6.12 Key Terms, Reflection, and Discussion

Data Quality. Data Quality Dimension. Data Integrity. Data Profiling. Accuracy. Completeness. Consistency. Timeliness. Uniqueness. Validity.

- Reflection: Profile your Data Project Build dataset informally against the six dimensions in Table 6.1. Which dimension is strongest, and which is weakest?
- Discussion: Should regulatory frameworks like the Data Privacy Act specify measurable data quality standards, or is a general, accurate, and up-to-date requirement sufficient?
- Discussion: Can an organization ever justify accepting lower data quality on one dimension in exchange for a gain on another? Give an example.

- Discussion: The microfinance case study shows completeness masking a real accuracy problem. What other pairs of dimensions do you think are most likely to mask each other in practice?

6.13 Activity and Data Project Build, Step 6

DATA PROJECT BUILD, STEP 6

Conduct a data quality assessment of your Data Project Build dataset. Rate it Low, Medium, or High on each of the six dimensions in Table 6.1, with one sentence of justification per dimension, and identify which dimension most urgently needs attention.

6.14 Chapter Self-Check

CHAPTER SELF-CHECK

- 1. The most widely used definition of data quality is:**
 - a) Zero missing values
 - b) Fitness for use
 - c) Maximum storage efficiency
 - d) Absence of duplicate records
- 2. A dataset where a customer's status disagrees between two systems is failing, which dimension?**
 - a) Completeness
 - b) Timeliness
 - c) Consistency
 - d) Uniqueness
- 3. Data profiling is best described as:**
 - a) A one-time cleansing step
 - b) Systematic analysis of a dataset's structure and quality before use
 - c) A method for encrypting data
 - d) A backup procedure
- 4. Which of the following is NOT one of the core data quality dimensions discussed in this chapter?**
 - a) Accuracy
 - b) Completeness
 - c) Profitability
 - d) Validity
- 5. Data quality dimensions can trade off against each other, meaning:**

a) They are unrelated to each other. b) Improving one dimension can sometimes reduce another. c) All dimensions always improve together. d) Only accuracy matters in practice

6.15 References

- Wang, R. Y., & Strong, D. M. (1996). Beyond accuracy: What data quality means to data consumers. *Journal of Management Information Systems*, 12(4), 5 to 33.
- Loshin, D. (2001). *Enterprise knowledge management: The data quality approach*. Morgan Kaufmann.
- Elmasri, R., & Navathe, S. B. (2016). *Fundamentals of database systems* (7th ed.). Pearson.
- Republic Act No. 10173 (2012). *Data Privacy Act of 2012*.

End of Chapter 6. Continue to Chapter 7: Data Cleansing and Validation Techniques

DATA CLEANSING AND VALIDATION TECHNIQUES

Repairing and Verifying Data Before It Is Trusted

“Cleansing does not create trustworthy data. It removes the specific, identified reasons a dataset could not yet be trusted.”

7.0 Chapter Overview

Chapter 6 taught you to diagnose data quality problems along specific dimensions. This chapter closes Part III by teaching you the practical techniques used to fix what that diagnosis surfaces: cleansing and validation. Where Chapter 6 asked whether data is fit for use, this chapter asks what to actually do once you have found that, along some dimension, it is not.

7.1 Learning Outcomes

- LO 7.1 Apply data cleansing techniques to address common quality defects.
- LO 7.2 Apply data validation techniques to prevent quality defects at the point of entry.
- LO 7.3 Evaluate the trade-offs involved in automated versus manual cleansing approaches.

7.2 Introduction

Cleansing and validation are complementary rather than interchangeable practices, and confusing them is a common early mistake. Validation is preventive: it prevents bad data from entering a system in the first place by applying rules at the point of collection or entry. Cleansing is corrective: it identifies and repairs quality defects already present in data that has already been collected and stored. A mature data quality program uses both, since no validation rule set can prevent every possible defect, and no amount of downstream cleansing fully substitutes for preventing errors at the source.

This chapter closes Part III by turning the diagnostic vocabulary in Chapter 6 into a practical repair technique. Where Chapter 6 taught you to say precisely which dimension a

dataset is failing, this chapter teaches you what actually to do about it, and closing this gap between diagnosis and action is exactly what separates a data quality assessment that produces a report from one that produces genuinely improved data.

7.3 Historical Background

Early data cleansing practice was almost entirely manual, involving staff reviewing printed reports for obvious errors, a labor-intensive approach that scaled poorly as data volumes grew. The development of rule-based and, later, statistical and machine-learning-assisted cleansing tools from the 1990s onward allowed organizations to automate the detection of many common defect patterns, duplicate records, inconsistent formatting, and outlier values at a scale manual review could never match.

Van den Broeck and colleagues' influential 2005 taxonomy of data cleansing, distinguishing detection, diagnosis, and editing as distinct stages, remains widely referenced in both academic and industry practice. This chapter's structure reflects that same underlying logic: first find defects systematically, then understand their cause, then correct them appropriately rather than through a single undifferentiated cleaning pass.

Notably, Van den Broeck's taxonomy originated in medical and public health research rather than in commercial data management specifically, developed to address data abnormalities in clinical and epidemiological datasets where an uncorrected error could directly affect a health-related conclusion. This origin helps explain the taxonomy's insistence on a diagnosis stage distinct from mere detection and correction; medical researchers could not treat cleansing as a purely mechanical exercise, since understanding why a data abnormality occurred was often itself clinically informative, a rigor this chapter's application to organizational data management deliberately preserves.

7.4 Definitions

KEY TERMS

Data cleansing: The process of detecting and correcting, or removing, corrupt, inaccurate, incomplete, or duplicate records from a dataset.

Data Validation: The process of checking data against defined rules at the point of entry or processing to prevent invalid data from being accepted.

Deduplication: the process of identifying and merging or removing duplicate records that refer to the same underlying entity.

Standardization: the process of converting data into a consistent format, such as unifying date formats or address structures, across a dataset.

Outlier Detection: The process of identifying data values that deviate significantly from the expected pattern, which may indicate an error or a genuine, meaningful anomaly.

7.5 Core Concepts and Frameworks




Figure 7.1. The Cleansing Cycle: Detect, Diagnose, Edit.

7.5.1 The Detect, Diagnose, Edit, and Cleanse Cycle

Effective cleansing follows a disciplined three-stage cycle. Detection systematically identifies where a dataset likely violates a quality rule, such as a duplicate record, an out-of-range value, or a missing required field, typically through automated profiling tools rather than manual inspection. Diagnosis investigates why the defect occurred, tracing it back to a root cause such as a flawed collection form or a system integration error, since the same correction approach is not appropriate for every cause. Editing applies the actual correction, whether

automated (e.g., standardizing a date format programmatically) or manual (e.g., a human reviewer resolving an ambiguous duplicate match).

7.5.2 Common Cleansing Techniques

Deduplication identifies records that likely refer to the same real-world entity despite minor differences in how they were recorded, such as two near-identical name and address entries, and merges or removes the redundant copies. Standardization converts inconsistent formats and dates recorded in multiple different patterns within the same field into a single consistent format. Missing value treatment addresses incomplete records by removing them, flagging them for manual follow-up, or, where appropriate, imputing a reasonable estimate. Outlier detection flags values falling far outside an expected range for human review, since an outlier may represent either an error worth correcting or a genuine, significant anomaly worth investigating further.

7.5.3 Validation Rules as Prevention

Validation applies rules at the point of data entry to reject or flag invalid values before they are stored. Format validation checks that a value matches an expected pattern, such as a valid email address structure. Range validation checks that a numeric value falls within a plausible range, rejecting values that are implausible or negative. Referential validation checks that a value correctly references an existing related record, such as an order referencing a customer ID that actually exists in the customer table. Well-designed validation prevents a substantial share of the defects that would otherwise require costly downstream cleansing.

7.6 Table 7.1, Cleansing Techniques and What They Fix

Technique	Primary Dimension Addressed	Example
Deduplication	Uniqueness	Merging two records for the same customer
Standardization	Consistency	Unifying date formats across a dataset
Missing Value Treatment	Completeness	Flagging or imputing a blank required field

Technique	Primary Dimension Addressed	Example
Outlier Detection	Accuracy	Reviewing an implausible transaction amount
Validation Rules	Prevention of multiple dimensions at entry	Rejecting an improperly formatted email at signup

Table 7.1. Matching cleansing techniques to the quality dimension they address.

7.7 Government Policy and Professional Standards

Systematic cleansing and validation practices directly support compliance with the Data Privacy Act, as the Act's data quality principle requires organizations to correct inaccurate personal data, particularly upon a data subject's request under their right to rectification. Organizations without a documented cleansing process struggle to demonstrate, if challenged by the National Privacy Commission, that they have reasonable procedures in place to maintain accurate personal data over time.

7.8 Philippine Context and Industry Practice

PHILIPPINE CONTEXT

Philippine organizations consolidating customer data across multiple legacy systems, a common scenario during digital transformation initiatives, frequently encounter severe deduplication challenges rooted in inconsistent Filipino name formatting, differing conventions for middle names, multiple surnames, and nicknames used interchangeably across different systems and points of collection. Effective deduplication in this context often requires fuzzy matching techniques tuned to Filipino naming conventions, rather than exact-match logic that works well in contexts with more standardized naming patterns.

A related complication specific to Philippine deduplication efforts involves married women's surnames, which may appear as a maiden name in older records and a married name in newer ones, sometimes with no clear internal system flag indicating they refer to the same individual. Deduplication logic that accounts only for straightforward misspellings or formatting variations, without specifically addressing this common

surname-change pattern, will systematically miss a meaningful share of genuine duplicates in many Philippine customer databases.

7.9 Case Study and Best Practices

CASE STUDY

Illustrative composite case.

A Philippine telecommunications provider's customer database contained an estimated twelve percent duplicate records, primarily due to customers registering multiple times across different service channels, retail stores, phone hotline, and mobile app, without a unified identity check. An automated deduplication pass using fuzzy name and address matching reduced duplicates to under two percent. Still, the remaining cases required manual review, since several apparent duplicates were, on investigation, genuinely different family members sharing a household address, illustrating why fully automated cleansing without any human review step risks incorrectly merging distinct real-world entities.

The provider also implemented validation rules at each of its three registration channels following the cleansing exercise, requiring a mobile number lookup against existing records before a new account could be created, directly addressing the root cause identified earlier in the detect, diagnose, and edit cycle. New duplicate registration rates fell by a comparable margin to the original cleansing pass within the first quarter after implementation, confirming Chapter 7's central claim in practice: validation applied at the point of entry prevents the same defect that cleansing can only ever repair after the fact.

- Best Practice: Combine automated detection with human review for ambiguous cases, particularly deduplication, where an incorrect automated merge can itself become a new data quality defect.

- Best Practice: Implement validation rules at the point of entry wherever possible; preventing a defect is almost always cheaper than cleansing it later.

7.10 Summary and Key Takeaways

Data cleansing corrects quality defects already present in stored data, following a detect, diagnose, and edit cycle. In contrast, data validation prevents defects at the point of entry through format, range, and referential rules. Both are necessary; validation alone cannot catch every possible defect, and cleansing alone cannot prevent those defects from recurring. This chapter closes Part III; Part IV now turns to processing and transforming data once its quality has been established.

- Common Pitfall: Relying entirely on automated deduplication without human review, risking incorrect merges of genuinely distinct records.
- Common Pitfall: Cleansing data repeatedly without ever implementing validation rules that would prevent the same defects from recurring.
- Common Pitfall: Correcting an outlier automatically without investigating whether it represents a genuine, meaningful anomaly rather than an error.

KEY TAKEAWAYS

- Validation is preventive, applied at entry; cleansing is corrective, applied to data already collected.
- The detect, diagnose, edit cycle structures effectively cleans rather than a single undifferentiated cleaning pass.
- Deduplication, standardization, missing value treatment, and outlier detection each address a specific quality dimension.

- Fully automated cleansing without human review risks introducing new errors, particularly in ambiguous deduplication cases.

7.11 Key Terms, Reflection, and Discussion

Data Cleansing. Data Validation. Deduplication. Standardization. Outlier Detection. Detect, Diagnose, Edit Cycle.

- Reflection: Based on your Chapter 6 quality assessment, which cleansing technique from Table 7.1 does your Data Project Build dataset most need?
- Discussion: When should an outlier be corrected as an error, and when should it be preserved as a genuine, meaningful anomaly? How would you decide?
- Discussion: This chapter argues that validation is almost always cheaper than cleansing. Can you think of a case where investing heavily in upfront validation would not be worth the cost?

7.12 Activity and Data Project Build, Step 7

DATA PROJECT BUILD, STEP 7

Design a cleansing and validation plan for your Data Project Build dataset. Identify at least two specific defects (real or anticipated) and specify which technique from Table 7.1 would address each, and draft two validation rules that would have prevented one of those defects at the point of entry.

7.13 Chapter Self-Check

CHAPTER SELF-CHECK

1. The key difference between validation and cleansing is:

- a) They are the same technique. b) Validation is preventive at entry; cleansing is corrective on stored data. c) Validation only applies to numbers. d) Cleansing only applies to duplicate records

2. The detect, diagnose, and edit cycle was formalized by:

- a) Codd (1970) b) Van den Broeck et al. (2005) c) Wang and Strong (1996)
d) Porter (1979)

3. Which technique specifically addresses duplicate records?

- a) Standardization b) Deduplication c) Range validation d) Outlier detection

4. Referential validation checks that:

- a) A value falls within a plausible numeric range. b) A value correctly references an existing related record. c) A date is formatted consistently. d) A field is not empty

5. Why is human review often still needed alongside automated deduplication?

- a) Automation is always wrong. b) Some apparent duplicates are genuinely distinct entities. c) Automated tools cannot process text. d) Deduplication is not a real technique

7.14 References

- Van den Broeck, J., Cunningham, S. A., Eeckels, R., & Herbst, K. (2005). Data cleaning: Detecting, diagnosing, and editing data abnormalities. *PLoS Medicine*, 2(10), e267.
- Deoras, S. (2018). 10 best data cleaning tools to get the most out of your data. *Analytics India Magazine*.
- Republic Act No. 10173 (2012). Data Privacy Act of 2012.

- Yellowfin BI. Data quality best practice guide.

End of Chapter 7. Continue to Chapter 8: ETL Processes and Data Transformation Techniques

PART IV

DATA PROCESSING AND TRANSFORMATION

Part IV covers how data moves from its source systems into a consistent, usable form through the ETL process, and how organizations automate that movement reliably at scale rather than depending on manual, error-prone execution.

ETL PROCESSES AND DATA TRANSFORMATION TECHNIQUES

Turning Raw, Trustworthy Data into Data Ready for Use

“Extraction moves data. Transformation is what actually makes it useful.”

8.0 Chapter Overview

With data collected, stored, and brought to an acceptable quality standard in Parts II and III, Part IV turns to processing: moving data from its source systems into the form and location where it can actually be used, and transforming it along the way. This chapter introduces the Extract, Transform, Load process, the backbone of most organizational data pipelines, and the transformation techniques applied within it.

8.1 Learning Outcomes

- LO 8.1 Explain and apply the Extract, Transform, Load (ETL) process.
- LO 8.2 Differentiate and apply appropriate data transformation techniques for analysis.
- LO 8.3 Compare data loading strategies and their appropriate use cases.

8.2 Introduction

Data rarely lives in its raw, collected form where it needs to be for analysis or reporting. A retail chain's sales data might originate across dozens of individual point-of-sale systems, each with its own format and update schedule, while the analytics team needs a single, consistent, up-to-date view. ETL (Extract, Transform, Load) is the standard process for bridging this gap: pulling data from its sources, reshaping it into a consistent and usable form, and depositing it into the destination system where it will be consumed.

Part IV opens with this chapter for a specific reason: processing sits between the quality assurance work of Part III and the security and analytics work of Parts V and VI, and it is easy to treat it as purely mechanical, a plumbing problem rather than a discipline with its own real

judgment calls. This chapter and Chapter 9 argue otherwise: how data is extracted, transformed, and loaded shapes what that data can later be trusted to mean, in ways every downstream analyst inherits, whether or not they were involved in the original pipeline design.

8.3 Historical Background

ETL emerged as a formalized practice alongside the data warehouse concept in the 1980s and 1990s, as organizations increasingly sought to consolidate operational data from multiple source systems into a single system optimized for reporting and analysis, rather than analyzing each operational system directly and risking both performance degradation of live systems and inconsistent, uncoordinated analysis. Dedicated ETL software tools matured through the 1990s and 2000s to automate what had initially been largely custom-coded, one-off integration scripts.

More recently, the rise of cloud data warehouses with far greater processing power has popularized an alternative sequencing, ELT, Extract, Load, Transform, in which raw data is loaded into the destination system first and transformed afterward using that system's own processing power, rather than transforming data in a separate staging step before loading. Both approaches remain in active use, and the choice between them depends on the specific technical environment and data volume an organization is working with.

The original motivation for ETL's transform-before-load sequencing was largely a practical constraint of the era: early data warehouse destination systems had limited processing power relative to dedicated transformation servers, making it more efficient to reshape data before loading than to have the warehouse itself perform that work. As destination systems grew dramatically more powerful, particularly with cloud data warehouses, this original constraint weakened considerably, which is precisely why ELT became a genuinely viable sequencing alternative rather than merely a theoretical one; the technology finally caught up with a sequencing option that had always been logically possible but was not previously practical.

8.4 Definitions

KEY TERMS

ETL (Extract, Transform, Load) is the process of extracting data from source systems, transforming it into a consistent, usable format, and loading it into a destination system.

ELT (Extract, Load, Transform) is a variant process that loads raw data into the destination system before transforming it, leveraging the destination system's own processing power.

Data Pipeline: The automated sequence of processes that moves data from source to destination, typically including extraction, transformation, and loading steps.

Data Transformation: The process of converting data from its original format or structure into a different format or structure suited to its intended use.

8.5 Core Concepts and Frameworks




Figure 8.1. The ETL Pipeline.

8.5.1 The Three Stages of ETL

Extraction pulls data from one or more source systems, which may be structured databases, unstructured file stores, or external APIs, and is typically the stage most sensitive to source system availability and format changes. Transformation reshapes the extracted data by applying the techniques described in Section 8.5.2 to conform to the destination system's

schema, format, and business rules. Loading writes the transformed data into its destination, whether a data warehouse for analytics, a data lake for broader storage, or another operational system, using one of the loading strategies covered in Section 8.5.3.

8.5.2 Core Data Transformation Techniques

Several transformation techniques recur across most ETL pipelines. Filtering removes records or fields not needed for the destination purpose, reducing volume and, often, exposure of unnecessary sensitive data. Aggregation summarizes detailed records into higher-level summaries, such as converting individual transactions into daily sales totals. Joining combines data from multiple sources into a single unified record, matching related information, such as customer details and their transaction history, that originated in separate systems. Normalization and denormalization restructure data's relational organization: normalization reduces redundancy to improve transactional efficiency, while denormalization introduces controlled redundancy to improve analytical query performance.

8.5.3 Data Loading Strategies

Batch loading transfers large volumes of data in bulk at scheduled intervals, typically suited to reporting needs that do not require up-to-the-minute data. Real-time or stream loading continuously loads and processes data as it is generated, providing immediate access to the latest information, at greater technical complexity and cost than batch approaches. Incremental loading transfers only new or modified data since the last load, rather than the entire dataset, improving efficiency but requiring a reliable mechanism to identify exactly what has changed. Full refresh loading replaces the entire destination dataset with a new copy each time, simpler to implement but far less efficient for large, frequently updated datasets.

8.6 Table 8.1, Data Loading Strategies Compared

Strategy	Best Suited For	Key Trade-Off
Batch Loading	Scheduled reporting, non-urgent data	Simplicity, but the data can be hours or days old
Real-Time / Stream Loading	Fraud detection, live dashboards	Immediate data, but higher technical complexity and cost

Strategy	Best Suited For	Key Trade-Off
Incremental Loading	Large, frequently updated datasets	Efficient, but requires reliable change-tracking
Full Refresh Loading	Small datasets, simplicity priority	Simple, but inefficient and slow at scale

Table 8.1. Choosing a data loading strategy.

8.7 Government Policy and Professional Standards

ETL processes that handle personal data must maintain traceability consistent with Data Privacy Act accountability requirements, since transformation steps that combine or aggregate data from multiple sources can inadvertently create new personal data risks, such as re-identifying individuals from data that was anonymized in its source system but becomes identifiable once joined with other data during transformation.

8.8 Philippine Context and Case Study

PHILIPPINE CONTEXT

Philippine organizations with limited technical infrastructure budgets often favor batch ETL processes over real-time streaming approaches, given the lower infrastructure and technical staffing costs of batch processing. This is a reasonable, deliberate trade-off for most reporting use cases. However, organizations in genuinely time-sensitive sectors, such as fraud detection in financial services, increasingly find the investment in real-time capability justified despite the higher cost.

The batch-versus-real-time decision is also shaped by talent availability in the Philippine labor market, specifically, since real-time streaming-pipeline expertise remains considerably scarcer and correspondingly more expensive to hire or retain than batch ETL expertise. A mid-sized Philippine organization evaluating this trade-off should weigh not only the technical infrastructure cost difference this chapter describes, but the practical difficulty of sustainably staffing whichever approach it chooses over the long term, since a sophisticated pipeline no one on staff can maintain is a liability rather than an asset.

CASE STUDY

Illustrative composite case.

A Philippine e-commerce platform initially used full-page refreshes for its nightly sales data warehouse update, reloading the entire multi-year transaction history every night, even when little had changed. As transaction volume grew, this process began taking over ten hours to complete, frequently failing to finish before the business day began. A redesigned pipeline using incremental loading, transferring only transactions from the previous day rather than the full history, reduced the nightly update to under twenty minutes, illustrating how a loading strategy appropriate at one data volume can become untenable as that volume grows.

Implementing incremental loading required solving a problem full refresh loading had never needed to address: reliably identifying exactly which records were new or modified since the previous run, since any gap in this change-tracking mechanism would silently produce an incomplete warehouse rather than the failed, but at least visible, timeout the full refresh approach had produced. The team's solution, a dedicated timestamp column updated automatically on every insert or modification, added a small amount of upfront schema design work in exchange for the reliability an incremental strategy depends on, a trade-off Table 8.1 identifies as the central cost of choosing incremental over full refresh loading.

8.9 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Choose a loading strategy based on actual data volume and freshness requirements, not by default, and revisit the choice as those requirements change.
- Best Practice: Document every transformation step applied to personal data, supporting both debugging and Data Privacy Act traceability requirements.
- Common Pitfall: Defaulting to full refresh loading indefinitely as data volume grows, without recognizing when incremental loading becomes necessary.

- Common Pitfall: Performing joins or aggregations that inadvertently re-identify individuals from data that was properly anonymized in its source.

8.10 Summary and Key Takeaways

ETL (extraction, transformation, and loading) is the standard process for moving data from source systems into a consistent, usable destination format, using techniques such as filtering, aggregation, and joining, while reshaping the data along the way. Loading strategy, batch, real-time, incremental, or full refresh, should be chosen deliberately based on data volume and freshness requirements rather than by default.

- Common Pitfall: Defaulting to full refresh loading indefinitely as data volume grows, without recognizing when incremental loading becomes necessary.
- Common Pitfall: Performing joins or aggregations that inadvertently re-identify individuals from properly anonymized source data.
- Common Pitfall: Leaving transformation logic undocumented, making it difficult to trace how a final figure was actually derived.

KEY TAKEAWAYS

- ETL extracts data from sources, transforms it into a usable form, and loads it into a destination system.
- Core transformation techniques include filtering, aggregation, joining, and normalization or denormalization.
- Loading strategy should match actual data volume and freshness needs, and should be revisited as those needs change.

- Transformations that join or aggregate personal data can inadvertently create new Data Privacy Act risks even from properly anonymized sources.

8.11 Key Terms, Reflection, and Discussion

ETL. ELT. Data Pipeline. Data Transformation. Filtering. Aggregation. Joining. Batch Loading. Incremental Loading.

- Reflection: What transformation steps would your Data Project Build dataset need before it was ready for analysis?
- Discussion: Is ELT's approach of transforming data after loading a genuine improvement over traditional ETL, or does it simply shift complexity to a different stage of the pipeline?
- Discussion: This chapter warns that joins and aggregations can re-identify individuals from anonymized data. How would you test whether a specific transformation created this risk before deploying it?

8.12 Activity and Data Project Build, Step 8

DATA PROJECT BUILD, STEP 8

Design an ETL plan for your Data Project Build dataset: specify the extraction source, at least two transformation steps it would require, and a justified loading strategy from Table 8.1.

8.13 Chapter Self-Check

CHAPTER SELF-CHECK

1. ETL stands for:

- a) Encrypt, Transfer, Load b) Extract, Transform, Load c) Evaluate, Test, Launch d) Extract, Transfer, Log

2. Which loading strategy transfers only new or modified data since the last load?

- a) Full refresh loading b) Batch loading c) Incremental loading d) Real-time loading

3. ELT differs from ETL primarily by:

- a) Skipping the extraction step entirely b) Transforming data after loading rather than before c) Never using a destination system d) Being unrelated to data pipelines

4. Combining customer details and transaction history from separate source systems into one record is an example of:

- a) Filtering b) Aggregation c) Joining d) Full refresh loading

5. A key ETL-related risk under the Data Privacy Act is:

- a) Slower reporting speed b) Re-identifying individuals through joins of otherwise anonymized data c) Higher storage cost d) Reduced data volume

8.14 References

- Kimball, R., & Ross, M. (2013). The data warehouse toolkit: The definitive guide to dimensional modeling (3rd ed.). Wiley.
- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- Apache Spark. Apache Spark documentation (latest edition).

End of Chapter 8. Continue to Chapter 9: Data Processing Tools and Automation

DATA PROCESSING TOOLS AND AUTOMATION

Scaling Data Pipelines Beyond What Manual Effort Can Sustain

“A pipeline that requires someone to remember to run it manually is not yet a pipeline. It is a task waiting to be forgotten.”

9.0 Chapter Overview

Chapter 8 taught you the ETL process and its transformation techniques conceptually. This chapter closes Part IV by surveying the tools organizations actually use to implement and automate that process at scale, as well as the automation principles that distinguish a reliable, production-grade pipeline from a fragile, manually triggered script.

9.1 Learning Outcomes

- LO 9.1 Identify major categories of data processing tools and their appropriate use cases.
- LO 9.2 Explain the principles and benefits of data pipeline automation.
- LO 9.3 Evaluate when automation investment is justified relative to manual or semi-manual processing.

9.2 Introduction

A data pipeline that works correctly once, run manually by a knowledgeable analyst, is not the same as one an organization can actually depend on. Dependability requires automation: the pipeline must run on its own schedule, handle routine errors gracefully, and alert someone when something genuinely requires human attention, without depending on a specific individual to remember to trigger it manually and be available to troubleshoot it by hand every time something goes wrong.

This chapter closes Part IV by asking a question Chapter 8 deliberately left open: once you know how to design an ETL pipeline conceptually, how do you actually make it run reliably, at scale, without requiring constant manual attention? The tools and automation principles here

are what separate a pipeline that worked once in a classroom exercise from one that an organization could genuinely put into production.

9.3 Historical Background

Early data processing tools were largely proprietary, expensive enterprise software targeted at large organizations, reflecting the high infrastructure and licensing costs of the era's computing resources. The 2000s saw the rise of powerful open-source alternatives, notably Apache Hadoop for distributed storage and processing of very large datasets, which dramatically lowered the cost barrier to large-scale data processing and made sophisticated pipelines accessible to organizations well beyond the largest enterprises.

Apache Spark's emergence in the early 2010s further advanced the field by offering substantially faster in-memory processing than Hadoop's original disk-based approach, and cloud providers' managed data processing services in the years since have continued this democratization trend, allowing organizations to access sophisticated, automated processing infrastructure without operating the underlying hardware themselves.

This progression from proprietary enterprise software to open-source frameworks to fully managed cloud services traces a consistent pattern worth naming explicitly. Each generation has shifted more of the underlying technical complexity away from the organizations using the tools and toward the tools themselves, allowing practitioners to focus increasingly on what a pipeline should do rather than the low-level mechanics of how it runs. This trajectory is part of why Section 9.5.3's judgment about when automation investment is justified has become more accessible to a broader range of organizations over time, even as the underlying distributed processing technology has grown more capable, not less.

9.4 Definitions

KEY TERMS

Pipeline Automation: The practice of configuring a data pipeline to run, monitor itself, and handle routine errors without requiring manual intervention for each execution.

Orchestration: The coordination of multiple interdependent pipeline tasks, ensuring they run in the correct order and handling dependencies between them.

Distributed Processing: Splitting a large data processing task across multiple computing nodes working in parallel, enabling processing at a scale that a single machine could not handle.

Data Pipeline Monitoring: The ongoing tracking of a pipeline's execution status, performance, and error conditions to ensure it continues operating correctly over time.

9.5 Core Concepts and Frameworks




Figure 9.1. Pipeline Orchestration.

9.5.1 Categories of Data Processing Tools

Distributed processing frameworks, such as Apache Hadoop and Apache Spark, handle very large-scale data processing by splitting work across many computing nodes running in parallel, which is essential for organizations processing data volumes beyond what a single machine can handle in a reasonable time. Orchestration tools coordinate multi-step pipelines, ensuring tasks run in the correct dependency order and handling retries when a step fails, rather than requiring each step to be triggered manually in sequence. Managed cloud data processing services increasingly bundle both capabilities into a single platform, reducing the technical infrastructure an organization must operate and maintain.

9.5.2 Why Automation Matters Beyond Convenience

Automation is often framed as merely a convenience, saving an analyst the trouble of running a process manually, but its more significant benefit is reliability. An automated pipeline runs on a consistent schedule regardless of staff availability, applies the same transformation logic identically every time rather than being subject to manual execution variation, and can be configured to alert the appropriate person immediately when a genuine error occurs, rather than allowing a silent failure to go unnoticed until a downstream report looks wrong days or weeks later.

9.5.3 When Automation Investment Is, and Is Not, Justified

Automation carries real setup and maintenance costs, and is not universally justified for every data process. A one-time data migration or a rarely run, low-stakes analysis may not justify the investment automation requires. A recurring pipeline feeding a regularly used report or operational system almost always does so, since the cumulative cost of repeated manual execution and the risk of a missed or inconsistent manual run typically exceeds automation's setup cost within a fairly short time.

9.6 Table 9.1, When to Automate

Factor	Favors Automation	Favors Manual/Ad Hoc
Frequency	Runs regularly (daily, weekly)	One-time or rarely run
Stakes	Feed decisions or operational systems	Low-stakes exploratory analysis
Consistency Need	Must run identically every time	Some execution variation is acceptable
Data Volume	Large enough to require distributed processing	Small enough for manual handling

Table 9.1. Deciding whether pipeline automation is justified.

9.7 Government Policy and Philippine Context

PHILIPPINE CONTEXT

Philippine organizations increasingly access distributed processing and orchestration capabilities through managed cloud platforms rather than operating on-premises Hadoop or Spark clusters directly, given the significant technical staffing required for such infrastructure. This trend has made sophisticated pipeline automation genuinely accessible to mid-sized Philippine organizations that would not previously have justified the capital and staffing investments required for on-premises infrastructure.

This accessibility shift carries a specific implication for MDM graduates entering the Philippine job market: the most valuable skill is increasingly not deep, low-level familiarity with any single distributed processing framework's internal configuration, but the judgment to select and configure the right managed service for a given organization's actual scale and budget, and the discipline to build proper monitoring and error handling around it regardless of which specific platform is chosen. Employers report that this judgment, more than platform-specific technical depth alone, is what distinguishes a genuinely effective junior data professional from one who has only completed platform-specific certification coursework.

9.8 Case Study and Best Practices

CASE STUDY

Illustrative composite case.

A Philippine healthcare network's monthly reporting process relied on a single analyst manually running a sequence of six scripts in a specific order each month, a process that took roughly two full working days and occasionally produced inconsistent results when a step was run out of order, or a script was updated without the others being updated to match. Automating this sequence with an orchestration tool, explicitly defining each step's dependencies and automatically flagging errors, reduced monthly

processing time to under 2 hours and eliminated sequencing errors, while also removing the organization's dependency on one analyst's availability each month.

The dependency risk this automation eliminated became concrete only after the fact: the analyst who had originally built and run the manual process took an extended medical leave shortly before the automation project began, and the organization discovered that no one else fully understood the correct execution order well enough to reproduce the monthly report reliably in her absence. This near-miss, more than any abstract efficiency argument, was what ultimately secured budget approval for the orchestration tool, illustrating how automation's reliability benefit often becomes visible to decision-makers only once its absence has already created a real, felt risk.

- Best Practice: Automate any pipeline that runs regularly and feeds a decision or operational system, prioritizing reliability over the upfront convenience of manual execution.
- Best Practice: Build monitoring and alerting into every automated pipeline from the start, rather than discovering a silent failure only when a downstream report looks wrong.

9.9 Summary and Key Takeaways

Data processing tools span distributed processing frameworks, orchestration tools, and increasingly bundled managed cloud services, and automation's central value is reliability, not merely convenience: consistent execution, graceful error handling, and timely alerting. Automation investment is justified for recurring, higher-stakes pipelines, but less so for one-time or low-stakes analysis. This chapter closes Part IV; Part V now turns to securing the data these pipelines process.

- Common Pitfall: Automating a pipeline without building in monitoring and alerting, so failures go unnoticed until a downstream report looks wrong.
- Common Pitfall: Investing in automation for a one-time, low-stakes process where the setup cost is never recovered.

- Common Pitfall: Depending on a single analyst's manual execution of a recurring pipeline, creating a hidden single point of failure.

KEY TAKEAWAYS

- Distributed processing frameworks like Hadoop and Spark handle data volumes beyond a single machine's capacity.
- Orchestration tools coordinate multi-step pipelines and their dependencies automatically.
- Automation's central value is reliability and consistency, not merely saved manual effort.
- Automation investment is best justified for recurring, higher-stakes pipelines, not universally for every data process.

9.10 Key Terms, Reflection, and Discussion

Pipeline Automation. Orchestration. Distributed Processing. Data Pipeline Monitoring. Apache Hadoop. Apache Spark.

- Reflection: Would the ETL pipeline you designed in Chapter 8 for your Data Project Build justify automation investment, using Table 9.1's criteria?
- Discussion: As managed cloud processing services lower the technical barrier to sophisticated pipelines, what skills remain essential for a data management professional to understand, even if they no longer configure the underlying infrastructure by hand?

- Discussion: The healthcare case study shows that automation's value became visible only after a near-miss. How might an organization make the case for automation investment before such a risk actually materializes?

9.11 Activity and Data Project Build, Step 9

DATA PROJECT BUILD, STEP 9

Using Table 9.1, assess whether your Data Project Build's ETL pipeline (Step 8) justifies automation investment. If yes, specify what monitoring and alerting you would build in; if no, justify why manual or ad hoc processing remains appropriate.

9.12 Chapter Self-Check

CHAPTER SELF-CHECK

1. The primary value of pipeline automation, beyond convenience, is:

- a) Lower electricity cost b) Reliability and consistent execution c) Eliminating the need for data quality checks d) Replacing the need for data governance

2. Apache Spark's key advantage over the original Hadoop approach was:

- a) Lower accuracy b) Faster in-memory processing c) No longer requiring distributed computing d) Eliminating the need for ETL

3. Orchestration tools primarily handle:

- a) Data encryption b) Coordinating multi-step pipeline dependencies and order c) Physical server cooling d) Customer relationship management

4. According to Table 9.1, which factor favors automation?

- a) A one-time, rarely run analysis b) A pipeline that feeds a regularly used operational system c) Very small data volume d) Low-stakes exploratory work

5. Distributed processing frameworks are most necessary when:

- a) Data volume exceeds a single machine's practical capacity b) Data is very small c) No transformation is needed d) Data is only used once

9.13 References

- Zaharia, M., et al. (2016). Apache Spark: A unified engine for big data processing. Communications of the ACM, 59(11), 56 to 65.
- Apache Hadoop. Apache Hadoop documentation (latest edition).
- White, T. (2015). Hadoop: The definitive guide (4th ed.). O'Reilly Media.

End of Chapter 9. Continue to Chapter 10: Data Security Threats, Encryption, and Access Control

PART V

DATA SECURITY AND PRIVACY

Part V addresses whether data survives contact with threats that should never have access to it, covering core security concepts, common threats, and the Philippine Data Privacy Act's requirements for the responsible handling of personal data.

DATA SECURITY THREATS, ENCRYPTION, AND ACCESS CONTROL

Protecting Data from Those Who Should Not Have It

“Every other chapter in this book assumes the data survives. This chapter is about making sure that the assumption holds.”

10.0 Chapter Overview

With data collected, stored, cleansed, and processed through the chapters so far, Part V turns to a question that, if answered badly, undermines everything built in the preceding chapters: Is this data actually protected from unauthorized access, alteration, or loss? This chapter introduces the core security concepts, common threats, and protective mechanisms, encryption and access control chief among them, that data management depends on.

10.1 Learning Outcomes

- LO 10.1 Explain the CIA triad and its role in data security.
- LO 10.2 Identify common threats and vulnerabilities to organizational data.
- LO 10.3 Apply encryption and access control concepts to protect data appropriately.

10.2 Introduction

Data security is sometimes treated as a narrow technical specialty, firewalls and antivirus software, separate from the broader data management practices this book has covered so far. This framing understates how deeply security is woven into every stage of the lifecycle already covered: a storage architecture chosen in Chapter 5 without security consideration, a pipeline built in Chapters 8 and 9 that moves data through unencrypted channels, each represents a security decision, made well or made poorly, whether or not anyone explicitly labeled it as one.

Part V opens with this chapter because the stakes of everything built in Parts I through IV, a governed, well-collected, quality-assured, processed dataset, are only realized if that data

survives contact with people and systems that should never have had access to it. A security failure does not merely damage the specific data affected; it can retroactively undermine the trust every earlier chapter's careful work was meant to establish.

10.3 Historical Background

Formal information security practice predates modern data management, with roots in military and government cryptography extending back centuries. Still, its application to organizational data security intensified sharply as computer networks, and later the internet, connected previously isolated systems and dramatically expanded the potential attack surface. High-profile data breaches from the 2000s onward, affecting both private companies and government agencies, shifted data security from a specialist IT concern into a boardroom and public-policy issue, directly motivating the wave of data protection legislation, including the Philippines' Data Privacy Act, that Chapter 11 examines in depth.

The CIA triad itself, despite its foundational status in this chapter and the field generally, has a somewhat informal history; unlike Codd's relational model or Porter's Five Forces in other disciplines, no single paper or author is credited with its formal introduction, and it appears to have emerged gradually through military and early computer security practice before becoming the standard teaching framework it is today. This informal origin has not diminished its usefulness. Still, it is worth knowing that the triad functions more as a widely adopted organizing convention than as a rigorously derived theoretical model, which is part of why some practitioners have proposed extending it with additional properties, such as authenticity or non-repudiation, without displacing the original three.

10.4 Definitions

KEY TERMS

CIA Triad, the foundational information security model comprising Confidentiality, Integrity, and Availability, the three properties that security measures aim to protect.

Encryption is the process of converting data into a coded format that is unreadable without a decryption key, protecting it both at rest and in transit.

Access Control: Mechanisms that restrict who can view or use specific data, commonly implemented through role-based permissions.

Malware, Software designed to damage, disrupt, or gain unauthorized access to a system, including viruses, ransomware, and spyware.

Multi-Factor Authentication (MFA) is a security practice that requires users to provide multiple independent factors of verification before gaining system access.

10.5 Core Concepts and Frameworks

The CIA Triad




Figure 10.1. The CIA Triad.

10.5.1 The CIA Triad

Confidentiality ensures that data is accessible only to authorized users, primarily through encryption and access controls. Integrity ensures that data remains accurate and unaltered except through authorized action, protected through mechanisms such as checksums, audit

logs, and the validation techniques covered in Chapter 7. Availability ensures that data remains accessible to authorized users when needed, protected through backup, disaster recovery planning, and defenses against availability attacks, such as distributed denial-of-service (DDoS) attacks. Effective security requires attention to all three properties; a system that is perfectly confidential but frequently unavailable, or perfectly available but easily altered by unauthorized parties, has failed at security just as thoroughly as one with no protections at all.

10.5.2 Common Threats and Vulnerabilities

Malware, software designed to damage or disrupt systems or gain unauthorized access, includes viruses, ransomware that encrypts an organization's own data and demands payment for its release, and spyware that covertly exfiltrates data. Phishing and social engineering attacks target human vulnerability rather than technical vulnerability, tricking authorized users into revealing credentials or granting access through deception rather than breaking through technical defenses directly. Insider threats, whether malicious or the result of simple negligence, remain among the most difficult threat categories to defend against, since insiders already hold legitimate access that external attackers must otherwise work to obtain.

10.5.3 Encryption and Access Control as Complementary Defenses

Encryption and access control address different, complementary aspects of confidentiality. Encryption protects data even if it falls into unauthorized hands, rendering it unreadable without the corresponding decryption key, and is applied to both data at rest (stored on disk) and data in transit (moving across a network). Access control, particularly role-based access control, restricts who can access the data in the first place, granting permissions aligned with job roles rather than granting broad access by default. Neither substitutes for the other: encrypted data, an unauthorized insider can still decrypt because they were mistakenly granted the key, offers little real protection, and tightly access-controlled but unencrypted data remains exposed if that access control is bypassed or the underlying storage is physically compromised.

10.6 Table 10.1, Core Data Security Measures

Measure	Primary Protection	Example
Encryption	Confidentiality of data at rest and in transit	End-to-end encrypted messaging

Measure	Primary Protection	Example
Role-Based Access Control	Confidentiality, limiting exposure by job role	Financial analyst access vs. marketing associate access
Data Masking	Confidentiality during testing or sharing	Showing only the last four digits of a card number
Multi-Factor Authentication	Confidentiality, preventing unauthorized login	Password plus SMS code for online banking
Backup and Disaster Recovery	Availability and integrity after loss or attack	Cloud-based backups for business continuity

Table 10.1. Core data security measures and what they protect against.

10.7 Government Policy and Professional Standards

The Data Privacy Act requires organizations to implement reasonable and appropriate organizational, physical, and technical security measures proportionate to the nature and sensitivity of the personal data they hold, a requirement the National Privacy Commission has interpreted to explicitly include encryption and access control among expected baseline measures for sensitive personal information. The Cybercrime Prevention Act (Republic Act No. 10175) separately criminalizes unauthorized access, data interference, and related offenses, providing a legal deterrent alongside the Data Privacy Act's compliance framework.

10.8 Philippine Context and Case Study

PHILIPPINE CONTEXT

Phishing and social engineering attacks targeting Philippine organizations have grown increasingly sophisticated, frequently impersonating legitimate government agencies, banks, or delivery services to trick employees or the public into revealing credentials or personal information. This trend makes employee security awareness training, alongside technical defenses such as MFA and access controls, an essential and often underinvested component of Philippine organizational security practice.

Impersonation of Philippine government agencies, specifically, tax authorities, social insurance programs, and national identification systems, has become common enough that public awareness campaigns now regularly warn citizens about it directly,

a signal of how significant the pattern has become at a national scale. Organizations that handle any data connected to these government touchpoints, such as employee tax records or benefits enrollment, are among the most common examples, and should treat impersonation of these specific agencies as a realistic, not merely hypothetical, threat when designing staff security awareness training.

CASE STUDY

Illustrative composite case.

A Philippine cooperative's finance officer received an email, apparently from the cooperative's president, requesting an urgent fund transfer to a new supplier account, a classic phishing pattern known as business email compromise. The officer, following an established verification protocol requiring a phone confirmation for any fund transfer request received only by email, called the president directly and discovered the email was fraudulent, the sender's address closely mimicking but not matching the president's actual address. The cooperative's verification protocol, a low-cost procedural control rather than an expensive technical one, prevented what could have been a significant financial loss.

A subsequent review of the cooperative's email security found that a technical control, an email authentication mechanism that would have flagged or blocked the spoofed sender address automatically, was available at essentially no additional cost through the cooperative's existing email provider but had never been enabled. The cooperative implemented this technical control alongside keeping its procedural verification requirement in place, rather than treating the successful procedural catch as evidence that no further investment was needed, precisely the layered defense this chapter recommends: the human verification step was not treated as a replacement for the available technical safeguard, but as a complement to it.

10.9 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Apply role-based access control by default, granting the minimum access necessary for a given job role rather than broad access that is later restricted.
- Best Practice: Combine technical defenses, encryption, MFA, access control, with procedural controls and staff security awareness training, since many successful attacks target human rather than technical vulnerability.
- Current Issue: Ransomware attacks targeting organizational data for extortion have grown substantially more common and more costly across sectors, including in the Philippines.
- Emerging Trend: Zero-trust security models, which require continuous verification for every access request rather than assuming trust based on network location alone, are increasingly replacing older perimeter-based security assumptions.

10.10 Summary and Key Takeaways

The CIA triad, confidentiality, integrity, and availability, frames what data security measures aim to protect, while encryption and access control serve as complementary core defenses against the most common threats: malware, phishing, social engineering, and insider risk. Philippine organizations operate under both the Data Privacy Act's security requirements and the Cybercrime Prevention Act's criminal deterrents, making security a legal as well as operational obligation.

- Common Pitfall: Investing heavily in technical defenses while neglecting staff security awareness training, leaving phishing and social engineering unaddressed.
- Common Pitfall: Encrypting data but granting broad, unrestricted access to the decryption key, undermining the protection encryption is meant to provide.
- Common Pitfall: Focusing security investment entirely on external threats while underestimating insider risk, whether malicious or negligent.

KEY TAKEAWAYS

- The CIA triad, confidentiality, integrity, and availability, frames the full scope of data security, not confidentiality alone.
- Encryption and access control are complementary, not interchangeable, defenses.
- Phishing and social engineering target human vulnerability and require procedural controls, not only technical ones.
- The Data Privacy Act and Cybercrime Prevention Act together create both compliance and criminal-deterrent obligations for Philippine organizations.

10.11 Key Terms, Reflection, and Discussion

CIA Triad. Encryption. Access Control. Malware. Phishing. Insider Threat. Multi-Factor Authentication.

- Reflection: Which CIA triad property, confidentiality, integrity, or availability, is most at risk for your Data Project Build dataset, and why?
- Discussion: Are Philippine organizations, particularly small and medium enterprises, investing adequately in staff security awareness training relative to technical defenses? What would change your assessment?
- Discussion: The cooperative case study succeeded through a procedural control rather than a technical one. Are procedural controls generally more or less reliable than technical controls, and why?

10.12 Activity and Data Project Build, Step 10

DATA PROJECT BUILD, STEP 10

Complete a security assessment for your Data Project Build dataset using Table 10.1: specify which security measures apply, identify the most significant threat from Section 10.5.2 that your dataset faces, and propose one technical and one procedural control to address it.

10.13 Chapter Self-Check

CHAPTER SELF-CHECK

1. The three properties in the CIA triad are:

- a) Confidentiality, Integrity, Availability b) Control, Investigation, Auditing c) Compliance, Insurance, Accountability d) Consistency, Integration, Automation

2. Ransomware is best classified as:

- a) A type of access control b) A type of malware c) An encryption technique d) A data quality dimension

3. Phishing attacks primarily exploit:

- a) Hardware failures b) Human vulnerability through deception c) Weak encryption algorithms only d) Database schema errors

4. Role-based access control grants permissions based on:

- a) Random assignment b) Job role c) Length of employment only d) Physical office location only

5. In the Philippines, unauthorized access to data is separately criminalized under:

- a) The Data Privacy Act only b) The Cybercrime Prevention Act (RA 10175) c) The Consumer Act d) The Corporation Code

10.14 References

- Andress, J. (2014). The basics of information security: Understanding the fundamentals of InfoSec in theory and practice (2nd ed.). Elsevier.
- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- Republic Act No. 10175 (2012). Cybercrime Prevention Act of 2012.
- National Privacy Commission. Guidelines on organizational, physical, and technical security measures.

End of Chapter 10. Continue to Chapter 11: Data Privacy, Compliance, and Regulatory Standards

DATA PRIVACY, COMPLIANCE, AND REGULATORY STANDARDS

Meeting the Legal and Ethical Obligations Data Carries

“Compliance is the floor a data management program must clear. It was never meant to be the ceiling.”

11.0 Chapter Overview

Chapter 10 covered the technical and procedural defenses that protect data from unauthorized access. This chapter closes Part V by examining the legal and regulatory obligations that shape how personal data must be handled, focusing on the Philippines' Data Privacy Act and the compliance practices organizations use to meet those obligations.

11.1 Learning Outcomes

- LO 11.1 Explain the core principles of the Philippine Data Privacy Act.
- LO 11.2 Identify individual employment and data subject rights regarding personal data.
- LO 11.3 Apply compliance practices to a real or hypothetical data management scenario.

11.2 Introduction

Every chapter in this book has referenced the Data Privacy Act at some point, in collection, in quality, in storage, in security, because privacy compliance is not a discrete stage of the data lifecycle, the way collection or storage is. It is a set of obligations that attaches to personal data throughout its entire lifecycle, from the moment of collection through eventual disposal. This chapter consolidates that thread into a single, coherent treatment of the law itself and what compliance with it actually requires in practice.

This chapter closes Part V by treating compliance as a distinct professional competency in its own right, not merely a legal detail attached to the technical work of earlier chapters. A data professional who understands encryption and access control (Chapter 10) but cannot explain a

data subject's rights under the law, or draft a basic Privacy Impact Assessment, holds only half of what Philippine organizations increasingly require of the role.

11.3 Historical Background

The Philippines' Data Privacy Act of 2012 was substantially influenced by earlier, more established data protection frameworks, particularly the European Union's data protection directive that preceded the later General Data Protection Regulation, reflecting a broader global trend toward comprehensive personal data protection legislation rather than the narrower, sector-specific privacy rules many countries, including the Philippines, previously had relied on. The Act's implementing rules and regulations, and the National Privacy Commission established to enforce it, followed in subsequent years, gradually building out the practical compliance infrastructure organizations now operate under.

It is worth noting that the Data Privacy Act's 2012 enactment predates the European Union's General Data Protection Regulation, which did not take effect until 2018, even though this chapter describes GDPR as a later, more established reference point in global data protection practice. This ordering reflects the fact that the Philippines drew primarily on the EU's earlier data protection directive, GDPR's regulatory predecessor, and subsequent Philippine implementing guidance has continued to draw on GDPR's more developed jurisprudence and practical guidance, even though the Philippine statute itself came first, illustrating how data protection law across jurisdictions has developed through ongoing mutual influence rather than in a strict, one-directional sequence.

11.4 Definitions

KEY TERMS

Personal Data, any information relating to an identified or identifiable individual, is the central subject matter protected by the Data Privacy Act.

Sensitive Personal Information, A specially protected category of personal data under the Data Privacy Act, including health, genetic, and certain other information warranting heightened protection.

Data Subject: The individual whose personal data is being collected, processed, or held by an organization.

Data Subject Rights: The specific rights the Data Privacy Act grants individuals over their own personal data, including access, correction, and erasure.

The National Privacy Commission (NPC) is the Philippine government body responsible for administering and enforcing the Data Privacy Act.

11.5 Core Concepts and Frameworks




Figure 11.1. Data Subject Rights.

11.5.1 The Data Privacy Act's Core Principles

Three core principles anchor the Act's requirements. Transparency requires that data subjects be informed of the nature, purpose, and extent of the processing of their personal data, which connects directly to the collection-stage consent obligations covered in Chapter 4. A legitimate purpose requires that processing be compatible with a declared, specified purpose, and not repurposed for unrelated uses without a fresh basis. Proportionality requires that

processing be adequate, relevant, and limited to what is necessary for the declared purpose, directly underpinning the retention and disposal discipline covered in Chapter 3.

11.5.2 Data Subject Rights

The Act grants individuals specific, enforceable rights over their own personal data. The right to be informed requires notice before or at the point of collection. The right of access allows individuals to obtain a copy of the personal data an organization holds about them. The right to rectification allows the correction of inaccurate data, which connects to Chapter 7's cleansing techniques. The right to erasure or blocking allows individuals to request deletion under specified conditions, which aligns with Chapter 3's lifecycle disposal stage. The right to data portability allows individuals to obtain and reuse their data across different services in certain circumstances.

11.5.3 Personal Data in the Employment Context

Employment relationships generate substantial personal data, often including sensitive personal information such as health records for medical clearances or benefits, that require particular compliance attention. Employers must establish a legitimate basis for the personal data they collect from employees, apply appropriate security measures proportionate to the sensitivity of the data, and respect employees' data subject rights, including the right to access their own employment records, subject to legitimate employer confidentiality interests, such as internal investigations.

11.6 Table 11.1, Data Subject Rights at a Glance

Right	What It Allows	Related Chapter
Right to Be Informed	Notice of collection purpose and extent	Chapter 4 (Collection)
Right to Access	Obtain a copy of one's own personal data	Chapter 5 (Storage)
Right to Rectification	Correct inaccurate personal data	Chapter 7 (Cleansing)
Right to Erasure / Blocking	Request deletion under specified conditions	Chapter 3 (Lifecycle)
Right to Data Portability	Obtain and reuse data across services	Chapters 8 to 9 (Processing)

Table 11.1. Core Data Privacy Act data subject rights.

11.7 Government Policy and Compliance Practice

Practical Data Privacy Act compliance typically requires several documented artifacts: a Privacy Impact Assessment identifying and mitigating privacy risks in a given system or process, a Privacy Manual documenting the organization's overall data protection policies and procedures, and appointment of a Data Protection Officer, introduced in Chapter 2's governance discussion, as the accountable compliance role. The National Privacy Commission has published guidance and, in cases of significant non-compliance, issued enforcement actions against organizations across sectors, establishing a growing body of practical compliance precedent.

11.8 Philippine Context and Case Study

PHILIPPINE CONTEXT

Philippine organizations across sectors, particularly those handling health, financial, or employment data, have accelerated investment in compliance following several well-publicized National Privacy Commission enforcement actions, shifting privacy compliance from a low-priority legal formality to an active operational concern with real reputational and financial consequences for non-compliance. Small and medium enterprises, however, still frequently lack the dedicated compliance resources larger organizations can invest, making efficient, practical compliance approaches particularly valuable for this segment.

For a resource-constrained Philippine organization, the most efficient starting point is rarely a comprehensive, all-at-once compliance overhaul, which can be prohibitively expensive to execute properly, but rather a prioritized approach beginning with the personal data inventory described in this chapter's case study. That inventory alone, even before every other compliance artifact this chapter mentions, is in place, immediately improves an organization's ability to respond to data subject requests and identify its most significant privacy risks, making it a disproportionately high-value first investment relative to its relatively modest cost.

CASE STUDY

Illustrative composite case.

A Philippine HR outsourcing firm received a data subject access request from a former client's employee, seeking a copy of all personal data the firm held about them. Without a documented process for handling such requests, the firm took several weeks and considerable manual effort to locate the relevant records scattered across multiple systems, well beyond the response timeline recommended by the Data Privacy Act. The firm's subsequent investment in a centralized personal data inventory, mapping which systems held which personal data, directly addressed this gap and reduced future access request response time from weeks to days.

Building the inventory also surfaced a finding the firm had not anticipated: several systems retained personal data belonging to individuals from client relationships that had ended years earlier, data the firm had no ongoing legitimate purpose to retain under the proportionality principle discussed in this chapter. The inventory exercise, therefore, accomplished two goals simultaneously, enabling faster future access-request response and identifying data that should have been disposed of under Chapter 3's lifecycle discipline, illustrating how the compliance work in this chapter and the lifecycle work in Chapter 3 are, in practice, the same underlying discipline applied from two different starting questions.

11.9 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Maintain a documented personal data inventory, mapping what personal data exists in which system, to enable a timely response to data subject access, rectification, and erasure requests.
- Best Practice: Conduct a Privacy Impact Assessment before launching any new system or process that will handle personal data, not retroactively after launch.

- Common Pitfall: Treating the Data Protection Officer appointment as sufficient compliance on its own, without the underlying documented policies, inventory, and processes the role depends on to function.
- Emerging Trend: Cross-border data transfer requirements are receiving increasing regulatory attention as more Philippine organizations adopt cloud services hosted outside the country.

11.10 Summary and Key Takeaways

The Data Privacy Act's core principles, transparency, legitimate purpose, and proportionality, and its data subject rights, access, rectification, erasure, and portability, among them, thread through every stage of the data lifecycle this book has covered, from collection through disposal. This chapter closes Part V; Part VI now turns to extracting and communicating insight from data that has been properly collected, secured, and governed.

- Common Pitfall: Treating the Data Protection Officer appointment as sufficient compliance, without the underlying inventory and documented processes the role depends on.
- Common Pitfall: Conducting a Privacy Impact Assessment only after a system has already launched, rather than before.
- Common Pitfall: Lacking a documented personal data inventory, making a timely response to data subject access requests impractical.

KEY TAKEAWAYS

- The Data Privacy Act rests on transparency, legitimate purpose, and proportionality as core principles.

- Data subject rights include access, rectification, erasure, portability, and the right to be informed.
- Practical compliance requires documented artifacts: a Privacy Impact Assessment, a Privacy Manual, and a personal data inventory.
- Privacy compliance is not a discrete lifecycle stage but an obligation threaded through every chapter of this book.

11.11 Key Terms, Reflection, and Discussion

Personal Data. Sensitive Personal Information. Data Subject. Data Subject Rights. Privacy Impact Assessment. Data Protection Officer.

- Reflection: Does your Data Project Build dataset contain personal data or sensitive personal information? If so, which data subject rights from Table 11.1 would be most relevant to it?
- Discussion: Should smaller organizations face a reduced compliance burden under the Data Privacy Act relative to large enterprises, or should the same standard apply regardless of organizational size?
- Discussion: The HR outsourcing case study found personal data retained past any legitimate purpose during its inventory exercise. Should this kind of accidental discovery be reported to the National Privacy Commission, even without an actual breach?

11.12 Activity and Data Project Build, Step 11

DATA PROJECT BUILD, STEP 11

Draft a brief Privacy Impact Assessment for your Data Project Build dataset: identify whether it contains personal or sensitive personal information, what data subject rights would apply, and one specific compliance practice from this chapter you would implement.

11.13 Chapter Self-Check

CHAPTER SELF-CHECK

1. The three core principles of the Data Privacy Act are:

- a) Speed, cost, quality b) Transparency, legitimate purpose, proportionality c) Encryption, access control, backup d) Accuracy, completeness, timeliness

2. The right allowing an individual to correct inaccurate personal data is called:

- a) Right to erasure b) Right to rectification c) Right to portability d) Right to be informed

3. The National Privacy Commission's role is to:

- a) Provide cloud storage services b) Administer and enforce the Data Privacy Act c) Design database schemas d) Approve business loans

4. A Privacy Impact Assessment should ideally be conducted:

- a) Only after a data breach occurs b) Before launching a new system that handles personal data c) Once every ten years, regardless of changes d) Only by the National Privacy Commission itself

5. Sensitive personal information under the Data Privacy Act includes:

- a) Only financial data b) Health and certain other specially protected categories c) All data, regardless of type d) Only publicly available data

11.14 References

- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- National Privacy Commission. Data Privacy Act implementing rules and regulations.
- National Privacy Commission. Guidelines on data subject rights and privacy impact assessments.

End of Chapter 11. Continue to Chapter 12: Data Analytics Fundamentals and Types

PART VI

DATA ANALYTICS, VISUALIZATION, AND THE FUTURE OF DATA MANAGEMENT

Part VI closes the data management journey: extracting insights through the four types of analytics, communicating those insights through honest visualization, and looking ahead to the trends and technologies reshaping the field.

DATA ANALYTICS FUNDAMENTALS AND TYPES

Turning Well-Managed Data into Answers

“Analytics does not create insight from nothing. It reveals insight that was already latent in data disciplined enough to hold it.”

12.0 Chapter Overview

Every prior chapter in this book has been, in a real sense, preparation for this one. Part VI opens by asking what all that collection, storage, quality, processing, and security work was ultimately for: extracting insights that support better decisions. This chapter introduces the four fundamental types of data analytics and the questions each is built to answer.

12.1 Learning Outcomes

- LO 12.1 Explain the concepts and applications of data analytics.
- LO 12.2 Differentiate between the four types of data analytics.
- LO 12.3 Select an appropriate analytics type for a given organizational question.

12.2 Introduction

A common and costly mistake is treating data analytics as a single, undifferentiated activity, running numbers to see what turns up, rather than a deliberate choice among distinct analytical approaches, each suited to a different kind of question. An organization that wants to know why sales declined last quarter needs a fundamentally different analytical approach than one that wants to predict which customers are likely to churn next quarter, even though both questions concern the same underlying sales data.

Part VI opens with this chapter because it represents a genuine culmination, not merely the next topic in sequence. Every governance, collection, quality, processing, and security practice covered in Parts I through V exists to produce data capable of supporting exactly the kind of analysis taught in this chapter and Chapter 13. If any of those earlier foundations were skipped

or done poorly, the limitation will surface here, in analysis that looks sophisticated but rests on data that cannot actually bear its weight.

12.3 Historical Background

Descriptive statistics, summarizing what has happened in a dataset, represent the oldest branch of data analysis, with roots extending back centuries in the development of statistical theory. What is now called business intelligence emerged over the twentieth century as organizations began systematically applying these statistical foundations to organizational decision-making, initially through relatively simple descriptive reporting, then gradually incorporating more sophisticated diagnostic, predictive, and eventually prescriptive techniques as computing power and statistical methodology advanced.

The rise of machine learning from the 2000s onward substantially expanded what predictive and prescriptive analytics could accomplish, allowing organizations to model far more complex patterns in far larger datasets than traditional statistical methods alone could practically handle. However, this expansion also raised the stakes of the data quality and governance foundations covered in Parts I through III, since sophisticated predictive models trained on poor-quality data reliably produce poor-quality, and sometimes actively harmful, predictions.

The progression from descriptive through diagnostic to predictive and prescriptive analytics, as described in this chapter, is sometimes presented as though each type simply superseded the one before it, but the historical reality, and the practical reality any organization actually faces, is additive rather than substitutive. Organizations with mature predictive and prescriptive analytics capabilities still rely heavily on descriptive reporting for routine operational monitoring; the four types coexist as a toolkit for different questions, not a ladder where reaching the top rung eliminates the need for the earlier ones. Table 12.1 and the analytics-type-selection guidance in this chapter are built specifically to reinforce this point.

12.4 Definitions

KEY TERMS

Data Analytics: The process of examining data to conclude, identify patterns, and support decision-making.

Descriptive Analytics, Analytics focused on summarizing historical data to understand what has happened.

Diagnostic Analytics, Analytics focused on identifying the causes behind patterns observed in historical data, answering why something happened.

Predictive Analytics, Analytics that use historical data and statistical or machine learning models to forecast what might happen in the future.

Prescriptive Analytics, Analytics that provide actionable recommendations based on predictive insights, answering what should be done.

12.5 Core Concepts and Frameworks

The Four Types of Data Analytics




Figure 12.1. The Four Types of Data Analytics.

12.5.1 The Four Types of Data Analytics

Descriptive analytics answers what happened, summarizing historical data to identify patterns and relationships, such as a website traffic report showing visitor counts over the past

month. Diagnostic analytics goes a step further to answer why it happened, digging into the underlying causes and contributing factors within the data, such as identifying that a traffic spike coincided with a specific marketing campaign. Predictive analytics uses historical data and statistical or machine learning models to predict future outcomes based on identified patterns. Prescriptive analytics answers the question of what should be done, providing actionable recommendations grounded in predictive insights, such as a specific inventory reorder quantity based on a demand forecast.

12.5.2 Matching Analytics Type to Organizational Question

Each analytics type serves a distinct organizational purpose, and effective analytics practice starts by identifying the type of question being asked before selecting a method. A monthly performance report is fundamentally a descriptive analytics task. An investigation into an unexpected drop in a key metric is a diagnostic analytics task. A demand forecast for next quarter is a predictive analytics task. A recommendation engine suggesting the next best action for a given customer is a prescriptive analytics task. Applying the wrong type, running only descriptive summaries when a diagnostic investigation is actually needed, is a common source of analytics that generates activity without generating genuine decision support.

12.5.3 Analytics Depends on the Foundations Built in Prior Parts

It is worth being explicit about a dependency this chapter inherits from the preceding parts of the book. Descriptive analytics that summarize inaccurate data produces misleading results. Diagnostic analytics investigating causes within inconsistent data may identify spurious, non-existent relationships. Predictive models trained on incomplete or biased data produce unreliable and potentially discriminatory forecasts. Prescriptive recommendations built on flawed predictions actively mislead decision-makers. Every technique in this chapter and Chapter 13 assumes the collection, quality, processing, and governance foundations from Parts I through V are actually in place.

12.6 Table 12.1, The Four Types of Data Analytics

Type	Question Answered	Example
Descriptive	What happened?	Monthly sales summary report

Type	Question Answered	Example
Diagnostic	Why did it happen?	Investigating the root cause of a sales decline
Predictive	What might happen?	Forecasting next quarter's demand
Prescriptive	What should be done?	Recommending an optimal inventory reorder quantity.

Table 12.1. The four types of data analytics and the questions each answers.

12.7 Government Policy and Professional Standards

Predictive and prescriptive analytics involving personal data raise distinct Data Privacy Act considerations beyond those covered in Chapter 11, particularly where automated decision-making significantly affects individuals, such as automated credit scoring or hiring recommendations. Organizations deploying such systems bear heightened responsibility to demonstrate that the underlying data and model do not produce discriminatory or otherwise harmful outcomes, an emerging area of both regulatory attention and professional ethical standards.

12.8 Philippine Context and Case Study

PHILIPPINE CONTEXT

Philippine organizations have historically concentrated their analytics investment in descriptive reporting, monthly dashboards, and summaries, with comparatively less adoption of diagnostic, predictive, and prescriptive analytics. However, this is shifting as data science and analytics talent, including MDM graduates, become more widely available across the Philippine industry. Sectors with strong regulatory reporting requirements, such as banking, have generally led this shift toward more sophisticated analytics maturity.

A practical implication for graduates entering organizations that are still primarily focused on descriptive analytics is that advancing an organization's analytics maturity is often better pursued incrementally than by immediately proposing an ambitious

predictive modeling initiative. Demonstrating value through a well-executed diagnostic analysis first, answering a specific 'why' question that the organization's existing descriptive reports have left unresolved, tends to build internal credibility and stakeholder trust that a more ambitious predictive or prescriptive proposal will later depend on.

CASE STUDY

Illustrative composite case.

A Philippine retail chain's analytics practice initially consisted entirely of descriptive monthly sales reports, which reliably showed that sales at several branches had been declining for two consecutive quarters but offered no explanation for the decline. A diagnostic analysis, cross-referencing the declining branches against local competitor openings, foot traffic data, and staffing levels, identified that the decline correlated most strongly with staffing shortages during peak hours rather than competitive pressure, a cause that the descriptive reports alone had never surfaced. This diagnostic finding directly informed a staffing schedule adjustment that a purely descriptive analytics practice would never have identified as the actual lever to pull.

Encouraged by this result, the chain's analytics team subsequently built a predictive model forecasting which branches were at elevated risk of a similar staffing-driven decline in future quarters, using the staffing and foot-traffic patterns the diagnostic analysis had identified as leading indicators. The predictive model's early warnings allowed the chain to adjust staffing proactively at flagged branches before a decline materialized, rather than only after two quarters of falling sales had already accumulated, illustrating how the four analytics types in Table 12.1 often build on each other sequentially within a single, maturing analytics practice rather than being applied in isolation.

12.9 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Identify which of the four analytics types actually matches the organizational question being asked before selecting a method or tool.
- Best Practice: Verify the data quality and governance foundations from earlier parts of this book before investing in predictive or prescriptive analytics built on top of them.
- Current Issue: Many organizations invest in predictive analytics capability before their descriptive and diagnostic practices are mature, producing sophisticated forecasts built on a shaky foundation of unexamined historical data.
- Emerging Trend: Generative AI tools are increasingly lowering the technical barrier to diagnostic and predictive analytics, though the underlying data quality dependency from Section 12.5.3 remains unchanged.

12.10 Summary and Key Takeaways

Data analytics spans four distinct types: descriptive, diagnostic, predictive, and prescriptive, each answering a different kind of organizational question, and effective practice begins by identifying which question is actually being asked before selecting a method. Every analytics technique this chapter covers depends directly on the foundations for collection, quality, processing, and governance established in the preceding eleven chapters.

- Common Pitfall: Running only descriptive summaries when the actual organizational need is a diagnostic investigation into causes.
- Common Pitfall: Investing in predictive analytics before descriptive and diagnostic practices are mature enough to support it.
- Common Pitfall: Deploying predictive or prescriptive models on personal data without assessing discriminatory or harmful outcomes.

KEY TAKEAWAYS

- Descriptive analytics answers what happened; diagnostic answers why; predictive answers what might happen; prescriptive answers what should be done.
- Matching analytics type to the actual organizational question prevents analytics activity that does not generate genuine decision support.
- The data quality and governance foundations bound analytics quality, which is built on, not by, analytical sophistication alone.
- Predictive and prescriptive analytics involving personal data raise heightened Data Privacy Act and ethical considerations, particularly for automated decisions affecting individuals.

12.11 Key Terms, Reflection, and Discussion

Data Analytics. Descriptive Analytics. Diagnostic Analytics. Predictive Analytics. Prescriptive Analytics.

- Reflection: What organizational question would your Data Project Build dataset be best suited to answer, and which of the four analytics types does that question require?
- Discussion: Should organizations be required to demonstrate their underlying data quality before deploying predictive or prescriptive analytics that significantly affect individuals? What would such a requirement look like in practice?
- Discussion: The retail case study moved from descriptive to diagnostic to predictive analytics over time. Is this sequencing necessary, or could an organization reasonably start with predictive analytics directly?

12.12 Activity and Data Project Build, Step 12

DATA PROJECT BUILD, STEP 12

Using Table 12.1, formulate one question your Data Project Build dataset could answer for each of the four analytics types, and identify which type would provide the most valuable insight for your dataset's core organizational purpose.

12.13 Chapter Self-Check

CHAPTER SELF-CHECK

1. Which analytics type answers the question, "What happened?"

- a) Diagnostic b) Descriptive c) Predictive d) Prescriptive

2. A demand forecast for next quarter is an example of:

- a) Descriptive analytics b) Diagnostic analytics c) Predictive analytics d) None of these

3. Prescriptive analytics is distinguished by:

- a) Only describing past events b) Providing actionable recommendations based on predictive insight c) Ignoring historical data entirely d) Being unrelated to decision-making

4. Why does analytics quality depend on earlier chapters in this book?

- a) It does not; analytics is independent of data quality. b) Analytics built on poor-quality data produce misleading or unreliable results. c) Analytics only uses data collected after Chapter 12. d) Analytics replaces the need for data governance

5. A diagnostic analysis of a sales decline seeks to answer:

- a) What will happen next quarter? b) Why did the decline happen? c) What should be done about it? d) How much revenue was lost, without explanation?

12.14 References

- Fox, P. (2018). Data analytics course. Rensselaer Polytechnic Institute.
- Tw.rpi.edu. Data analytics course materials.
- IBM. SPSS Statistics documentation (latest edition).

End of Chapter 12. Continue to Chapter 13: Data Visualization and Communicating Insights

DATA VISUALIZATION AND COMMUNICATING INSIGHTS

Making Analysis Understandable to the People Who Must Act on It

“An analysis no one understands is, for practical purposes, an analysis that was never done.”

13.0 Chapter Overview

Chapter 12 taught you to choose the right analytical approach for a given question. This chapter closes the analytics-focused portion of Part VI by addressing a separate, equally critical skill: communicating what that analysis found in a way decision-makers can actually understand and act on. A brilliant analysis presented poorly produces no better organizational outcome than no analysis at all.

13.1 Learning Outcomes

- LO 13.1 Utilize data visualization principles to present analytical findings effectively.
- LO 13.2 Select appropriate chart types for different data and communication purposes.
- LO 13.3 Apply best practices for communicating data insights to non-technical audiences.

13.2 Introduction

Data visualization is often treated as a final cosmetic step, making a finished analysis look presentable, rather than as an integral part of the analytical process itself. This undersells its actual function. A well-chosen visualization does not merely decorate a finding; it often reveals a pattern that would remain invisible in a table of raw numbers, and a poorly chosen one can actively mislead an audience even when the underlying analysis was entirely correct.

This chapter closes the analytics portion of Part VI by addressing a skill Chapter 12 deliberately left aside: even a perfectly executed diagnostic or predictive analysis accomplishes

nothing organizationally if the person who needs to act on it cannot understand what it found. Communication is not a lesser, softer skill layered on top of technical analysis; it is what determines whether analytical work actually changes a decision.

13.3 Historical Background

Statistical graphics have a documented history extending back centuries, with William Playfair's late-eighteenth-century line and bar charts among the earliest systematic applications of visual representation to numerical data. Edward Tufte's influential work from the 1980s onward gave the field a more rigorous design vocabulary, emphasizing principles such as maximizing the ratio of data-representing ink to non-data decoration and avoiding visual distortions that misrepresent the underlying data's actual proportions.

The rise of dedicated business intelligence and visualization software from the 2000s onward, and the more recent growth of self-service visualization tools that do not require specialized technical skill to operate, has dramatically expanded who within an organization can produce visualizations, a democratization that has increased both the volume of visualization produced and, without attention to the design principles this chapter covers, the volume of poorly designed and even misleading visualization in circulation.

Playfair's original innovations are worth appreciating specifically because they were genuinely novel inventions, not refinements of existing practice; before his work, there was no established convention for representing quantitative economic data as a line chart or bar chart at all, and his charts had to argue implicitly for their own usefulness to an audience with no prior visual vocabulary for interpreting them. Modern readers take chart literacy largely for granted, but this chapter's emphasis on clear labeling and honest design can be understood as a direct descendant of the same underlying challenge Playfair faced: a visualization is only as useful as its audience's ability to correctly interpret what it shows.

13.4 Definitions

KEY TERMS

Data Visualization: The graphical representation of data and information, designed to make patterns, trends, and outliers visually apparent.

Data-Ink Ratio: Edward Tufte's principle that the proportion of a visualization's ink devoted to representing actual data, rather than decoration, should be maximized.

Dashboard: A visual display consolidating multiple related metrics or visualizations into a single view for ongoing monitoring.

Chart Junk: Tufte's term for unnecessary decorative elements in a visualization that do not convey data and can distract from or distort the actual information.

13.5 Core Concepts and Frameworks

Choosing the Right Chart Type




Figure 13.1. Choosing the Right Chart Type.

13.5.1 Choosing the Right Chart Type

Different chart types suit different analytical purposes, and choosing incorrectly can obscure or distort the very pattern a visualization is meant to reveal. Line charts suit data showing change over time, such as monthly revenue trends. Bar charts suit comparison across discrete categories, such as sales by region. Scatter plots suit examining the relationship between two numerical variables, such as advertising spend against resulting sales. Pie charts, despite their popularity, suit only a narrow purpose, showing proportional parts of a whole, and are frequently misapplied to comparisons or trends they communicate poorly.

13.5.2 Design Principles for Honest, Clear Visualization

Beyond selecting the correct chart type, several design principles separate an honest, clear visualization from a misleading or confusing one. Axes should generally begin at zero for bar charts, since a truncated axis can dramatically exaggerate the apparent size of a real but modest

difference. Color should be used purposefully, distinguishing categories or highlighting a specific finding, not applied decoratively in ways that create visual noise without conveying information. Labels and titles should make a visualization interpretable without requiring the viewer to consult a separate explanation, stating clearly what is being shown and, ideally, the specific insight the chart is meant to convey.

13.5.3 Communicating to Non-Technical Audiences

Analysts frequently overestimate how much technical context a non-technical audience, such as an executive or non-technical client, already holds, leading to visualizations and explanations dense with jargon and statistical detail that obscure rather than clarify the actual finding. Effective communication for this audience typically leads with the finding and its business implication, this metric declined and here is what it means for the decision at hand, rather than leading with methodology, and uses the simplest visualization that accurately conveys the pattern rather than the most technically sophisticated one available.

13.6 Table 13.1, Matching Chart Type to Purpose

Chart Type	Best Suited For	Common Misuse
Line Chart	Trends over time	Used for unordered categories
Bar Chart	Comparison across categories	Truncated axis exaggerating differences
Scatter Plot	Relationship between two variables	Used when only one variable is of interest
Pie Chart	Proportional parts of a single whole	Used for trends or many-category comparisons

Table 13.1. Choosing the right chart type for common analytical purposes.

13.7 Government Policy and Professional Standards

Visualizations built from personal or sensitive data carry the same Data Privacy Act obligations as the underlying data itself, particularly around aggregation choices that might inadvertently make individuals identifiable within a small-group breakdown, even when no individual-level record is directly displayed. Professional data visualization practice, reflected in standards from organizations and authors such as Stephanie Evergreen and Cole Nussbaumer

Knafllic, increasingly treats this kind of privacy-aware aggregation as a core design consideration, not an afterthought separate from visual design itself.

13.8 Philippine Context and Case Study

PHILIPPINE CONTEXT

Philippine government open data initiatives increasingly publish public dashboards on health, economic, and disaster-risk indicators, making visualization literacy a genuinely practical civic as well as professional skill for Philippine data professionals. Effective dashboard design for these public-facing use cases carries particular responsibility, since a poorly designed public dashboard can misinform a broad public audience far beyond the harm a single internal organizational report might cause.

Disaster-risk dashboards specifically, given the Philippines' frequent exposure to typhoons and other natural hazards, illustrate the highest-stakes version of this responsibility: a visualization that understates risk through poor design choices, an inappropriately compressed color scale, an unclear severity legend, can directly and measurably affect how residents in an affected area respond to a warning. Data professionals working on public-facing hazard communication should treat the design principles in this chapter as safety-relevant, not merely aesthetic, given what is actually at stake when such a dashboard is misread.

CASE STUDY

Illustrative composite case.

A Philippine provincial health office's COVID-19 case dashboard initially displayed daily case counts on a bar chart with a truncated y-axis, making relatively small day-to-day fluctuations appear as dramatic spikes and causing repeated public alarm and confusion that generated significant call volume to the office. Redesigning the chart with a zero-based axis, alongside a rolling seven-day average line to smooth genuine day-to-day noise, produced a visualization that accurately reflected the actual trend and substantially reduced public confusion, without changing anything about the underlying data itself.

The health office's communications team later noted that the redesign's success depended on more than the technical axis correction alone; the rolling average line proved equally important, since even a correctly scaled zero-based axis still displayed enough natural daily noise, lower reporting on weekends, delayed results from certain testing sites, to invite misreading by a public audience unfamiliar with those operational patterns. Combining an honest scale with a smoothing technique addressed both problems at once, illustrating that honest visualization sometimes requires more than avoiding a single well-known distortion like axis truncation; it requires anticipating how an unfamiliar audience will actually read the chart.

13.9 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Choose chart type based on the specific analytical purpose, comparison, trend, relationship, or proportion, from Table 13.1, not by default or personal preference.
- Best Practice: Lead with the finding and its implication when communicating to non-technical audiences, reserving methodological detail for those who specifically request it.
- Common Pitfall: Truncating a bar chart's axis in a way that exaggerates a real but modest difference, whether intentionally or through inattention to default software settings.
- Emerging Trend: Interactive dashboards allowing audiences to explore data themselves are increasingly supplementing, though rarely fully replacing, static reports for ongoing monitoring use cases.

13.10 Summary and Key Takeaways

Data visualization is an integral part of the analytical process, not a cosmetic final step, and choosing the correct chart type and applying honest design principles, zero-based axes, purposeful color, clear labeling, determines whether a visualization reveals genuine insight or

actively misleads its audience. Effective communication to non-technical audiences leads with findings and implications, not methodology. This chapter closes the analytics portion of Part VI; Chapter 14 closes the book with a look toward the field's future.

- Common Pitfall: Truncating a bar chart's axis by default software setting, unintentionally exaggerating a modest real difference.
- Common Pitfall: Using a pie chart for a trend or a many-category comparison it is poorly suited to communicate.
- Common Pitfall: Leading a non-technical presentation with methodology rather than the finding and its business implication.

KEY TAKEAWAYS

- Chart type should match analytical purpose: line for trends, bar for comparison, scatter for relationships, pie only for simple proportions.
- Zero-based axes and purposeful color use distinguish honest visualization from misleading visualization.
- Effective communication to non-technical audiences leads with findings and implications, not methodology.
- Visualization design carries Data Privacy Act implications when aggregation choices risk re-identifying individuals.

13.11 Key Terms, Reflection, and Discussion

Data Visualization. Data-Ink Ratio. Dashboard. Chart Junk. Zero-Based Axis.

- Reflection: For the analytics questions you formulated in Chapter 12's Data Project Build step, which chart type from Table 13.1 would best communicate each finding?

- Discussion: Is a truncated axis ever justifiable, for example, to show meaningful variation in a naturally narrow-range metric, or should zero-based axes be treated as an inviolable rule?
- Discussion: The health dashboard case study needed both an honest axis and a smoothing technique to avoid misleading the public. Can you think of a visualization that would still mislead even with both of these corrections applied?

13.12 Activity and Data Project Build, Step 13

DATA PROJECT BUILD, STEP 13

Sketch or describe a dashboard concept for your Data Project Build dataset, specifying at least three visualizations, their chart types, and the specific finding or question each is designed to communicate to a non-technical audience.

13.13 Chapter Self-Check

CHAPTER SELF-CHECK

1. Edward Tufte's data-ink ratio principle argues that visualizations should:

- a) Maximize decorative elements b) Maximize the proportion of ink devoted to actual data c) Always use pie charts d) Avoid color entirely

2. A truncated bar chart axis most commonly causes:

- a) Faster rendering b) Exaggeration of real but modest differences c) Improved accuracy d) Automatic data cleansing

3. Which chart type is best suited to showing proportional parts of a single whole?

- a) Line chart b) Scatter plot c) Pie chart d) Bar chart with truncated axis

4. When communicating to a non-technical audience, best practice is to:

- a) Lead with detailed methodology b) Lead with the finding and its business implication c) Use the most complex visualization available d) Avoid visualizations entirely

5. Chart junk refers to:

- a) Corrupted data files b) Unnecessary decorative elements that do not convey data c) A type of bar chart d) A data storage format

13.14 References

- Tufte, E. R. (2001). The visual display of quantitative information (2nd ed.). Graphics Press.
- Knaflic, C. N. (2015). Storytelling with data: A data visualization guide for business professionals. Wiley.
- Few, S. (2012). Show me the numbers: Designing tables and graphs to enlighten (2nd ed.). Analytics Press.
- Evergreen, S. D. (2016). Effective data visualization: The right chart for the right data. SAGE Publications.

End of Chapter 13. Continue to Chapter 14: Emerging Trends and Technologies in Data Management

EMERGING TRENDS AND TECHNOLOGIES IN DATA MANAGEMENT

Capstone, Completing and Presenting Your Data Project Build

“The specific tools this chapter describes will change. The discipline the rest of this book taught you will not.”

14.0 Chapter Overview

This capstone chapter closes the book by looking forward to the technologies and trends reshaping data management, from artificial intelligence to evolving privacy regulation, while bringing together every prior Data Project Build step into a final, complete portfolio project. It also brings together the disciplines from every prior chapter into a closing reflection on what data management as a profession actually asks of you.

14.1 Learning Outcomes

- LO 14.1 Explore emerging trends and technologies shaping the future of data management.
- LO 14.2 Evaluate the implications of artificial intelligence for data management practice.
- LO 14.3 Synthesize the complete Data Project Build into a presentable data management portfolio.

14.2 Introduction

Every chapter in this book has built toward this moment: a governed, collected, stored, quality-assured, processed, secured, and analyzed dataset, ready to support a real organizational decision. This final chapter addresses two remaining questions. First, what is changing in the field you are entering, so that the specific tools and platforms you have learned do not quietly become the limit of your understanding as they inevitably evolve. Second, how do you bring

together everything built across the preceding thirteen chapters into a single, coherent portfolio project you can present?

It is worth being direct about what this final chapter is and is not. It is not a guarantee that the specific tools and vendor names mentioned here will still be current by the time you enter the workforce, and treating them as the point would miss the chapter's actual purpose. What this chapter can more honestly offer is a way of thinking about change in the field, durable principles applied to a shifting technical landscape, that should remain useful long after any specific platform mentioned here has been superseded.

14.3 Historical Background

Data management as a discipline has undergone several distinct technological eras across the roughly five decades since Codd's relational model, each expanding what organizations could feasibly do with data and each raising, not lowering, the importance of the governance and quality foundations covered in Parts I and III of this book. The relational database era gave organizations reliable, structured storage. The big data era, driven by tools such as Hadoop and Spark, enabled organizations to process unstructured data at a massive scale. The current era, driven substantially by artificial intelligence and machine learning, is enabling organizations to automatically extract predictive and generative insights from that data at a speed and scale no earlier era could match.

The pattern across all three eras is worth stating plainly, since it is the closing argument this entire book has been building toward. Each new technological era did not diminish the importance of the discipline taught in earlier chapters; it increased it. Bigger, faster, more automated tools amplify the consequences of both good and poor underlying data management practice, meaning the governance, quality, and security foundations built in Parts I, III, and V of this book matter more, not less, in an AI-driven era than they did in the relational database era Codd's original work made possible.

14.4 Definitions

KEY TERMS

Artificial Intelligence (AI) in Data Management: The application of machine learning and related techniques to automate data management tasks, including quality detection, classification, and governance enforcement.

Data Mesh is an emerging organizational and architectural approach that distributes data ownership across individual business domains rather than centralizing it in a single team.

MLOps is the practice of applying disciplined operational and governance practices to machine learning models, analogous to how DevOps applies them to software.

Data Fabric is an architectural approach that provides a unified layer of data access and integration across an organization's distributed data sources.

14.5 Emerging Trends Shaping the Field




Figure 14.1. Emerging Trends in Data Management.

Artificial intelligence is reshaping data management practice in two directions simultaneously: AI is increasingly used as a tool within data management itself, automating

data quality detection, classification, and even governance policy enforcement at a scale manual review cannot match, while at the same time, AI systems themselves depend more critically than ever on the data management discipline this book has taught, since a flawed underlying dataset produces a flawed model regardless of how sophisticated that model's architecture is.

Several architectural and organizational trends are also reshaping how data management is practiced at scale. Data mesh challenges the traditionally centralized data team model by distributing data ownership to individual business domains, each responsible for the quality and governance of its own data products, echoing this book's Chapter 2 emphasis on named, accountable data stewardship, applied at a larger organizational scale. MLOps extends the disciplined lifecycle and governance thinking that this book has applied throughout to data, collection, quality, security, and the machine learning models increasingly built on that data. Regulatory evolution continues as well, with data protection frameworks worldwide, and the Philippines' own Data Privacy Act enforcement practice, continuing to mature and tighten.

14.6 Table 14.1, Emerging Trends at a Glance

Trend	What It Changes	Relevance to This Book
AI-Assisted Data Quality	Automates detection of quality defects at scale	Extends Chapters 6 to 7's cleansing and validation practice
Data Mesh	Distributes data ownership to business domains	Extends Chapter 2's governance and stewardship principles
MLOps	Applies lifecycle discipline to machine learning models	Extends Chapter 3's lifecycle thinking to models, not just data
Evolving Privacy Regulation	Tightens and expands data protection obligations.	Extends Chapter 11's compliance practice into an ongoing, not static, obligation

Table 14.1. Emerging trends reshaping data management practice.

14.7 Government Policy and the Future of Compliance

Philippine data protection regulation should be expected to continue evolving, likely toward more specific guidance on AI-driven automated decision-making, cross-border data transfer, and sector-specific requirements, mirroring the trajectory of data protection law in more mature regulatory jurisdictions. Data management professionals entering the field now

should expect the compliance landscape covered in Chapter 11 to keep changing throughout their careers, making the underlying principles, transparency, legitimate purpose, and proportionality more durable and valuable to internalize than any specific current rule.

14.8 Philippine Context and Emerging Opportunity

PHILIPPINE CONTEXT

The Philippines' growing business process outsourcing and shared services sector, alongside expanding government open data initiatives and a maturing private-sector data protection compliance landscape, is generating sustained demand for data management professionals who understand the full discipline this book covers, not narrow technical skills in isolation. Graduates of programs such as the MDM who can speak credibly to governance, quality, security, and analytics together, rather than only one specialty, are positioned to lead this growing field rather than staff it.

This breadth is worth naming as a genuine competitive advantage in the Philippine job market, specifically, since many existing practitioners in Philippine organizations arrived at data-related roles through a single narrow specialty, a database administrator who moved into governance, a security specialist who moved into compliance, without the deliberately integrated foundation this book has built across fourteen chapters. Graduates who can move fluently across governance, quality, security, and analytics in a single conversation, rather than deferring to a different specialist for each, offer organizations a rarer and correspondingly more valuable combination of skills than the current Philippine talent market typically produces.

14.9 Case Study: The Complete Data Project Build

CASE STUDY

Illustrative composite case, synthesizing the full book.

An MDM graduate student begins this course with a publicly available dataset on regional agricultural production, selected in Chapter 1, with only a general interest in

food security policy. By Chapter 2, the dataset has a defined data owner and steward and a documented data dictionary. By Chapter 3, a lifecycle map specifies its retention as indefinite public research data given its non-personal nature. By Chapter 5, it is recommended for storage in a cloud-based data warehouse suited to its structured, moderate-volume format. By Chapter 7, a cleansing pass has resolved inconsistent regional naming conventions across years of data. By Chapter 9, an automated pipeline refreshes the dataset from its public source quarterly. By Chapter 11, the student has confirmed that the dataset contains no personal data and documented this determination. By Chapter 13, three dashboard visualizations reveal a diagnostic finding: production declines in specific regions correlate with a specific irrigation infrastructure gap, a gap that the original descriptive reports alone had never surfaced. The student's final presentation traces every one of these steps back to a specific chapter's framework, exactly the assembled, evidence-based credibility this book has aimed to teach throughout.

It is worth being explicit about what made this particular student's presentation credible to its evaluating panel, since it is exactly the pattern this book has argued for throughout. Every claim in the final presentation traced to a specific, earlier piece of documented work: the storage recommendation came from Chapter 5's architecture criteria, not from whichever platform the student happened to be personally familiar with; the privacy determination came from Chapter 11's actual assessment process, not from an unexamined assumption that public data is automatically exempt from consideration; the diagnostic finding came from Chapter 13's honest visualization principles, not from a chart optimized to look impressive. None of these individual pieces was extraordinary on its own; their credibility came from being assembled honestly, in the sequence this book has taught, rather than compiled under presentation-deadline pressure.

14.10 Best Practices, Current Issues, and Emerging Trends

- Best Practice: Treat specific tools and platforms as replaceable, and the underlying principles, governance, quality dimensions, and lifecycle thinking as the durable core of your professional competence.
- Best Practice: Revisit your organization's, or your own, data management practices periodically as both technology and regulation evolve, rather than treating any single implementation as permanent.
- Current Issue: The pace of AI adoption in many organizations is outstripping the data governance maturity needed to deploy it responsibly, a gap that this book's integrated treatment of both topics is meant to help close.

14.11 Summary and Key Takeaways

This final chapter, and this book, close on a deliberately forward-looking note: the specific technologies, distributed processing frameworks, cloud platforms, and AI tools will continue to change throughout your career, but the underlying discipline this book has taught, governance, quality, security, and analytics built on trustworthy data, does not expire when any particular tool is superseded by the next one.

- Common Pitfall: Deploying AI-driven data tools without the governance and quality maturity needed to trust their output.
- Common Pitfall: Treating a specific current platform or vendor as a permanent skill, rather than an application of durable underlying principles.
- Common Pitfall: Assuming evolving regulation will remain static, rather than periodically revisiting compliance practices as the legal landscape matures.

KEY TAKEAWAYS

- AI is reshaping data management both as a tool within the discipline and as a practice that depends more critically than ever on the discipline's foundations.
- Data mesh, MLOps, and evolving privacy regulation extend this book's governance and lifecycle principles to new organizational and technical contexts.
- Philippine demand for data management professionals rewards understanding the full discipline, not narrow specialization alone.
- Durable professional competence rests on principles, not on any single current tool or platform.

14.12 Key Terms, Reflection, and Discussion

AI in Data Management. Data Mesh. MLOps. Data Fabric.

- Reflection: Looking back across all fourteen chapters, which single framework from this book do you expect to use most often in your future career?
- Discussion: Is data mesh's distributed ownership model, extending accountability to individual business domains, a genuine improvement over centralized data governance, or does it risk recreating the data silo problems Chapter 6 described?
- Discussion: Looking back across all fourteen chapters, which single framework from this book do you expect to use most often in your future career, whether as a data professional, a manager, or a policymaker?

14.13 Final Activity: Data Project Build, Step 14 (Capstone Presentation)

DATA PROJECT BUILD, STEP 14 (FINAL)

Compile all thirteen prior Data Project Build steps into a complete data management portfolio document, and deliver a final presentation tracing your dataset through the full lifecycle this book has covered: governance and ownership, collection and storage, quality and cleansing, processing, security and privacy, and analytics and visualization. Evaluation will follow a rubric assessing governance clarity, quality rigor, security and compliance awareness, and clarity of the final analytical insight presented.

14.14 Chapter Self-Check

CHAPTER SELF-CHECK

1. Data mesh primarily addresses which earlier chapter's concept, at a larger organizational scale?

- a) Chapter 13's visualization principles b) Chapter 2's governance and stewardship c) Chapter 9's automation tools only d) Chapter 4's collection methods only

2. MLOps applies disciplined lifecycle and governance practices to:

- a) Physical office security b) Machine learning models c) Marketing budgets d) Employee onboarding

3. According to this chapter, what remains durable even as specific tools change?

- a) The exact software vendor used b) The underlying principles of governance, quality, security, and analytics c) The specific programming language used d) Nothing; everything must be relearned

4. AI's relationship to data management, per this chapter, is best described as:

- a) AI replaces the need for data management entirely. b) AI is both a tool within data management and dependent on its foundations. c) AI is unrelated to data quality d) AI eliminates all data privacy concerns

5. The Philippine sector cited as generating sustained demand for data management professionals includes:

- a) Only agriculture b) Business process outsourcing and shared services c) Only government d) None; demand is declining

14.15 References

- DAMA International. (2017). DAMA-DMBOK: Data management body of knowledge (2nd ed.). Technics Publications.
- Dehghani, Z. (2022). Data mesh: Delivering data-driven value at scale. O'Reilly Media.
- Republic Act No. 10173 (2012). Data Privacy Act of 2012.
- National Privacy Commission. Emerging technology and privacy guidance (latest edition).

End of Chapter 14. End of Book. Congratulations on completing your Data Project Build.

List of Abbreviations

CIA	Confidentiality, Integrity, Availability (security triad)
DAMA	Data Management Association International
DMBOK	Data Management Body of Knowledge
DPO	Data Protection Officer
ELT	Extract, Load, Transform
ETL	Extract, Transform, Load
IDPS	Intrusion Detection and Prevention System
MDM	Master in Data Management (program)
MFA	Multi-Factor Authentication
MLOps	Machine Learning Operations
NPC	National Privacy Commission
NoSQL	Not Only SQL (non-relational database)
PIA	Privacy Impact Assessment
PSA	Philippine Statistics Authority
RA	Republic Act
RACI	Responsible, Accountable, Consulted, Informed
RBAC	Role-Based Access Control
SQL	Structured Query Language

List of Figures

Figure 1.1. From Data to Insight

Figure 2.1. Data Governance Roles

Figure 3.1. The Data Lifecycle

Figure 4.1. Structured vs. Unstructured Data

Figure 5.1. Centralized vs. Distributed Storage

Figure 6.1. Core Data Quality Dimensions

Figure 7.1. The Cleansing Cycle: Detect, Diagnose, Edit

Figure 8.1. The ETL Pipeline

Figure 9.1. Pipeline Orchestration

Figure 10.1. The CIA Triad

Figure 11.1. Data Subject Rights

Figure 12.1. The Four Types of Data Analytics

Figure 13.1. Choosing the Right Chart Type

Figure 14.1. Emerging Trends in Data Management

List of Tables

Table 1.1. Data vs. Information vs. Insight

Table 2.1. Data Governance Roles at a Glance

Table 3.1. The Six Lifecycle Stages

Table 4.1. Structured vs. Unstructured Data

Table 5.1. Choosing a Storage System

Table 6.1. Core Data Quality Dimensions

Table 7.1. Cleansing Techniques and What They Fix

Table 8.1. Data Loading Strategies Compared

Table 9.1. When to Automate

Table 10.1. Core Data Security Measures

Table 11.1. Data Subject Rights at a Glance

Table 12.1. The Four Types of Data Analytics

Table 13.1. Matching Chart Type to Purpose

Table 14.1. Emerging Trends at a Glance

Table B.1. Consolidated Philippine Legal and Regulatory Reference

Table C.1. Index of Major Frameworks Introduced in This Book

Table C.2. Course Outcome Alignment

Glossary

The following terms are introduced and defined throughout the fourteen chapters of this book. Each entry lists the chapter in which the term is first defined; refer to that chapter's Definitions section (Standard Chapter Template, item 5) for full context.

Data Management (*Ch. 1*) The practices, roles, and technologies used to collect, store, secure, govern, and derive value from data throughout its useful life.

Data Governance (*Ch. 1*) The framework of policies, roles, and decision rights determining how data is managed and by whom.

Data Steward (*Ch. 2*) An individual accountable for the quality, definition, and appropriate use of a specific data domain.

Data Owner (*Ch. 2*) The individual with ultimate decision authority over a given data domain.

Data Dictionary (*Ch. 2*) A centralized, documented reference defining every significant data field an organization uses.

Data Lifecycle (*Ch. 3*) The stages data passes through, from creation or collection to eventual archival or deletion.

Data Retention Policy (*Ch. 3*) A documented rule specifying how long a category of data must or may be kept before disposal.

Primary Data (*Ch. 4*): Data collected directly from a source or activity for a specific purpose.

Structured Data (*Ch. 4*) Data organized according to a predefined model or schema, typically rows and columns.

Unstructured Data (*Ch. 4*): Data that does not follow a predefined model or schema, such as text, images, or video.

Data Warehouse (*Ch. 5*) A centralized repository optimized for storing structured, processed data for analysis and reporting.

Data Lake (*Ch. 5*) A storage repository holding vast amounts of raw data, structured and unstructured, in native format.

Data Quality (*Ch. 6*) The degree to which data is fit for the person or process using it.

Data Profiling (*Ch. 6*) The systematic analysis of a dataset's structure, content, and quality characteristics.

Data Cleansing (*Ch. 7*): Detecting and correcting or removing corrupt, inaccurate, incomplete, or duplicate records.

Data Validation (*Ch. 7*): Checking data against defined rules at the point of entry to prevent invalid data from being accepted.

ETL (*Ch. 8*): Extract, Transform, Load; the standard process for moving data from sources into a usable destination.

Data Pipeline (*Ch. 8*): The automated sequence of processes that moves data from source to destination.

Pipeline Automation (*Ch. 9*): Configuring a data pipeline to run and handle routine errors without manual intervention.

CIA Triad (*Ch. 10*): The foundational security model of Confidentiality, Integrity, and Availability.

Encryption (*Ch. 10*): Converting data into a coded format that is unreadable without a decryption key.

Personal Data (*Ch. 11*): Any information relating to an identified or identifiable individual.

Data Subject Rights (*Ch. 11*): The rights the Data Privacy Act grants individuals over their own personal data.

Data Analytics (*Ch. 12*): The process of examining data to draw conclusions and support decision-making.

Predictive Analytics (*Ch. 12*): Analytics using historical data and models to forecast what might happen.

Data Visualization (*Ch. 13*): The graphical representation of data designed to make patterns and trends apparent.

Dashboard (*Ch. 13*): A visual display consolidating multiple related metrics into a single view.

Data Mesh (*Ch. 14*): An emerging approach distributing data ownership to individual business domains.

Appendix A: The Complete Data Project Build Worksheet Packet

This appendix consolidates all fourteen Data Project Build steps referenced throughout the book into a single worksheet packet. Students may use this as a standalone companion document, completing each step as they progress through the corresponding chapter, and compiling the results into the final capstone portfolio and presentation required in Chapter 14.

Step 1 (Ch. 1) Dataset selection and organizational question statement.

Step 2 (Ch. 2) Governance role mapping and initial data dictionary entries.

Step 3 (Ch. 3) Lifecycle map and retention period justification.

Step 4 (Ch. 4) Source and type classification, plus collection statement.

Step 5 (Ch. 5) Recommended storage system and architecture.

Step 6 (Ch. 6) Data quality assessment across the six core dimensions.

Step 7 (Ch. 7) Cleansing and validation plan with two defects addressed.

Step 8 (Ch. 8) ETL plan with extraction, transformation, and loading strategy.

Step 9 (Ch. 9) Automation justification and monitoring plan.

Step 10 (Ch. 10) Security assessment and technical and procedural controls.

Step 11 (Ch. 11) Privacy Impact Assessment summary.

Step 12 (Ch. 12) Four analytics-type questions and selected priority insight.

Step 13 (Ch. 13) Dashboard concept with at least three justified visualizations.

Step 14 (Ch. 14) Final compiled data management portfolio and capstone presentation.

Appendix B: Philippine Legal and Regulatory Quick Reference

The following Philippine laws and government bodies are referenced throughout this book and are consolidated here for quick reference.

Law / Program	Relevance	Chapter
RA 10173: Data Privacy Act	Core personal data protection framework	1, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13
RA 10175: Cybercrime Prevention Act	Criminalizes unauthorized access and data interference	10
National Privacy Commission (NPC)	Administers and enforces the Data Privacy Act	1, 2, 3, 6, 7, 11
DAMA-DMBOK	International data management body of knowledge standard	1, 2, 6, 14

Table B.1. Consolidated Philippine legal and regulatory reference.

Appendix C: Index of Key Frameworks and Course Outcome Alignment

This appendix indexes the major named frameworks and theories introduced throughout the book, and maps each of the book's six Parts to the MDMN 202 Course Outcomes (CO1–CO7) established in the official course syllabus.

Framework / Theory	Originator(s)	Chapter
Relational Model	Codd (1970)	1, 5
DAMA-DMBOK Governance Framework	DAMA International (2017)	1, 2
RACI Structure	Standard governance practice	2
Data Lifecycle Model	Records management tradition	3
Primary/Secondary Source Distinction	Statistical survey methodology	4
Data Quality Dimensions	Wang & Strong (1996); Loshin (2001)	6
Detect, Diagnose, Edit Cycle	Van den Broeck et al. (2005)	7
ETL / ELT Process	Data warehouse tradition	8
Orchestration and Automation Principles	Distributed systems practice	9
CIA Triad	Classical information security theory	10
Data Privacy Act Core Principles	RA 10173 (2012)	11
Four Types of Analytics	Business intelligence tradition	12
Data-Ink Ratio	Tufte (2001)	13
Data Mesh	Dehghani (2022)	14

Table C.1. Index of major frameworks introduced in this book.

Course Outcome	Description	Primary Parts/Chapters
CO1	Solid foundation in principles and practices of data management	Part I (Ch. 1–3)

Course Outcome	Description	Primary Parts/Chapters
CO2	Apply data collection and storage techniques for structured and unstructured data.	Part II (Ch. 4–5)
CO3	Ensure data quality and integrity using appropriate strategies and tools	Part III (Ch. 6–7)
CO4	Implement data governance policies and comply with regulatory standards	Parts I, IV, V (Ch. 2, 8–9, 11)
CO5	Secure data against common threats and safeguard user privacy	Part V (Ch. 10–11)
CO6	Analyze and visualize data to derive valuable insights.	Part VI (Ch. 12–13)
CO7	Explore emerging trends and technologies shaping the future of data management.	Chapter 14

Table C.2. Course Outcome alignment, per the official MDMN 202 syllabus.

Appendix D: Suggested Course Calendar

The following 14-week calendar maps one chapter to one instructional week and is suitable for a standard one-semester MDMN 202 offering. Institutions following the official syllabus's 54-hour, module-based structure may instead pace two chapters per module block, aligning with the midterm and final examination points indicated in the course learning plan.

Week	Chapter Topic	Deliverable
Week 1	Chapter 1: Understanding Data Management	Data Project Build Step 1 due
Week 2	Chapter 2: Key Components and Data Governance	Data Project Build Step 2 due
Week 3	Chapter 3: The Data Lifecycle	Data Project Build Step 3 due
Week 4	Chapter 4: Data Sources, Types, and Collection Methods	Data Project Build Step 4 due
Week 5	Chapter 5: Data Storage Systems and Architectures	Data Project Build Step 5 due
Week 6	Chapter 6: Dimensions and Challenges of Data Quality	Data Project Build Step 6 due
Week 7	Chapter 7: Data Cleansing and Validation Techniques	Data Project Build Step 7 due
Week 8	Chapter 8: ETL Processes and Data Transformation Techniques	Data Project Build Step 8 due
Week 9	Chapter 9: Data Processing Tools and Automation	Data Project Build Step 9 due
Week 10	Chapter 10: Data Security Threats, Encryption, and Access Control	Data Project Build Step 10 due
Week 11	Chapter 11: Data Privacy, Compliance, and Regulatory Standards	Data Project Build Step 11 due
Week 12	Chapter 12: Data Analytics Fundamentals and Types	Data Project Build Step 12 due
Week 13	Chapter 13: Data Visualization and Communicating Insights	Data Project Build Step 13 due
Week 14	Chapter 14: Emerging Trends and Technologies in Data Management	Data Project Build Step 14 due

Table D.1. Suggested 14-week course calendar.

Suggested Further Reading

Students seeking deeper treatment of specific topics covered in this book may consult the following foundational texts, each cited within the relevant chapter:

- DAMA International. (2017). DAMA-DMBOK: Data management body of knowledge (2nd ed.). Technics Publications. The definitive data governance and management reference (Chapters 1–2, 6).
- Silberschatz, A., Korth, H. F., & Sudarshan, S. (2020). Database system concepts (7th ed.). McGraw-Hill, Comprehensive treatment of database storage and architecture (Chapter 5).
- Kimball, R., & Ross, M. (2013). The data warehouse toolkit (3rd ed.). Wiley., The definitive ETL and dimensional modeling reference (Chapter 8).
- Andress, J. (2014). The basics of information security (2nd ed.). Elsevier, Foundational treatment of the CIA triad and security practice (Chapter 10).
- Tufte, E. R. (2001). The visual display of quantitative information (2nd ed.). Graphics Press, Foundational data visualization design theory (Chapter 13).
- Knafllic, C. N. (2015). Storytelling with data: A data visualization guide for business professionals. Wiley, Practical, business-oriented visualization guidance (Chapter 13).

Notes on Sources and Academic Integrity

All case studies in this book, labeled “illustrative composite case,” are constructed scenarios designed to demonstrate the application of a chapter's frameworks. They do not depict real, identifiable individuals, organizations, or events, and any resemblance to actual organizations is coincidental. Where this book references real laws, government agencies, or well-documented technical standards by name, such as cited Republic Acts or the DAMA-DMBOK, these references are factual and drawn from publicly available information, distinct from the illustrative composite cases used for pedagogical demonstration elsewhere in the same chapters.

All theoretical frameworks, models, and direct citations are attributed to their original scholarly sources in each chapter's References section, formatted in APA 7th edition style. Students using this book for their own Data Project Build assignments should apply the same citation discipline to any external sources, data, or frameworks they draw upon beyond this textbook, consistent with their institution's academic integrity policy.

This edition has been formatted to the Philippine Technical Institute Inc. house style, including its citation and reference conventions. A dedicated reference-accuracy audit, verifying every citation against its source and confirming full APA 7th edition compliance, remains a planned next step in this book's production process and has not yet been completed as of this printing. Instructors and reviewers using this edition for coursework should treat citations as provisional pending that audit, and are welcome to report any citation concerns directly to the author.

About the Author



Dr. Cristeta G. Tolentino, DIT, is the Dean of the College of Computing Sciences (CCS) of PSU Lingayen Campus and a faculty member of the Graduate School of Pangasinan State University – Open University Systems (PSU-OUS). She earned her Bachelor of Science in Information Technology (BSIT), Master of Science in Information Technology (MSIT), and Doctor in Information Technology (DIT), equipping her with extensive expertise in computing, information systems, and technology-driven education.

Her professional interests include computing education, information systems, curriculum development, research, and digital transformation. She is an active member of the Philippine Society of Information Technology Educators (PSITE), the Philippine Computer Society (PCS), and the Council of Information Technology Deans and Heads (CITEDH), demonstrating her commitment to advancing information technology education and professional excellence.