

National Data Management Guidelines

Data Storage and Operations



Table of Contents

1. Data Storage and Operations Domain Introduction

- 1.1 Domain Overview
- 1.2 Target Audience
- 1.3 Abbreviations and Glossary
- 1.4 Readability Guide
- 1.5 Stakeholders Involved
- 1.6 Key Success Factors
- 1.7 Compliance with National Data Standards
- 1.8 Data Storage and Operations Process Map

2. Data Storage & Operations Policy

- Overview
- 2.1 Data Storage & Retention Policy Creation
- 2.2 Data Storage & Retention Policy Publication

3. Data Storage & Operations Planning

- Overview
- 3.1 Baseline Data Storage Capability Assessment
- 3.2 Data Storage & Usage Requirements Forecasting
- 3.3 Target Data Storage Infrastructure Architecture
- 3.4 Data Storage & Operations Roadmap

4. Data Storage & Operations Implementation

- Overview
- 4.1 Database Access Control
- 4.2 Data Backup & Recovery
- 4.3 Disaster Recovery & Business Continuity
- 4.4 Data Storage & Operations Roadmap Activation

5. Data Storage & Operations Performance & Change Management

- Overview
- 5.1 Data Operations SLAs
- 5.2 Database Performance Monitoring
- 5.3 Database Configuration Management
- 5.4 Database Maintenance and Updates
- 5.5 QDTI

6. Appendix

- 6.1 Templates and Frameworks

1. Data Storage and Operations Domain Introduction



1.1 Domain Overview

The Data Storage & Operations domain focuses on ensuring the secure, scalable, and resilient storage, access, and recovery of data across the Entity's digital infrastructure—supporting business continuity, and data integrity.

Through structured assessment, planning, implementation, and performance management, entities can modernize their storage environments, align with Qatar National Data Standards, and maintain operational readiness in the face of evolving business and technology demands.

This domain enables entities to transition from fragmented, legacy storage practices toward a unified, cloud-aligned, policy-driven practice — ensuring consistent service delivery, reduced risk of data loss, and seamless recoverability in the event of failure or disruption.

Organized across 3 phases — Planning, Implementation, and Operation— the Data Storage & Operations domain provides a comprehensive approach to assess existing capabilities, design future-ready infrastructure, operationalize best practices, and continuously enhance system performance and reliability.

By following this approach, entities ensure their storage and operations environments are technically sound, strategically aligned, and fully compliant with national standards — while supporting digital transformation, data availability, and long-term infrastructure sustainability.

Key Components of the Data Storage and Operations Domain

Data Storage & Operations Policy

This phase establishes the foundational policy for managing the secure, efficient, and compliant storage and retention of data across the Entity. It ensures all storage activities are guided by consistent policy direction, aligned with national standards, and endorsed by internal governance bodies.

- **Data Storage & Retention Policy Creation:** Defines Entity-wide principles, rules, and roles for storing, retaining, and disposing of data based on classification, business needs, and regulatory requirements.
- **Data Storage & Retention Policy Publication:** Guides the internal communication, endorsement, and awareness of the approved policy to promote adoption and accountability.

Data Storage & Operations Planning

This phase provides a structured foundation for understanding the Entity's current storage landscape, projecting future needs, designing a future-ready architecture, and organizing implementation through a sequenced roadmap. It ensures the Entity's storage and operations decisions are data-driven, cloud-aligned, and compliant with national standards—enabling cost-efficiency, performance, and long-term scalability.

- **Baseline Data Storage Capability Assessment:** Establishes a clear inventory and performance profile of existing storage systems, infrastructure utilization, and migration readiness.
- **Data Storage & Usage Requirements Forecasting:** Projects short- and long-term storage growth, backup needs, and resource requirements to support future operations.
- **Target Data Storage Infrastructure Architecture:** Designs the Entity's future-state architecture to support flexible, scalable, and regulations-aligned data storage across cloud and hybrid environments.
- **Data Storage & Operations Roadmap:** Defines the prioritized initiatives, resources, timelines, and governance mechanisms required to implement the storage strategy effectively.

Data Storage & Operations Implementation

This phase operationalizes the Entity's data storage and operations by executing secure access controls, reliable backup procedures, disaster recovery plans, and the approved roadmap. It ensures that planning outputs are transformed into functioning solutions that maintain service continuity, protect data integrity, and comply with national standards.

- **Database Access Control:** Implements role-based access control (RBAC) to manage and audit who can access databases, what permissions they have, and how access is granted or revoked.
- **Data Backup & Recovery:** Establishes system-specific backup strategies and recovery procedures to ensure critical data can be restored within acceptable timeframes.
- **Disaster Recovery & Business Continuity:** Documents recovery priorities and executable procedures through structured playbooks that enable rapid service restoration during disruptions.
- **Data Storage & Operations Roadmap Activation:** Launches execution of prioritized initiatives, monitors progress through milestones, and ensures alignment with governance processes and national priorities.

Data Storage & Operations Performance & Change Management

This phase ensures the ongoing reliability, efficiency, and compliance of the Entity's storage and database operations through structured monitoring, configuration governance, update management, and compliance assessment. It enables continuous improvement by detecting performance issues early, managing system changes responsibly, and maintaining audit-ready records of all configurations and software assets.

- **Data Operations SLAs:** Defines service-level expectations for system availability, performance, and recoverability, and establishes mechanisms for monitoring and enforcing them.
- **Database Performance Monitoring:** Tracks key performance indicators such as query speed, uptime, and resource utilization to identify and resolve bottlenecks proactively.
- **Database Configuration Management:** Documents and governs database settings to ensure consistency across environments, traceability of changes, and compliance with policy.
- **Database Maintenance and Updates:** Applies vendor-recommended patches, tracks software assets, and documents justifications for deferred updates to maintain system security and supportability.
- **QDTI:** Evaluating the Entity's adherence to the National Data Storage and Operations Standards as part of the broader QDTI framework.


1.2 Target Audience

This handbook is intended for the Data Management and Security Director. It is also suitable for team leads and other individuals responsible for data storage, infrastructure operations, and database management. The handbook serves as a resource, offering practical guidance to help establish the essential components of the Entity's data storage and operations program—setting the foundation for resilient, secure, and policy-aligned data infrastructure.

1.3 Abbreviations and Glossary

Abbreviation/Term	Definition / Long Name
AD	Active Directory
API	Application Programming Interface – a method for systems to exchange data and functionality
Backup Frequency	How often database backups are taken (e.g., daily, weekly)
BCP	Business Continuity Plan – a documented strategy to continue operations during disruptions
CMDB	Configuration Management Database – repository that stores configuration details
Configuration Parameter	A defined setting in a database that governs behavior, performance, or security
DBMS	Database Management System – software used to store, retrieve, and manage data
DR	Disaster Recovery – the process to restore IT operations after a failure
DSO	Data Storage & Operations
IOPS	Input/Output Operations Per Second – measures storage performance
ITSM	IT Service Management – tools and processes for delivering IT services (e.g., ServiceNow)
KPI	Key Performance Indicator – a measurable value that indicates performance over time
LDAP	Lightweight Directory Access Protocol – a protocol for accessing directory services
Max Execution Time	The maximum time allowed for a specific transaction or query to complete
Patch	A software update issued by a vendor to fix bugs, improve performance, or address security
Query Execution Time	The time it takes for a database to process and return a query result
RPO	Recovery Point Objective – the maximum tolerable period of data loss
RTA	Recovery Time Actual – the actual time taken to restore service
RTO	Recovery Time Objective – the maximum tolerable downtime for a system
Sandbox	A testing or experimentation environment isolated from production
SIEM	Security Information and Event Management – tools that collect and analyze security data
Uptime	The percentage of time a system is operational and accessible
Version Control	The process of managing changes to software or configuration over time

1.4 Readability Guide



Note: Should an activity only contain one sub activity, this sub activity will be in the format of "1.1.A" instead of "1.1.1.A"

1.6 Stakeholders Involved

Below are the stakeholders and actors who contribute to the different processes outlined in the upcoming chapters.

Stakeholder	Role Description
Data Governance Office (DGO)	A Data Governance Office (DGO) is the centralized national body in the National Planning Council responsible for facilitating and coordinating data governance efforts within all associated government entities. It serves as the principal body overseeing the management of data as a strategic asset, ensuring that data-related processes align with the national objectives and comply with relevant policies and regulations.
Data Management & Security Director	The Data Management and Security Director is responsible for leading and overseeing the development and implementation of the data agenda within the government entity, ensuring alignment with strategic objectives and compliance with security regulations. This role promotes best practices in data management while ensuring robust security measures are in place to safeguard the Entity's data assets. The director bridges business and IT to support data-driven decision-making, elevate data as a strategic asset, and ensure data is managed and secured across its lifecycle. Additionally, the director fosters a culture of data security and literacy within the entity, supervises the value realization agenda for data, and ensures alignment with national security policies and standards, as well as the vision of executive leadership.
Data Management and Governance Section Head	The Data Management and Governance Section Head is responsible for leading, managing, and overseeing both data management and governance initiatives within the government entity. This role focuses on establishing and implementing best practices to ensure data governance and management align with the Entity's business requirements, strategic objectives, and compliance with national data standards. The Section Head collaborates with various teams across the Entity to develop and operationalize frameworks, policies, and processes for data governance and management. This includes improving data quality, accessibility, and protection, while also ensuring data lifecycle management practices are in place. The Section Head fosters a data-driven culture and ensures that data is treated as a strategic asset to drive decision-making and operational excellence.

Stakeholder	Role Description
Data Architecture and Analytics Section Head	The Data Architecture and Analytics Section Head oversees the data architecture and analytics agenda/initiatives within the government entity, ensuring alignment with strategic objectives and national data management standards. This role is responsible for designing and maintaining a secure, scalable data infrastructure that supports data integration, management, and accessibility, as well as Overseeing data modeling, database design, and system integration, developing advanced analytics, data visualization, and reporting solutions to enable informed decision-making. The Section Head collaborates with departments to identify key performance indicators (KPIs), analyze trends, and provide actionable insights that optimize processes, enhance service delivery, and support policy-making. By combining expertise in data architecture and analytics, this role drives a culture of data-driven innovation and operational excellence across the Entity.
Information Security and Risk Management Section Head	The Information Security and Risk Management Section Head is responsible for overseeing the Entity's cybersecurity agenda, integrating risk management and compliance into governance frameworks. This role ensures the protection of information assets by implementing robust security measures, monitoring risks, and fostering a culture of security awareness. The Section Head collaborates across departments to align security initiatives with organizational goals and regulatory requirements, ensuring resilience against cyber threats and compliance with national and international standards.
Data Governance Specialist	The Data Governance Specialist oversees the frameworks, policies, and procedures of governance and ensures their implementation to guarantee the effective management and utilization of data and resources within the Entity. This role includes developing and maintaining governance structures, leading compliance initiatives, and fostering a culture of accountability and transparency.
Data Quality Specialist	The Data Quality Specialist is responsible for overseeing data quality management processes within the entity. This role involves implementing data quality standards, monitoring data quality, and ensuring data consistency and effective usage across the Entity. The Data Quality Specialist collaborates with various teams to address data issues and enhance overall data quality.
Data Strategy Expert	The Data Strategy Expert is tasked with developing and promoting a data management and analytics strategy to achieve broad acceptance across the government entity. The expert identifies innovative ways to utilize data to achieve strategic objectives and enhance effective decision-making. Additionally, they introduce frameworks and practices to deliver accurate, reliable, and actionable data insights.

Stakeholder	Role Description
Data Owners	Data Owners are responsible for their department's (or sub-department's) data and its quality throughout its lifecycle across systems and ownership boundaries. They are accountable for supporting, guiding, reviewing, correcting, and/or approving the work of the data stewards managing their data to ensure their datasets are maintained to the highest possible standards.
Data Stewards	Data Stewards are responsible for identifying datasets related to relevant subdomains for inventory purposes and monitoring data quality metrics within and across entities. They are nominated as a business representative with a deep understanding of a specific data subdomain and can identify functional data requirements and needs.
Metadata Specialist	The Metadata Specialist is responsible for managing the Entity's data catalog and metadata assets to enhance data visibility, accessibility, and governance. This role focuses on developing and maintaining metadata standards, ensuring metadata quality, and supporting integration with the data catalog. The specialist plays a key role in enabling data discovery, supporting data-driven initiatives, and ensuring alignment with the Entity's data governance and management strategies.
Data Compliance Specialist	The Data Compliance Specialist ensures the Entity's data practices comply with national regulations, organizational policies, and international standards. This role manages data privacy, security, and governance compliance, conducts audits, mitigates non-compliance risks, and promotes awareness of regulatory requirements. The specialist works with stakeholders to embed compliance controls into data processes, ensuring adherence to legal and ethical standards.
Data Architect	The Data Architect is responsible for designing and implementing the Entity's overall data architecture, including logical and physical data models. This involves creating secure, scalable data structures that align with business needs, security requirements, and national and international data management and governance standards. The Architect ensures proper data integration, quality, and accessibility across systems, working with teams to maintain consistency and accuracy, and support data-driven decision-making to optimize data as a strategic asset.
Data Visualization Officer	The Data Visualization Officer is responsible for building dashboards and visual reports based on data, which are used to provide insights to leadership for decision-making.

Stakeholder	Role Description
Data Engineer	The Data Engineer transforms data into a format that can be easily analyzed and utilized. The Data Engineer works closely with other stakeholders to ensure data quality is maintained and implemented systems adhere to the highest possible standards.
Business Data Analyst	The Business Data Analyst is responsible for analyzing data from multiple sources to provide actionable insights that drive strategic and operational decisions within the Entity. This role involves designing reports and dashboards to effectively communicate findings, collaborating with cross-functional teams to ensure data accuracy and relevance, and aligning analytics with business objectives. The Analyst plays a key role in enabling data-driven decision-making and identifying opportunities to improve processes and outcomes.

1.6 Key Success Factors

Strategic Infrastructure Alignment and Leadership Support: Ensure that data storage and operations initiatives are tightly aligned with the Entity's digital transformation efforts and have the support of executive leadership. Strategic alignment facilitates resource prioritization, risk management, and long-term sustainability.

Resilient and Scalable Architecture: Design and implement a flexible storage and operations architecture that supports high availability, future data growth, and hybrid cloud adoption. Scalable, well-structured environments reduce downtime and support modernization efforts.

Regulations-Driven Governance and Compliance: Embed national standards, cloud-first policies, and classification-based controls into every layer of storage and operations. Strong governance ensures regulatory compliance, data protection, and operational consistency.

Proactive Monitoring and SLA Enforcement: Establish real-time monitoring systems and service-level tracking mechanisms to detect, report, and remediate performance issues quickly. SLA enforcement is essential for ensuring reliability and accountability.

Standardized Configuration and Change Control: Maintain consistency across environments through documented database configuration, version control, and structured change review. This reduces risks from misconfigurations and supports auditability.

Cross-Functional Collaboration and Ownership: Foster collaboration between IT operations, security, data governance, and business units to ensure shared ownership and alignment of responsibilities. Clear roles and coordinated execution drive operational maturity.

Continuous Improvement through Performance Insights: Use performance trends, SLA breaches, and system usage data to drive ongoing refinement of infrastructure, architecture, and policy. Data-driven decisions ensure systems stay optimized and responsive to changing needs.

The following chapters will provide detailed guidance in each of these areas, offering actionable steps to establish robust data storage and operations practices and drive operational efficiency within the Entity.

1.7 Compliance with National Data Standards

National Data Standards		QDKC
Control	Specification	Activity
DSO.1.1: Data Storage & Operations Policy	DSO.1.1.1 Data Storage & Retention Policy Creation F	2.1 Data Storage & Retention Policy Creation
	DSO.1.1.2 Data Monetization Capability Assessment F	2.2 Data Storage & Retention Policy Publication
DSO.1.2: Data Storage & Operations Planning	DSO.1.2.1 Baseline Data Storage Capability Assessment F	3.1 Baseline Data Storage Capability Assessment
	DSO.1.2.2 Data Storage and Usage Requirements Forecasting F	3.2 Data Storage and Usage Requirements Forecasting
	DSO.1.2.3 Target Data Storage Infrastructure Architecture F	3.3 Target Data Storage Infrastructure Architecture
	DSO.1.2.4 Data Storage & Operations Roadmap F	3.4 Data Storage & Operations Roadmap
DSO.2.1: Data Storage & Operations Implementation	DSO.2.1.1 Database Access Control F	4.1 Database Access Control
	DSO.2.1.2 Data Backup & Recovery F	4.2 Data Backup & Recovery
	DSO.2.1.3 Disaster Recovery & Business Continuity F	4.3 Disaster Recovery & Business Continuity
	DSO.2.1.4 Data Storage & Operations Roadmap Activation A	4.4 Data Storage & Operations Roadmap Activation
DSO.3.1: Data Storage & Operations Performance and Change Management	DSO.3.1.1 Data Storage SLAs A	5.1 Data Storage SLAs
	DSO.3.1.2 Data Storage SLAs VA	
	DSO.3.1.3 Database Performance Monitoring A	5.2 Database Performance Monitoring
	DSO.3.1.4 Database Configuration Management A	5.3 Database Configuration Management
	DSO.3.1.5 Database Maintenance and Updates A	5.4 Database Maintenance and Updates

Legend: Capability Level



Foundational





Advanced



Very Advanced

1.8 Domain Level Process Map



2. Data Storage and Operations Policy



2. Data Storage and Operations Policy

What is a Data Storage and Operations Policy and why does the Entity need it?

A Data Storage & Operations Policy is a comprehensive framework that defines the principles, rules, and guidelines for managing data storage across its lifecycle, including real-time, near-real-time, and remote data. It outlines processes for data retention, access control, secure disposal, backup and recovery, and compliance with legal and regulatory requirements. This policy ensures that the entity's data storage practices align with national standards and organizational objectives.

The entity requires a Data Storage & Operations Policy to establish compliance with legal, regulatory, and contractual obligations while promoting standardization and security across its data storage operations. By defining processes for retention, secure disposal, and disaster recovery, the policy mitigates risks associated with data breaches, unauthorized access, and accidental loss. Additionally, it provides a structured approach to optimizing resource usage, ensuring that storage infrastructure supports scalability and business continuity. Through clear roles, responsibilities, and operational guidelines, the policy aligns storage practices with the entity's broader strategic goals and national data management standards.

The remaining sections of this chapter will detail the activities of the below process that will guide you on executing Data Storage & Operations Policy:



Guiding Principles for Data Storage and Operations Policy

1. **Adopt a Risk-Based Approach:** Prioritize data storage practices based on their criticality and associated risks. Implement enhanced security measures for high-risk or sensitive data types.
2. **Comprehensive Coverage of Data Lifecycle:** Address all stages of the data lifecycle, including storage, retention, access, recovery, and disposal. Clearly define retention periods and secure disposal mechanisms based on data classification.
3. **Ensure Accessibility and Transparency:** Publish and socialize the policy across the entity to drive adoption. Use accessible digital channels to ensure stakeholders are informed and aware of their roles in storage operations.
4. **Implement Robust Access Controls:** Enforce role-based access control (RBAC) to limit access to storage systems based on job functions. Regularly review access permissions to prevent unauthorized access and ensure compliance with privacy regulations.
5. **Support Disaster Recovery and Continuity:** Include clear processes for backup and recovery, ensuring the entity can recover critical data in the event of failure, disaster, or accidental loss. Regularly test backup systems for effectiveness.

2.1 Data Storage & Retention Policy Creation

The Entity should establish a Data Storage & Retention Policy to ensure the secure and efficient storage of different types of data such as real-time data, near real-time data, and offline data.

Process - Data Storage & Retention Policy Creation




Figure 1: Data Storage & Retention Policy Creation Process Map

Legend: Process Step Decision Process Start Process End Document Off-Page Reference

2.1.A

Build Collaborative Working Group: To ensure the policy reflects operational realities, governance needs, and is not developed in isolation, the Entity must establish a collaborative working group of internal stakeholders and subject matter experts. This working group supports policy drafting by providing feedback, identifying operational considerations, and aligning the policy with Entity-specific contexts.

- i. **Establish the Collaborative Working Group:** Form a core team responsible for developing and reviewing the policy.
 - a. Led by the Data Governance Specialist, with oversight from the Data Management & Governance Section Head.
 - b. Include members from Data Governance, Data Architecture, Strategy, IT, compliance, and stewardship roles.
- ii. **Define Scope and Responsibilities of the Working Group:** Ensure the group's role is clear and time-bound.
 - a. Review proposed policy topics and structure.
 - b. Provide inputs during each phase of policy drafting and review.
 - c. Participate in validation and ensure alignment with Entity needs.
- iii. **Schedule Working Sessions:** Create a structured timeline for engagement.
 - a. Schedule policy drafting workshops and checkpoint reviews.
 - b. Ensure participation is documented and inputs are logged.

2.1.B

Assess Current Data Storage & Retention Policies: Before developing or updating the Data Storage & Retention Policy, the Entity must assess its current state. This includes reviewing any existing policies to identify gaps, overlaps, and misalignment with the Entity's data strategy, operational practices, and the National Data Standards. This assessment forms the foundation for effective and relevant policy development.

- i. **Conduct a Structured Policy Review:** Identify and evaluate any existing policies related to Data Storage & Retention.
 - a. If exists, review current documentation for:
 - Data Storage & Retention policy overview, purpose, scope, objectives, applicability, validity, policy statements, roles and responsibilities
 - Supporting Data Storage & Retention standards or procedures that may inform or support future policy development
 - b. Assess whether the policy content is outdated, incomplete, duplicated, or missing.
- ii. **Evaluate Alignment with Strategic and National Requirements:** Assess how well the current policy (if exists) supports key governance needs. Review for:
 - **Relevance to Strategic Objectives:** Does the policy align with the Entity's data strategy?
 - **Compliance with National Standards:** Does it reflect alignment with the current National Data Policy and National Data Standards?
 - **Clarity and Accessibility:** Is the policy easy to read, use, and apply?
- iii. **Document Findings and Key Gaps:** Summarize the outcomes of the assessment in preparation for step 2.1.C.
 - a. List key areas where new or updated policy content is needed.
 - b. Identify risks or inconsistencies that should be addressed (e.g., lack of version control, no update schedule).
 - c. Share findings with the collaborative working group for validation.



Key Message / Information

Building a working group ensures the policy is not developed in isolation but reflects the real governance needs of the Entity.

2.1.C
(1/2)

Determine Data Storage & Retention Policy Updates: Based on the findings from the policy assessment and input from the working group, the Entity must outline the specific updates or new content areas required in the Data Storage & Retention Policy. This step ensures that policy development is focused, complete, and responsive to both internal and national governance requirements.

2.1.C
Cont.
(2/2)

- i. **Define the Scope of Policy Creation/Update:** Translate assessment findings into a structured update plan.
 - a. Identify sections of the policy that need to be:
 - Created from scratch
 - Updated for alignment with strategy or standards
 - Removed due to redundancy or irrelevance
 - b. Determine whether the update (if any) qualifies as a major revision (e.g., structural change, expanded scope) or minor revision (e.g., wording adjustments).
- ii. **Prioritize Content Areas for Inclusion:** Ensure the policy covers all required Data Storage & Retention topics.
 - a. Reference Data Storage and Operations domain within the National Data Policy and National Data Standards to prioritize content including but not limited to:
 - Data storage and retention principles
 - Information system prioritization
 - Data backup and recovery
 - Storage conditions
 - b. Map these areas against identified gaps and stakeholder feedback.
- iii. **Create a Policy Update Action Plan:** Develop a high-level roadmap for the policy drafting process.
 - a. Outline key tasks, responsible roles, and delivery timelines.
 - b. Align this plan with the collaborative working group's availability and input schedule.
 - c. Use the action plan to guide coordination with stakeholders and support tracking of progress throughout the development cycle.

2.1.D
(1/3)

Develop/Update Data Storage & Retention Policy: Using the agreed-upon creation/update scope and structure, the Entity must develop or revise the Data Storage & Retention Policy to ensure the secure and efficient storage of different types of data such as real-time data, near real-time data, and offline data. The policy must be clear, actionable, aligned with strategic objectives, and compliant with the national data standards.

- i. **Draft Policy Content Using Approved Structure:** The Data Governance Specialist leads the drafting exercise using the standard Data Storage & Retention policy template.
 - a. Use the Data Storage & Retention Policy Template to structure the document.
 - b. The policy should include, at a minimum:
 - Overview, purpose, scope, and objectives
 - Policy applicability, validity, and review process
 - Guidelines aligned with the National Data Management Policy and Standards
 - Detailed policy statements covering:
 - Data storage and retention principles
 - Information systems prioritization
 - Database access rules

- Data backup and recovery
- Storage conditions based on approved regulations
- Retention periods of data
- Secure disposal and deletion of data
- Required actions in the even of an accidental permanent loss of data

- Roles and responsibilities
- References

Asset Guidance

The **Data Storage & Retention Policy Template** is a sample table of contents the Entity can refer to when drafting its policy. It includes the topics required by the standards, as well as additional recommended topics to consider. It also gives a sample of policy statements for one topic from the table of content to help entities draft the policy document.

Illustrative

Table of Contents	
1 Introduction.....#	2 Data Storage and Operations
1.1 Overview.....#	2.1 Data Storage & Retention Principles
1.2 Purpose.....#	1. All data shall be stored in secure, compliant, and scalable storage systems, ensuring accessibility, integrity, and availability across all operational environments.
1.3 Scope & Applicability.....#	2. Retention periods for data shall be defined based on its classification, business requirements, and regulatory obligations, ensuring data is retained only for as long as it is required.
1.4 Objectives.....#	3. Sensitive and critical data shall be encrypted both in transit and at rest to protect against unauthorized access and ensure compliance with relevant security standards.
2 Data Storage and Retention Policy.....#	4. Data storage systems shall be regularly audited and monitored to identify potential vulnerabilities, ensure capacity planning, and maintain optimal performance.
2.1 Data Storage and Retention Principles.....#	5. Archival processes shall be implemented for inactive or historical data, ensuring secure storage and efficient retrieval when necessary.
2.2 Information Systems Prioritization.....#	6. Secure disposal processes shall be enforced to permanently delete data that has surpassed its retention period, in compliance with legal and regulatory standards.
2.3 Database Access Rules.....#	7. Data replication and backup strategies shall be implemented to ensure business continuity and quick recovery from disruptions, with regular testing of recovery processes.
2.4 Data Backup and Recovery.....#	8. Data stored across distributed systems shall conform to organizational policies, ensuring consistency, synchronization, and centralized governance.
2.5 Storage Conditions based on Approved Regulations.....#	9. Retention policies shall be reviewed and updated periodically to align with changes in regulatory requirements, business needs, and technological advancements.
2.6 Data Retention Periods.....#	10. All stakeholders shall adhere to approved data storage and retention practices, with clear accountability assigned for managing, archiving, and disposing of data assets.
2.7 Secure Disposal and Deletion of Data.....#	
2.8 Accidental Permanent Loss of Data.....#	
3 Periodic Policy Review and Updates.....#	
3.1 Revision Frequency.....#	
3.2 Revisions.....#	
3.3 Archival Procedure.....#	
4 Roles within the Scope of this Policy.....#	
5 Related Legislation, References, and Documents.....#	
6 Document Control.....#	

*** (Sample policy statements for a topic) ***

ii. **Apply Clear and Actionable Language:** Ensure the policy is practical, enforceable, and accessible to all stakeholders.

- Use direct, concise language tailored to the Entity's context.
- Ensure each policy statement is tied to a strategic or operational outcome.
- Emphasize responsibilities, process expectations, and compliance triggers.

iii. **Prepare the Draft for Preliminary Review:** Finalize the draft as a working version for internal feedback.

- Share with the collaborative working group.
- Confirm the draft addresses all previously identified policy gaps.
- Prepare a feedback summary log for tracking suggestions in the next step.

2.1.D
Cont.
(3/3)**Key Message / Information**

Draft the policy using clear and concise language that reflects the Entity's specific needs while accommodating diverse operational scenarios. Each policy statement should be directly tied to broader business objectives, demonstrating how it supports strategic goals and organizational priorities. Emphasize actionable and enforceable language to enable consistent implementation and ensure compliance.

2.1.E

Initial Review for Draft Data Storage & Retention Policy: Once the policy draft is complete, the Entity must conduct an internal review to validate alignment with the National Data Standards, legal and regulatory obligations, and strategic objectives. This review ensures the policy is robust, compliant, and ready for broader publication steps.

- i. **Conduct Governance and Strategy Review:** The draft is reviewed by key stakeholders for policy integrity and relevance.
 - a. The Data Management & Governance Section Head and Data Strategy Expert assess the draft for:
 - Strategic alignment with the Entity's data governance objectives
 - Compliance with National Data Policy and Standards
 - Consistency with related internal policies
 - b. Confirm the policy structure follows the approved template and includes all required content areas.
- ii. **Validate Legal, Regulatory, and Operational Compliance:** Ensure the policy supports risk management and audit readiness.
 - a. Identify any areas requiring clarification or legal alignment.
 - b. Ensure the policy does not conflict with regulatory or operational obligations.
 - c. Highlight areas where governance enforcement or change management will be required.
- iii. **Capture and Document Feedback:** Record feedback from reviewers in a structured format to inform final refinements.
 - a. Document all comments and suggestions in a feedback log.
 - b. Prioritize revisions based on severity (e.g., compliance issues vs. language improvement).
 - c. Share the feedback log with the Data Governance Specialist for implementation in 2.1.F.

**Key Message / Information**

This review phase ensures the draft policy is credible, compliant, and strategically sound—before formalization or publication.

2.1.F

Incorporate Feedback: Following the internal review, the Entity must revise the draft Data Storage & Retention Policy to reflect feedback received from governance, strategy, legal, and operational stakeholders. This step ensures the final version is complete, accurate, and ready for formal approval and publication.

- i. **Review and Prioritize Feedback Log:** Work through the feedback log collected in 2.1.E in a structured and transparent way.
 - a. Categorize feedback by type (e.g., compliance, structure, clarity, alignment).
 - b. Prioritize revisions based on criticality, starting with compliance and strategy gaps.
 - c. Flag any unresolved or disputed comments for follow-up discussion with the Data Management & Governance Section Head.
- ii. **Apply Revisions to the Draft Policy:** Update the policy draft to reflect accepted changes.
 - a. Revise content in the official policy document, maintaining formatting consistency.
 - b. Track major changes using version control or change highlights.
 - c. Ensure all policy components (e.g., roles, responsibilities, review cycle) are fully addressed.
- iii. **Finalize the Policy Document for Approval:** Prepare the updated policy for formal submission.
 - a. Confirm that all feedback has been addressed or closed.
 - b. Update metadata such as version number, revision date, and approval status.
 - c. Share the final version with the Data Management & Governance Section Head for submission into the Entity's governance approval workflow.

RACI Matrix - Data Storage & Retention Policy Creation

Indicative

Sr No.	Activity	Data Management & Security Director	Data Management and Governance Section Head	Data Governance Specialist	Collaborative Working Group	Data Strategy Expert	Legal Team
2.1.A	Build Collaborative Working Group	A	R	C	I	I	
2.1.B	Assess Current Data Storage & Retention Policies		A	R	C	C	
2.1.C	Determine Data Storage & Retention Policy Updates		A	R	C	C	
2.1.D	Develop/Update Data Storage & Retention Policy		A	R	C	C	C
2.1.E	Initial Review for Draft Data Storage & Retention Policy	A	R	I	C	R	C
2.1.F	Incorporate Feedback		A	R	C		

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Build Collaborative Working Group of SMEs



Draft Data Storage & Retention Policy

2.2 Data Storage & Retention Policy Publication

This activity will guide the Entity on the steps to publish its data policies.

Process – Data Storage & Retention Policy Publication




Figure 2: Data Storage & Retention Policy Publication Process Map

Legend: Process Step Decision Process Start Process End Document Off-Page Reference

2.2.A
(1/2)

Review the Final Draft of the Data Storage & Retention Policy: Before proceeding to formal approval and publication, the Entity must conduct a final review of the completed Data Storage & Retention Policy. This step ensures that the policy aligns with the Entity’s data management strategy, governance principles, and business objectives, and that any remaining gaps are identified and addressed.

- i. **Confirm Strategic and Operational Alignment:** The Data Management & Security Director reviews the final draft to ensure policy alignment with enterprise priorities. Assess whether the policy:
 - Reflects the Entity’s data governance framework and strategic direction
 - Supports operational effectiveness and clear role accountability
 - Aligns with internal governance principles
- ii. **Identify Gaps, Misalignments, or Ambiguities:** Ensure the policy content is comprehensive and implementation-ready.
 - a. Flag missing elements, conflicting statements, or unclear policy language.
 - b. Identify any inconsistencies with related documents.
 - c. Provide written feedback on areas requiring revision or clarification.

2.2.A
Cont.
(2/2)

- iii. **Document Review Outcome and Next Steps:** Ensure feedback is structured and ready for incorporation in the next phase.
- a. Share feedback with the Data Governance Specialist in a structured format.
 - b. Record major comments for traceability (e.g., in a review tracker or comment log).
 - c. Confirm that feedback will be addressed in 2.2.B before publication proceeds.

**Key Message / Information**

This review step ensures the policy is accurate, aligned, and strategically sound before it enters the final approval and publication process.

2.2.B

Incorporate Director Feedback: Following the strategic review by the Data Management & Security Director, the Data Governance Specialist must revise the draft policy to reflect the provided feedback. This step ensures the policy addresses all strategic and governance-related concerns before moving into auditability review.

- i. **Review and Categorize Feedback:** Organize the feedback received from the director review for systematic implementation.
- Categorize comments (e.g., structural, content, compliance-related)
 - Prioritize critical feedback, especially those that impact alignment with strategic direction or compliance.
 - Clarify ambiguous feedback, if needed, through direct engagement with the reviewer.
- ii. **Apply Revisions to the Policy Draft:** Update the policy content to incorporate accepted feedback.
- a. Edit the draft to address feedback while preserving clarity and consistency.
 - b. Ensure changes strengthen strategic alignment, enforceability, and policy coherence.
 - c. Update document versioning and maintain a changelog of revisions made.
- iii. **Prepare the Policy for Auditability Review:** Ready the revised document for submission into the internal review and audit process.
- a. Confirm that all priority feedback has been addressed and documented.
 - b. Finalize the revised version for handover to the next stage.
 - c. Maintain a summary of revisions to support auditability and approval traceability.

2.2.C

Conduct an Auditability Review: Before final approval, the Entity must ensure the policy meets audit readiness criteria. This includes confirming that the policy is version-controlled, aligned with internal regulatory and operational standards and procedures, and compliant with documentation standards required for future audits and inspections. Feedback is gathered to address gaps in accountability, compliance, or performance tracking, if necessary.

2.2.D

Incorporate Internal Audit Feedback: Integrate feedback and recommendations from the internal audit review to address any remaining compliance, accountability, or documentation gaps. This step ensures the policy meets all auditability requirements and is fully prepared for final approval and formal publication.

2.2.E

Share Data Storage & Retention Policy with DGO: The Entity must submit the finalized draft of the Data Storage & Retention Policy to the DGO for formal review and feedback. This step ensures the policy aligns with national governance frameworks, regulatory expectations, and recognized data management best practices.

- i. **Submit the Finalized Policy to the DGO:** The policy is formally submitted through the appropriate channel for DGO review.
 - a. Send the finalized draft via the designated email or submission platform (as required at the respective point in time).
 - b. Confirm version number, submission date, and the responsible point of contact.
- ii. **Provide Supporting Context and Summary:** The submission must include a brief explanation to facilitate DGO review.
 - a. Summarize the policy scope, purpose, and any recent updates.
 - b. Highlight key areas where the policy aligns with national standards or fills previously identified governance gaps.
- iii. **Track the Submission and Prepare for Follow-Up:** Maintain traceability and readiness to act on feedback.
 - a. Log the submission in the Entity's governance repository or tracking system.
 - b. Note expected response timelines and plan for feedback incorporation in Step 2.2.F.

**Key Message / Information**

Engaging the DGO in the policy review process reinforces national alignment and promotes consistency across government-wide data storage and operations initiatives.

2.2.F
(1/2)

Prepare for Publication: Once all feedback has been incorporated, the Entity must finalize the policy for formal publication. This step ensures the document is properly versioned, archived, and ready to be communicated and distributed across the organization.

- i. **Finalize Policy Content and Format:** Ensure the policy is complete and ready for publishing.
 - a. Confirm all internal and DGO feedback has been fully addressed.
 - b. Apply final formatting, ensure clarity and consistency, and validate that the structure follows the approved policy template.
 - c. Update metadata fields including version number, status, effective date, and next review cycle.

2.2.F
Cont.
(2/2)

- ii. **Archive and Register the Final Version:** Secure the policy in Entity records and ensure traceability.
 - a. Save the finalized document in the Entity's governance repository or policy management system.
 - b. Maintain version history and a changelog, including approval and feedback documentation.
 - c. Confirm access permissions are in place ahead of publication and communication.
- iii. **Notify Stakeholders of Upcoming Publication:** Alert relevant teams that the policy is finalized and will be published in 2.2.G.
 - a. Share a summary message with data owners, stewards, and governance leads.
 - b. Prepare communication materials (e.g., email, announcement note) for rollout.



Key Message / Information

Preparing the policy for publication ensures it is complete, traceable, and properly maintained, enabling a smooth transition to Entity-wide communication and adoption.

2.2.G
(1/2)

Publish and Socialize the Data Storage & Retention Policy: Once finalized, the Entity must publish the policy in an accessible location and actively communicate its purpose, scope, and expectations to relevant stakeholders. This step ensures the policy is not only available, but also understood, adopted, and applied consistently across the organization.

- i. **Publish the Policy in the Official Repository:** Ensure centralized access and version control.
 - a. Upload the final policy to the Entity's internal policy repository or governance portal.
 - b. Tag the document with metadata such as owner, version, status, effective date, and next review date.
 - c. Confirm appropriate access permissions are configured.
- ii. **Communicate Policy Availability and Key Messages:** Promote awareness and usage of the policy across roles and departments.
 - a. Use the Policy Communication Email Template to inform staff of the policy release.
 - b. Include the policy's purpose, scope, and summary of key expectations.
 - c. Tailor communications by role if necessary (e.g., Stewards, Data Owners, IT, Legal).
- iii. **Reinforce Policy Adoption Through Awareness Channels:** Ensure the policy is understood and referenced in practice.
 - a. Present the policy in onboarding, training sessions, or governance briefings.
 - b. Link the policy in supporting materials such as guidelines, templates, or dashboards.
 - c. Gather informal feedback or usage indicators to monitor engagement.

Asset Guidance

The **Data Management Policy Communication Email Template** is a sample announcement email the Entity can use to notify staff about the publication of the policy and provide key information.

Illustrative

Subject: [IMPORTANT] New [Domain Name] Policy Published – Please Review

Dear [Recipient Name/Team],

We are pleased to announce the publication of the new [Domain Name] Policy, which is now in effect as of [Date]. This policy is designed to [briefly explain the purpose, e.g., "ensure the effective management of metadata and data catalog operations across the entity"] and aligns with the National Data Standards.

What is the Policy About?

The [Domain Name] Policy governs the management, oversight, and compliance of [specific domain] activities across the entity. It outlines critical guidelines and procedures, including:

- [Key Point 1, e.g., "Data Management and Organizational Structure"]
- [Key Point 2, e.g., "Version control and security measures"]
- [Key Point 3, e.g., "Metadata quality and governance"]

This policy has been developed in collaboration with the Data Committee and relevant stakeholders, ensuring compliance with the national data management standards.

How Does This Affect You?

- **Action Required:** Please review the policy and ensure that your team is aligned with its requirements. The policy affects [mention departments or roles impacted, e.g., "all data stewards, IT staff, and department heads"]. Compliance with this policy is essential to ensure adherence to the National Data Management Standards, which are critical to the entity's data governance and operational success.
- **Changes to Processes:** [If applicable, briefly describe any changes to current workflows or responsibilities as outlined by the policy.]

Where Can You Access the Policy?

The policy is available on our internal [intranet/portal] and can be accessed through the following link:
[Link to Policy Document]

Next Steps

- **Review Deadline:** Please ensure that you and your team have reviewed the policy by [Date].
- **Training/Workshops:** [If applicable] We will be hosting a [training/workshop] on [Date] to assist teams in understanding and applying the new policy.
- **Ongoing Support:** For questions, feedback, or support, please contact [Support Contact] at [email/phone number].

Thank you for your attention to this important matter and for your cooperation in ensuring compliance with the new policy. Together, we can maintain a strong framework for [mention domain, e.g., "data catalog and metadata management"] across the entity.

Best regards,
[Your Name]
[Your Position]
[Your Department]
[Contact Information]



Key Message / Information

Publishing the policy is not the end of the process. It must be socialized effectively so that stakeholders understand how it applies to their roles and responsibilities.

Process Risks: Develop and Publish Data Storage & Retention Policy

Risk	Severity	Mitigation
1. Lack of engagement from business SMEs due to other demanding responsibilities.	High	Data Management and Governance Section Head to secure management commitment for the key SMEs required to develop the Data Storage & Retention policy.
2. No standardized or existing process to publish the policy.	Medium	Data Management and Governance Section Head to coordinate with the compliance office to establish a centralized process and repository for the publication and maintenance of the Data Storage & Retention Policy.

RACI Matrix - Data Storage & Retention Policy Publication

Indicative

Sr No.	Activity	DGO	Data Management & Security Director	Data Management and Governance Section Head	Data Governance Specialist	Data Strategy Expert	Internal Audit
2.2.A	Review the Final Draft of the Data Storage & Retention Policy		A/R	C			
2.2.B	Incorporate Director Feedback		C	A	R		
2.2.C	Conduct an Auditability Review			I	I		A/R
2.2.D	Incorporate Internal Audit Feedback			A	R	C	
2.2.E	Share Data Storage & Retention Policy with DGO	I	A	R			
2.2.F	Prepare for Publication			A	R	C	
2.2.G	Publish and Socialize Data Storage & Retention Policy		A	R		C	

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Publish Data Storage & Retention Policy

3. Data Storage and Operations Planning



3. Data Storage and Operations Planning


What is Data Storage and Operations Planning and why does the Entity need it?

The Data Storage & Operations Planning process provides a structured pathway for assessing, forecasting, and designing the Entity's future data storage environment in line with national priorities. It bridges policy and practice by enabling the Entity to evaluate its current capabilities, anticipate future demands, and design a flexible, secure, and cloud-aligned architecture to support evolving business and technology needs.

This process ensures that the Entity has a clear understanding of its current infrastructure footprint, projected storage growth, and the architectural components required to support continuity, scalability, and compliance. It empowers both technical and business teams to collaborate on defining capacity, resilience, and modernization needs.

Through activities such as baseline capability assessment, capacity forecasting, architectural design, and roadmap development, the Entity can reduce risk, eliminate fragmentation, and ensure readiness for digital transformation. These planning efforts enable smarter investment decisions, improved data availability, and stronger alignment with national data standards and security policies.

The remaining sections of this chapter will detail the below activities that guide the Entity through executing the Data Storage & Operations Planning process.



Guiding Principles for Data Storage and Operations Planning

1. **Design for Scalability and Resilience:** Architect storage infrastructure that can grow with the Entity's needs while withstanding performance and availability pressures.
2. **Prioritize Business Continuity:** Ensure that backup, disaster recovery, and failover strategies are central to all planning decisions to protect mission-critical data.
3. **Optimize for Cloud Readiness:** Adopt a cloud-first posture by selecting platforms and tools that are compatible with national cloud policies and endorsed CSPs.
4. **Plan for Interoperability:** Design storage to integrate with data platforms, analytics systems, and national digital infrastructure to support unified data access and governance.

3.1 Baseline Data Storage Capability Assessment

This activity guides the Entity in retrieving or refreshing its maturity assessment and establishes a baseline view of the Entity's current data storage infrastructure across physical, virtual, and cloud environments.

3.1.1 Data Storage and Operations Maturity Assessment

To build an effective roadmap, the Entity must first assess its current state. This activity guides the Entity in retrieving or refreshing its maturity assessment for the Data Storage and Operations (DSO) domain using the Data Management Current Maturity Assessment Framework and Data Management Maturity Model. The assessment identifies current strengths, gaps, and improvement priorities across governance, processes, and tools, and serves as the foundation for roadmap planning.

Process – Data Storage and Operations Maturity Assessment




Figure 3: Data Storage and Operations Maturity Assessment Process Map

Legend: Process Step Decision Process Start Process End Document Off-Page Reference

3.1.1.A (1/2)

Retrieve the Latest Data Storage and Operations Maturity Assessment. The Entity must retrieve its most recent maturity assessment to establish a baseline understanding of current capabilities across the DSO domain. This step is essential for determining whether an update or reassessment is required and provides the foundation for roadmap planning.

- i. **Request Existing Assessment Report:** The Data Architect collaborates with the Data Strategy Expert to locate and retrieve the latest data storage and operations maturity assessment.
 - Engage the Data Governance Specialist to obtain the most recent maturity assessment report.
 - Confirm that the report includes maturity ratings across all DSO dimensions.
 - Validate that the report was formally conducted and documented using the approved Data Management Current Maturity Assessment Framework.

3.1.1.A
Cont.
(2/2)

- ii. **Validate the Currency of the Assessment:** The Data Architect and Data Strategy Expert jointly evaluate whether the existing assessment is still valid and representative of current capabilities.
- Confirm that the assessment was completed within the past 12 months.
 - Assess whether any significant organizational, tooling, or process changes have occurred since the report.
 - If the report is outdated or misaligned with the Entity's current state, outline that a refresh will be required.
 - Note any open or unaddressed recommendations that may still be relevant and require integration into the new roadmap.



Key Message / Information

Ensure that the maturity assessment covers both business and technical dimensions. If the assessment is fragmented or was conducted informally, treat it as insufficient and proceed to a formal refresh.

A current and validated maturity assessment is essential to building a grounded and targeted improvement roadmap. Without it, any planning activity risks being disconnected from the entity's operational reality and compliance posture.

3.1.1.B
(1/2)

Refresh the Data Storage and Operations Maturity Assessment (If Required): If the Entity's existing DSO maturity assessment is outdated, incomplete, or unavailable, a full reassessment must be conducted using the **Data Management Current Maturity Assessment Framework** and the supporting **Data Management Maturity Model** (referenced in Activity 2.2 within the Data Management Strategy and Governance domain guidelines). This ensures that the planning efforts in the roadmap are grounded in an accurate and current view of the Entity's storage and operations capabilities.

- i. **Initiate the Refresh Process:** The Data Architecture and Analytics Section Head triggers the reassessment and mobilizes key stakeholders.
- Formally launch the maturity refresh process by assigning ownership to the Metadata Specialist.
 - Confirm participation of relevant stakeholders including Data Stewards, Data Owners, Data Custodians, IT, and business users.
 - Communicate the scope and purpose of the reassessment to all contributors.
- ii. **Apply the Data Management Current Maturity Assessment Framework:** The Entity must use the **Data Management Maturity Model** to evaluate DSO current-state capabilities.
- Review and respond to each of the assessment dimensions (Organization, Operations, Storage Monitoring, etc.) using the 1–5 scale.
 - Ensure contributors refer to actual practices, tools, and governance documentation, and not perceived maturity.
 - Encourage team discussions and workshops to calibrate responses and promote honest evaluation.

3.1.1.B
Cont.
(2/2)

iii. **Consolidate Input and Finalize Scores.** Leveraging the **Data Management Current Maturity Assessment Framework**, the Data Architect consolidates all responses, supporting evidence, and justifications.

- Validate consistency across responses and resolve conflicting input where necessary.
- Assign a final maturity level (1–5) for each dimension, including rationale.
- Prepare or update the maturity heatmap for the DSO domain, and document relevant observations.



Key Message / Information

Encourage transparency and use collaborative working sessions to resolve subjectivity and misalignment between technical and business contributors.

3.1.1.C
(1/2)

Analyze the Maturity Assessment Results: Once the DSO maturity assessment has been completed or refreshed, the Entity must analyze the results to understand current strengths, weaknesses, and capability gaps. This analysis forms the foundation for identifying high-priority improvement areas and preparing a focused roadmap.

- i. **Review Maturity Scores Across the DSO Dimensions:** The Data Architect and Data Strategy Expert jointly evaluate the final maturity scores to identify trends and performance patterns.
 - Review assigned maturity levels for all DSO dimensions using the 1–5 scale.
 - Identify areas of high maturity that reflect institutionalized and standardized practices.
 - Identify low-performing areas (e.g., Archiving, Storage Monitoring, etc.) that could hinder effective governance.
- ii. **Conduct Thematic Gap Analysis:** Group maturity results into logical clusters (e.g., governance, tooling, processes) to derive broader insights.
 - Categorize assessment results by functional theme (including but not limited to):
 - Governance and ownership
 - Storage
 - Operations
 - Identify cross-cutting gaps that appear across multiple dimensions (e.g., lack of organization affecting storage monitoring).
 - Highlight dependencies. For example, weak role clarity (low score in “Organization”) may contribute to poor storage management (low score in “Operations”).
- iii. **Identify Quick Wins and Structural Gaps:** Segment findings based on the level of effort and impact to prioritize next steps.

3.1.1.C
Cont.
(2/2)

- **Quick wins:** improvements that can be implemented with minimal effort (e.g., assigning roles).
- **Structural gaps:** issues that require strategic investment or change management (e.g., tool implementation, entity-wide policy enforcement).
- Recommend areas that should be prioritized in the roadmap due to compliance implications or high strategic value (e.g., defining a data archiving strategy and creating an operations SLA framework).



Key Message / Information

This analysis step is critical in converting maturity data into actionable insight. The quality of the roadmap depends on how well the Entity understands the underlying causes and impacts of its current maturity gaps.

3.1.1.D
(1/2)

Document Assessment Insights and Recommendations: The Entity must consolidate the results of the DSO maturity assessment into a structured executive-level slide. This summary is intended to present a high-level overview of the current maturity state, key findings, and prioritized opportunities for improvement. It will serve as a key input for roadmap planning and internal/external alignment.

- i. **Populate the Data Maturity Summary Slide:** The Data Architect prepares a summary slide (as referenced in the Data Management Current Maturity Assessment Framework in Activity 2.2 within the Data Management Strategy and Governance domain guidelines) that captures the summary maturity of the DSO domain.
 - a. **Maturity Score (Gauge):**
 - Calculate the overall average maturity score across all DSO dimensions.
 - Use the 1–5 scale and visually represent the result in the dial/gauge format.
 - Select the appropriate maturity level label (Initial, Limited, Evolving, etc.).
 - b. **Organization, Storage & Operations Processes, and Tools :**
 - Assign indicative scores across Organization, Storage & Operations Processes, and Tools based on maturity dimension grouping:
 - Organization: e.g., Organization, Operations
 - Storage & Operations Processes: e.g., Archiving, Storage Monitoring
 - Tools: e.g., Computing
 - c. **Key Insights & Findings:**
 - Summarize major pain points or maturity gaps discovered.
 - Highlight inconsistencies in governance, lack of tooling, low coverage, or siloed practices.
 - Use brief, high-impact bullet points appropriate for leadership-level review.
- ii. **Identify Key Improvement Opportunities:** This section translates findings into action-oriented strategic recommendations.

3.1.1.D
Cont.
(2/2)

- Define the (e.g., 2-4) highest-value opportunities for improvement.
- Ensure these align with major maturity gaps (e.g., introducing ownership, adopting a tool, formalizing change control).
- Frame opportunities in language suited for executive endorsement and roadmap alignment.

iii. **Validate with Key Stakeholders:** Before finalization, the summary slide must be reviewed by the Data Architecture and Analytics Section Head and Data Strategy Expert.

- Circulate the draft slide for review and feedback.
- Confirm that insights and opportunities are clearly and accurately captured.
- Refine language and formatting to ensure clarity, precision, and alignment with the rest of the Data Management Current Maturity Assessment Framework reporting pack.



Key Message / Information

Use concise, action-oriented phrasing; each bullet point should communicate a real insight or decision driver. Avoid vague recommendations and tie each suggestion to a specific weakness and pathway to resolution.

3.1.1.E

Finalize and Share the Maturity Assessment with DGO: If the DSO maturity assessment was refreshed, then the Entity must consolidate the findings and formally share the updated assessment (as part of the broader data management maturity assessment) with the DGO for alignment.

i. **Consolidate DSO Final Materials for Submission:** The Data Strategy Expert prepares the formal assessment package for DGO review.

- DSO domain level maturity scores
- DSO maturity heatmap
- DSO assessment insights and recommendations

ii. **Submit Assessment Package to the DGO:** The Data Architecture and Analytics Section Head oversees submission and initiates engagement with the DGO.

- Send the materials to the official DGO point-of-contact through the designated submission channel.
- Request a formal review and optionally ask for advisory feedback or validation of scoring.

iii. **Incorporate DGO Feedback and Finalize Record:** Once the DGO has reviewed the submission, incorporate any clarifications or adjustments required.

- Address DGO comments or requested edits in the assessment summary slide or supporting documents.
- Finalize all materials and archive the submission internally.
- Ensure the validated version is reflected in the Entity's central data governance records.

RACI Matrix – Data Storage and Operations Maturity Assessment

Indicative

Sr No.	Activity	DGO	Data Management and Security Director	Data Architecture and Analytics Section Head	Data Strategy Expert	Data Governance Specialist	Data Architect	Data Owner
3.1.1.A	Retrieve the Latest Data Storage and Operations Maturity Assessment		A	I	R	C	C	I
3.1.1.B	Refresh the Data Storage and Operations Maturity Assessment (If Required)		A	C	R	C	R	C
3.1.1.C	Analyze the Maturity Assessment Results		A	C	R	C	R	C
3.1.1.D	Document Assessment Insights and Recommendations		A	C	C		R	
3.1.1.E	Finalize and Share the Maturity Assessment with DGO	I/C	A	R	R	C	C	

RACI Matrix Legend: **R**: Responsible **A**: Accountable **C**: Consulted **I**: Informed

Based on this Activity, the Entity can



Refresh the Data Management Maturity Assessment



Analyze the Maturity Assessment Results

3.1.2 Technical Baseline Data Storage Assessment

This activity establishes a comprehensive baseline view of the Entity's current data storage infrastructure across physical, virtual, and cloud environments. It involves inventorying all storage systems, measuring utilization and performance, and classifying each system by its portability. The goal is to create a validated snapshot of existing capabilities to inform future planning, identify infrastructure risks, and support cloud-readiness decisions in alignment with the National Data Standards and Qatar's Cloud First Policy.

Process – Technical Baseline Data Storage Assessment




Figure 4: Technical Baseline Data Storage Assessment Process Map

Legend: ■ Process Step ◆ Decision ● Process Start ● Process End □ Document ■ Off-Page Reference

3.1.2.A
(1/2)

Initiate the Data Storage Capability Assessment: Establish a structured and coordinated launch of the assessment by defining scope, aligning stakeholders, and preparing tools. This foundational step ensures that the Entity begins with clarity, ownership, and consistency across all storage systems. It is led by the Data Architecture and Analytics Section Head, with support from the Data Architect, Data Custodian, and Data Steward.

- i. **Define assessment objectives and scope:** Establish a clear, shared understanding of what will be assessed and why it matters.
 - a. **Determine assessment purpose and strategic alignment:**
 - Ensure the baseline supports storage optimization, modernization, and compliance with Cloud First Policy
 - Link the output to follow-on activities in forecasting (3.2), target architecture (3.3), and roadmap development (3.4)
 - b. **Define coverage boundaries across infrastructure and systems:**
 - Include all storage systems—physical, virtual, hybrid, and cloud—in scope
 - Prioritize high-risk, high-cost, and mission-critical systems for detailed profiling
 - c. **Establish timelines and milestones:**
 - Set clear start and end dates, with checkpoints for validation and stakeholder review
 - Ensure alignment with broader data strategy or infrastructure initiatives
- ii. **Identify stakeholders and assign responsibilities:** Form a cross-functional assessment team.
 - a. **Designate the initiative lead:**
 - The Data Management & Governance Section Head will oversee coordination and progress tracking
 - b. **Assign role-specific responsibilities**
 - **Data Custodians:** Document physical and technical storage details (e.g., capacity, utilization, location, system type)
 - **Data Stewards:** Validate classification fields (e.g., portability, improvement areas), ensure consistency with metadata and governance rules
 - **Data Architects:** Provide oversight on alignment with current architecture and future transformation plans
 - **IT Infrastructure Leads:** Support data collection through CMDBs, monitoring tools, or export reports.

3.1.2.A
Cont.
(2/2)

c. Engage oversight bodies

- Notify the Data Committee and DGO for visibility and governance checkpoints during the assessment lifecycle

iii. **Establish assessment frameworks, tools, and templates:** Set a standard and structured approach using official documentation tools.

a. Leverage the “Baseline Storage Assessment Report” template issued by NPC:

Asset Guidance

The **Baseline Storage Assessment Report** can be leveraged by the Entity to document the current state of its data storage infrastructure. *Illustrative*

Baseline Storage Assessment Report						
The Baseline Storage Assessment Report Template is designed to document the current state of an organization's data storage infrastructure.						
Field	Description				Example	
Storage System Name	The name or identifier of the storage system.				Storage_Server_01	
Location	The physical or virtual location where the storage system resides.				Data Center A / Cloud - AWS	
System Type (Physical/Virtual)	Specifies whether the storage system is physical hardware or a virtualized storage environment.				Physical / Virtual	
Total Capacity (TB)	The total storage capacity of the system in terabytes (TB).				50 TB	
Used Capacity (TB)	The amount of storage currently being used in TB.				35 TB	
Available Capacity (TB)	The remaining unused storage capacity in TB.				15 TB	
Utilization Percentage	The percentage of total storage that is currently being used.				0.7	
Peak Usage (%)	The highest recorded storage utilization percentage over a given period.				0.85	
Baseline Usage (%)	The average storage utilization under normal operational conditions.				0.65	
Portability Classification	Classification of the system based on its ability to be migrated or modernized.				Legacy / Virtualize-able / Cloud-Eligible	
Key Risks	Identified risks associated with the storage system, such as high utilization, outdated hardware, or security vulnerabilities.				High storage consumption may lead to performance degradation	
Improvement Opportunities	Potential actions or upgrades to enhance storage efficiency, scalability, or security.				Upgrade to SSD for faster access; Migrate to cloud storage.	
Notes	Additional observations, comments, or contextual information.				Scheduled for expansion in Q3 2025.	

Storage System Name	Location	System Type (Physical/Virtual)	Total Capacity (TB)	Used Capacity (TB)	Available Capacity (TB)	Utilization Percentage

Peak Usage (%)	Baseline Usage (%)	Portability Classification	Key Risks	Improvement Opportunities	Notes

b. Define supporting criteria and thresholds for consistent assessment:

- Establish indicators for risk (e.g., >85% peak usage, legacy platforms, outdated hardware)
- Standardize use of terms for portability and improvement actions

c. Set up a secure workspace for collaboration:

- Use shared folders or repositories (e.g., SharePoint, ServiceNow, or secure drives) to consolidate input, track updates, and store backups
- Ensure access rights are provisioned appropriately for Custodians, Stewards, and Analysts

3.1.2.B
(1/2)

Conduct Data Storage Inventory: Collect a comprehensive, structured inventory of all data storage systems—both physical and virtual—across the Entity. This step focuses on identifying what exists, where it resides, how it is configured, and its current usage. The Data Custodian is the primary executor of this activity, with quality assurance support from the Data Steward, and technical validation by the IT Team.

- i. **Identify and log all storage systems across the Entity:** Create a unified list of all storage assets using the Baseline Storage Assessment Report template.

- a. **Conduct discovery workshops or CMDB extractions:**

- Pull system information from configuration management databases (CMDB), storage monitoring tools, or direct admin access
- Cross-reference against documented IT systems, business applications, and infrastructure maps

- b. **Document each storage system in the Baseline Storage Assessment Report:**

- Fill in system-level fields:
 - Storage System Name
 - Location (on-prem, data center, cloud)
 - System Type (physical or virtual)

- c. **Validate coverage with application owners and infrastructure teams:**

- Ensure the list is complete by verifying it with business units and technical teams
- Include special-use environments (e.g., test beds, disaster recovery, backup servers)

- ii. **Capture capacity, usage, and performance metrics:** Populate the template with accurate technical data for each system.

- a. **Record current and potential storage metrics:**

- Use monitoring tools or reports to document:
 - Total Capacity (TB)
 - Used Capacity (TB)
 - Available Capacity (TB)

- b. **Calculate utilization percentages:**

- Include:
 - Utilization Percentage (used ÷ total)
 - Peak Usage (%) over a defined window (e.g., last 6 months)
 - Baseline Usage (%) under normal operating conditions

3.1.2.B
Cont.
(2/2)

c. Tag systems showing abnormal or high-risk trends:

- Flag any systems with:
 - 85% peak utilization
 - Rapidly growing consumption rates
 - Imbalance between storage capacity and usage efficiency

iii. **Assess functional and business context of each system:** Connect technical storage data with business relevance.

a. Map storage systems to business functions or applications:

- Indicate whether the storage supports mission-critical services, reporting, archiving, etc
- Document interdependencies with upstream/downstream systems

b. Tag each system with a Portability Classification:

- Select one of the following:
 - Legacy — Cannot be easily virtualized or migrated
 - Virtualize-able — Eligible for virtualization but not cloud-native
 - Cloud-Eligible — Compatible with cloud storage models per Cloud First Policy

c. Coordinate with Data Stewards to validate business criticality and data relevance:

- Ensure classification is accurate from a governance and usage perspective
- Highlight systems with restricted or sensitive data that may impact migration decisions

3.1.2.C
(1/2)

Assess Infrastructure Utilization and Performance: Analyze the technical performance and utilization of all storage infrastructure components across operational environments. This step provides visibility into actual workloads, helps identify underutilized or overburdened resources, and forms the basis for optimization or migration planning. The Data Custodian leads the data collection and monitoring setup, with validation support from the IT Team, and the Data Architect for infrastructure mapping and impact analysis.

i. **Collect system-level performance and load metrics:** Gather peak and baseline performance data for all storage environments.

a. Measure key infrastructure indicators across storage systems:

- CPU usage
- Memory load
- Disk I/O throughput and latency
- Network I/O bandwidth and traffic patterns

3.1.2.C
Cont.
(2/2)**b. Record both baseline and peak performance levels:**

- Use monitoring dashboards or server logs to track operating norms
- Identify seasonal or cyclic spikes in resource demand (e.g., month-end reports, year-end closures)

c. Use consistent monitoring windows for comparability:

- Recommend 30–90 day monitoring windows to capture representative usage
- Align timeframes across systems to normalize analysis

ii. Analyze utilization efficiency and system load trends: Identify systems operating beyond or below expected capacity thresholds.**a. Flag systems showing underutilization or capacity strain:**

- Highlight those with <30% baseline usage or >85% peak utilization
- Identify ghost systems or abandoned storage that consumes power or backup cycles

b. Detect anomalies or negative performance trends:

- Compare current metrics against historical patterns using tools like Grafana, Nagios, or SolarWinds
- Look for growing latency, increasing load times, or deteriorating availability

c. Document systems requiring attention in the assessment report:

- Use the “Key Risks” field to capture storage saturation, outdated technology, or thermal issues
- Include performance constraints that affect application responsiveness or business continuity

iii. Validate and benchmark system performance contextually: Ensure infrastructure analysis is aligned with real business requirements.**a. Correlate performance data with business-critical systems and workloads:**

- Prioritize systems that support 24/7 services, real-time data access, or financial processing
- Document the impact of performance degradation on operations

b. Consult with business and application owners as needed:

- Validate load profiles for peak events (e.g., billing cycles, citizen service surges)
- Verify perceived performance bottlenecks versus actual technical issues

c. Identify infrastructure modernization or reallocation opportunities:

- Recommend actions such as rebalancing workloads, retiring outdated systems, or upgrading hardware
- Feed insights into the next step (3.1.D) to assess suitability for cloud or virtualization transition

3.1.2.D
(1/2)

Classify Information Systems by Portability: Evaluate each storage system and associated application or platform for its ability to be migrated, modernized, or virtualized. This classification enables better planning for cloud enablement, infrastructure upgrades, and future storage architecture decisions. The Data Custodian leads this classification process using technical criteria.

- i. **Define portability criteria and decision logic:** Establish a uniform method to classify systems across the entity.
 - a. **Use the template classification scheme (Legacy / Virtualize-able / Cloud-Eligible):**
 - **Legacy:** Systems not compatible with virtualization or cloud due to obsolete software, monolithic architecture, or vendor lock-in
 - **Virtualize-able:** Systems that can be hosted in virtual environments but are not inherently cloud-native
 - **Cloud-Eligible:** Systems that can be migrated to cloud platforms with minimal modification or are already compatible
 - b. **Align classification logic with Cloud First Policy requirements:**
 - Confirm compatibility with endorsed CSP platforms (e.g., AWS, Azure, GCP)
 - Consider support for containerization, API integration, and multi-tenancy features
 - c. **Define technology and operational triggers for classification:**
 - Include factors such as OS compatibility, system modularity, hardware dependency, middleware stack, and licensing conditions
- ii. **Evaluate each system against classification logic:** Apply the decision criteria systematically across the full storage inventory.
 - a. **Review supporting infrastructure and application stack per system:**
 - Include data from system architecture diagrams, deployment manifests, and dependency maps
 - Use application inventory and CMDB sources for validation
 - b. **Conduct joint reviews with IT, and application owners:**
 - Validate real-world operability, hosting limitations, and interoperability risks
 - Identify recent attempts or blockers for migration, if applicable
 - c. **Update the “Portability Classification” field in the Baseline Storage Assessment Report:**
 - Build on the previously captured inventory by adding the final portability classification per system

3.1.2.D
Cont.
(2/2)

- Tag systems that require further migration assessment or pose technical limitations using the “Notes” or “Improvement Opportunities” fields

iii. **Use classification outcomes to inform strategic planning:** Translate technical classification into actionable guidance for the roadmap.

a. **Highlight systems eligible for cloud transition in short to medium term:**

- Prioritize Cloud-Eligible and Virtualize-able systems in future activities (e.g., 3.3 and 3.4)
- Estimate potential cost savings and operational benefits

b. **Flag legacy systems requiring containment, replacement, or retirement:**

- Include these in the “Key Risks” field with a justification of limitations
- Consider sunset strategies, encapsulation with middleware, or offloading non-critical workloads

c. **Prepare classification summary charts for the Data Committee:**

- Include counts or percentages by classification type
- Visualize system distribution across portability tiers to support strategic alignment and decision-making

3.1.2.E
(1/2)

Document and Report Baseline Findings: Finalize and consolidate the assessment results into a structured, reviewed, and approved baseline report. This step ensures the inventory is complete, validated, and ready to inform forecasting, target architecture, and roadmap planning. The Data Governance Specialist leads the reporting process with contributions from the Data Custodian, Data Steward, and Data Architect.

i. **Consolidate and validate inputs from all sub-steps:** Ensure the Baseline Storage Assessment Report is complete, consistent, and technically sound.

a. **Review data completeness across all required fields in the template:**

- Confirm that all systems have entries for storage location, capacity, usage, and portability classification
- Ensure “Key Risks” and “Improvement Opportunities” are populated for systems flagged in 3.1.C or 3.1.D

b. **Cross-check for consistency and duplicates:**

- Validate naming conventions, location references, and system IDs with infrastructure teams
- Resolve duplication or inconsistency across entries related to multi-environment systems

c. Validate business and technical alignment:

- Data Stewards verify classification relevance and metadata accuracy
- Data Custodians verify infrastructure-level metrics and tagging accuracy
- Data Architects review for completeness in coverage and consistency with architecture views

ii. Summarize key insights and risk findings: Translate technical records into insights for planning and governance.

a. Develop summary views for stakeholder briefings:

- Create visuals showing total capacity distribution, storage usage by type, and portability classification breakdown
- Include count of systems by classification (Legacy / Virtualize-able / Cloud-Eligible)

b. Highlight storage-related risks and pain points:

- Summarize systems with critical usage saturation, legacy lock-in, or known availability issues
- Use visuals or tables to show where the most urgent infrastructure interventions are needed

c. Identify strategic opportunities and quick wins:

- Flag systems suitable for reallocation, retirement, or immediate cloud migration
- List opportunities for consolidation, optimization, or sunset planning

iii. Finalize, present, and store the report: Secure endorsement and ensure the report is officially archived and accessible.

a. Submit the final report to the Data Committee for endorsement:

- Prepare an executive summary and visual dashboard to support decision-making
- Include recommendations for deeper analysis or migration scoping where needed

b. Notify stakeholders and initiate handover to downstream activities:

- Share the report and findings with teams leading Activities 3.2, 3.3, and 3.4
- Conduct a handover session to ensure continuity of assumptions and insights

RACI Matrix - Technical Baseline Data Storage Assessment

Indicative

Sr No.	Activity	Data Committee	Data Management & Security Director	Data Architecture and Analytics Section Head	Data Governance Specialist	Data Architect	Data Custodian	Data Steward	IT Team
3.1.2.A	Initiate the Data Storage Capability Assessment		A	R	C	C	C	C	
3.1.2.B	Conduct Data Storage Inventory			A	C		R	C	C
3.1.2.C	Assess Infrastructure Utilization and Performance			A	C	C	R		C
3.1.2.D	Classify Information Systems by Portability			A	C	C	R		C
3.1.2.E	Document and Report Baseline Findings	C	I	A	R	C	C	C	

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Conduct Data Storage Inventory



Document and Report Baseline Findings

3.2 Data Storage and Usage Requirements Forecasting

This activity establishes a comprehensive baseline view of the Entity’s current data storage infrastructure across physical, virtual, and cloud environments. It involves inventorying all storage systems, measuring utilization and performance, and classifying each system by its portability. The goal is to create a validated snapshot of existing capabilities to inform future planning, identify infrastructure risks, and support cloud-readiness decisions in alignment with the National Data Standards and Qatar’s Cloud First Policy.

Process – Data Storage and Usage Requirements Forecasting




Figure 5: Data Storage and Usage Requirements Forecasting Process Map

Legend: Process Step Decision Process Start Process End Document Off-Page Reference

3.2.A
(1/2)

Establish Forecasting Scope and Parameters: Define a clear scope, timeframe, and forecasting framework to ensure accurate and relevant storage demand projections. This foundational step ensures that the forecasting activity aligns with business priorities, system coverage, and the needs of downstream planning. The Data Architecture and Analytics Section Head leads this step with support from the Data Architect, Data Custodian, Business Data Analysts, and relevant System Owners.

- i. **Define systems and data types to include in the forecast:** Establish a shared understanding of what will be projected and why.
 - a. **Identify systems in scope for storage forecasting:**
 - Include systems with high growth potential, critical business impact, or pending modernization
 - Focus on systems already classified in 3.1.D as Virtualize-able or Cloud-Eligible
 - b. **Specify types of data to be included in projections:**
 - Structured (e.g., databases, transaction logs), semi-structured (e.g., XML, JSON), and unstructured (e.g., video, PDFs)
 - Consider data domains such as operational, analytical, archival, and system-generated
 - c. **Exclude low-relevance systems based on business input:**
 - Filter out systems marked for retirement or deemed stagnant with no expected growth
 - Confirm exclusions with business owners to avoid blind spots
- ii. **Define forecasting time horizon and granularity:** Determine the period and resolution of projections.
 - a. **Set standard forecasting periods in alignment with the National Data Standards:**
 - First-year projection (initial capacity estimate)
 - Multi-year outlook (typically 3–5 years forward)
 - b. **Choose forecasting granularity per use case:**
 - Monthly or quarterly granularity for mission-critical systems
 - Annual granularity for low-growth or archival systems
 - c. **Align forecast windows with budgeting and strategy cycles:**
 - Sync timelines with financial planning, procurement schedules, and infrastructure upgrade cycles
 - Confirm dependencies with the Entity's IT and Finance departments

3.2.A
Cont.
(2/2)

iii. **Select methods, tools, and reference data for forecasting:** Equip the team with the models and baselines to use.

a. **Use historical usage data from 3.1 as the primary baseline:**

- Leverage peak, average, and trend data to model realistic growth curves
- Account for known data inflow events (e.g., new service launches, regulatory mandates)

b. **Choose forecasting methods appropriate to system type:**

- Linear trend models for steady-growth systems
- Step-change or scenario-based models for digital transformation programs or AI/IoT integrations

c. **Leverage “Storage Forecasting Report” template to document projections:**

- Align all stakeholders on using the standardized template
- Review the fields in the template (e.g., growth rates, year-by-year estimates, peak usage considerations) to ensure shared understanding of expectations
- Communicate when and how the template will be populated in the upcoming forecasting steps (3.2.B onward)

Asset Guidance

The **Storage and Usage Requirements Forecasting Template** can be leveraged by the Entity to estimate future storage capacity needs.

Illustrative

Storage Forecasting Report						
The Storage Forecasting Report Template is designed to provide a structured approach for estimating future storage capacity needs.						
Field	Description			Example		
Storage System Name	The name or identifier of the storage system being analyzed.			Storage_Server_01		
Current Capacity (TB)	The total storage capacity available for the system.			100 TB		
Current Usage (TB)	The amount of storage currently in use.			70 TB		
Projected Growth Rate (%)	The expected annual growth rate of storage usage.			15%		
Estimated Capacity Requirement (Year 1) (TB)	The projected storage requirement for the first year based on growth trends.			80 TB		
Estimated Capacity Requirement (Year 3) (TB)	The projected storage requirement for three years ahead.			106 TB		
Estimated Capacity Requirement (Year 5) (TB)	The projected storage requirement for five years ahead.			140 TB		
Peak Usage Considerations	Any expected spikes in storage usage due to seasonal demands or business operations.			Q4 sales data surge		
Budget Estimate (CapEx & OpEx)	Estimated budget needed for storage expansion, considering hardware and cloud costs.			\$50,000 (CapEx), \$10,000/year (OpEx)		
Forecasting Assumptions	Key assumptions used in the projections, such as expected data growth trends and business expansion.			Data growth will remain constant at 15% per year.		
Constraints & Risks	Potential risks or limitations that could impact storage growth.			Budget restrictions may delay upgrades.		
Mitigation Strategies	Possible solutions to address risks, such as optimizing storage usage or adopting hybrid storage solutions.			Implement data archiving to reduce primary storage load.		
Notes	Additional observations or remarks regarding the storage forecasting analysis.			Migration to cloud storage planned for Q2.		

Storage System Name	Current Capacity (TB)	Current Usage (TB)	Projected Growth Rate (%)	Estimated Capacity Requirement (Year 1) (TB)	Estimated Capacity Requirement (Year 3) (TB)	Estimated Capacity Requirement (Year 5) (TB)

Peak Usage Considerations	Budget Estimate (CapEx & OpEx)	Forecasting Assumptions	Constraints & Risks	Mitigation Strategies	Notes

3.2.B

Estimate Initial Storage Capacity Requirements: Develop Year 1 storage capacity projections for each in-scope system based on current utilization and expected short-term growth. This step forms the baseline of the multi-year forecast and ensures readiness for near-term investments or cloud transitions. The Data Custodian leads this activity with support from the Data Architect, Business Data Analyst, and system/application owners, and in coordination with the Finance and IT teams.

- i. **Establish a current-state baseline for each system:** Document the present capacity and usage figures as the starting point for projection.
 - a. **Use current capacity and utilization data from the 3.1 Baseline Assessment Report:**
 - Confirm fields such as Current Capacity (TB) and Current Usage (TB)
 - Validate against monitoring dashboards or infrastructure reports
 - b. **Ensure consistency with system boundaries and data types:**
 - Verify that forecasted units (e.g., TB) include space needed for both operational and metadata growth
 - Include all environments where applicable (e.g., production, dev/test, DR)
 - c. **Document these values in the “Storage Forecasting Report” template under the corresponding fields:**
 - Use Storage System Name, Current Capacity (TB), and Current Usage (TB) as the initial reference line for each row
 - Tag each system with a unique identifier or code for traceability
- ii. **Estimate Year 1 growth rates and projected usage:** Forecast the incremental increase in data volume for the upcoming year.
 - a. **Work with system and business owners to identify Year 1 growth drivers:**
 - Consider service expansion, digital onboarding of new users, increased reporting, new sensor/data streams, or policy mandates
 - Use historical growth trends where available (e.g., monthly storage increase rates)
 - b. **Determine the projected growth rate (%) for each system:**
 - Use linear extrapolation, percentage-based assumptions, or business-provided estimates
 - Document growth assumptions clearly in the “Forecasting Assumptions” field
 - c. **Populate the Projected Growth Rate (%) and Estimated Capacity Requirement (Year 1) fields:**
 - Calculate Estimated Requirement = Current Usage × (1 + Growth Rate)
 - Include systems where growth may trigger crossing thresholds or procurement triggers (e.g., >80% capacity)

3.2.C
(1/2)

Project Multi-Year Growth Trends: Extend storage forecasting to cover a 3-to-5-year horizon, capturing both gradual data expansion and anticipated step changes. This forward-looking view supports long-term budgeting, procurement, and capacity planning decisions. The Data Architect and Business Data Analyst co-lead this activity, supported by Data Custodians, IT, and Finance teams.

- i. **Define baseline and scenario assumptions for multi-year projections:** Ensure future capacity estimates are grounded in documented business and technical inputs.
 - a. **Review historical growth rates and business trends:**
 - Analyze past growth patterns for each system to establish a trend baseline
 - Reference metrics from 3.1 and Year 1 forecasts completed in 3.2.B
 - b. **Identify forward-looking business, regulatory, or technology drivers:**
 - Examples include: national initiatives, digital transformation programs, new service platforms, increasing sensor/IoT feeds, or legal mandates to retain data
 - Incorporate known cloud migration plans that may impact on-prem storage demand
 - c. **Document forecasting assumptions in the Storage Forecasting Report:**
 - Use the “Forecasting Assumptions” field to clarify rationale behind Year 3 and Year 5 estimates
 - Tag projections as “baseline”, “accelerated”, or “conservative” where scenario modeling is used
- ii. **Estimate capacity requirements for Year 3 and Year 5:** Generate forward-looking storage projections for each in-scope system.
 - a. **Choose a consistent forecasting method for extrapolation:**
 - Use compound annual growth rate (CAGR) or tiered step functions depending on the system’s nature
 - For dynamic systems, factor in business drivers that create nonlinear growth
 - b. **Populate Estimated Capacity Requirement (Year 3) and (Year 5) fields in the template:**
 - Derive these from Current Usage and Projected Growth Rate (%) or other business-defined models
 - Ensure units and methodology are consistent across systems
 - c. **Include Peak Usage Considerations for known usage surges:**
 - Example: annual census data loading, Q4 transactional loads, or disaster simulation exercises
 - Use this field to capture events that may temporarily exceed average projected growth

3.2.C
Cont.
(2/2)

iii. **Assess impact on infrastructure and cloud readiness:** Translate multi-year projections into strategic signals for capacity and architectural planning.

a. **Identify systems likely to exceed infrastructure thresholds or outgrow current configurations:**

- Flag any systems projected to reach 90%+ capacity within 3 years
- Consider whether scaling on-prem is viable or if cloud or hybrid strategies should be initiated

b. **Coordinate with architecture and procurement teams for alignment:**

- Use projections to inform design of target architecture in Activity 3.3
- Share expected future needs with procurement or cloud credit planning functions

c. **Record constraints or challenges in the “Constraints & Risks” field:**

- Include limitations related to budget, data residency, system interoperability, or vendor contracts
- Recommend any required mitigation actions or planning interventions

3.2.D
(1/2)

Identify Redundancy, Backup, and Compliance Storage Needs: Forecast storage capacity required for backup, disaster recovery, redundancy, and regulatory compliance, in addition to primary data growth. These estimates ensure holistic infrastructure planning and alignment with business continuity, data protection, and legal mandates. The Data Custodian leads this step, with input from Information Security, Legal, and the Data Architect.

i. **Define categories of non-primary storage to be forecasted:** Break down additional capacity needs into distinct use cases.

a. **Identify backup storage requirements for operational recovery:**

- Include full backups, incremental snapshots, transaction logs, and database recovery images
- Consider different retention periods for daily, weekly, and monthly backup tiers

b. **Estimate redundancy and replication storage requirements:**

- Capture needs for RAID, SAN mirrors, high availability (HA) clusters, and active-active failover
- Include storage consumed in geographically redundant DR environments

c. **Include compliance and archival storage obligations:**

- Consider legal or regulatory mandates to retain datasets for fixed periods (e.g., 5 or 10 years)
- Factor in data immutability or WORM (write-once-read-many) storage where applicable

3.2.D
Cont.
(2/2)

- ii. **Estimate volume impact of each storage category:** Translate each category into a concrete storage forecast per system or environment.
 - a. **Determine storage multipliers based on policy and architecture:**
 - Example: if daily full backups are retained for 7 days and weekly snapshots for 4 weeks, backup multiplier = 2× primary usage
 - Apply different models for hot, warm, and cold data
 - b. **Consult IT Security and Business Continuity policies to align assumptions:**
 - Validate recovery point objectives (RPO) and retention periods per system criticality
 - Ensure estimates reflect current backup scope and DR design
 - c. **Document additional storage needs in the Storage Forecasting Report:**
 - Use the “Forecasting Assumptions” and “Constraints & Risks” fields to explain calculations and edge cases
 - Where possible, separate redundancy and compliance drivers from core growth projections
- iii. **Evaluate optimization and risk mitigation strategies:** Identify ways to manage the added capacity demand without overprovisioning.
 - a. **Identify opportunities to reduce or consolidate backup volumes:**
 - De-duplicate backup files, adopt tiered retention, or offload older backups to archival platforms
 - Consider cloud-based backup-as-a-service (BaaS) models for long-term retention
 - b. **Review regulatory obligations to optimize compliance storage:**
 - Avoid over-retention of datasets by aligning storage timeframes with actual legal mandates
 - Explore technologies like cold storage, object storage, or nearline cloud tiers
 - c. **Capture all recommended actions in the “Mitigation Strategies” field:**
 - Example: “Implement deduplication and migrate compliance archive to cloud cold tier”
 - Ensure risks and trade-offs are clearly communicated to architecture and finance teams

3.2.E
(1/2)

Estimate Storage-Related Budget Requirements: Translate forecasted storage capacity needs into financial estimates to support procurement planning, and infrastructure budgeting. This step ensures the Entity aligns capacity growth with available funding and cost models. The Architecture and Analytics Section head leads the budgeting process in consultation with the Finance Department, Data Custodians, Data Architects, and Procurement and IT teams.

- i. **Break down forecasted capacity for budget planning purposes:** Structure capacity data to support cost modeling and funding allocations.
 - a. **Use the Year 1, Year 3, and Year 5 estimates captured in previous forecasting steps:**
 - Reference the fields already populated in 3.2.B and 3.2.C
 - Ensure figures are accurate and reflect total required storage including operational, backup, and compliance needs
 - b. **Segment forecasted capacity by deployment environment when relevant:**
 - Identify whether storage is planned for on-premise, public cloud, or hybrid platforms
 - Reference this in the “Forecasting Assumptions” or “Constraints & Risks” fields if not explicitly modeled
 - c. **Use these values to support financial estimates within the same report:**
 - Maintain traceability between volume forecasts and corresponding cost estimates
 - Ensure all projections align with known system boundaries and infrastructure strategies
- ii. **Apply appropriate cost models to derive budget estimates:** Calculate estimated spend using standardized pricing structures.
 - a. **Use cost benchmarks provided by finance or vendors:**
 - For on-premise: cost per TB/year for SAN, NAS, or direct-attached storage
 - For cloud: subscription pricing for object storage, archival tiers, or IaaS volumes
 - b. **Account for additional cost drivers:**
 - Licensing fees, hardware refresh cycles, storage-related support costs, or cloud API/request charges
 - Include redundancy-related costs such as replication, mirroring, or regional failover storage
 - c. **Populate the “Estimated Budget Requirements” field in the Storage Forecasting Report:**
 - Include projected costs for Year 1, Year 3, and Year 5 based on capacity forecasts

3.2.E
Cont.
(2/2)

- Use the “Forecasting Assumptions” field to clarify cost bases, multipliers, or pricing source

iii. **Align budgeting outputs with procurement:** Ensure financial projections support real-world decision-making.

a. **Review estimates with Finance and Procurement teams:**

- Confirm alignment with budgeting calendars and capital vs. operational expenditure models
- Identify whether cloud storage can be funded through cloud credits or subscription budgets

b. **Document risks, constraints, and mitigation strategies in the report:**

- Use “Constraints & Risks” to flag budget uncertainties, delayed procurement, or underestimated needs
- Use “Mitigation Strategies” to outline fallback actions (e.g., phased deployment, cold tier storage migration)

RACI Matrix - Data Storage and Usage Requirements Forecasting

Indicative

Sr No.	Activity	Data Management & Security Director	Data Architecture and Analytics Section Head	Data Architect	Data Custodian	Business Data Analyst	System Owners	Finance Department	IT Team	Procurement Team	Information Security Team	Legal Team
3.2.A	Establish Forecasting Scope and Parameters	A	R	C	C	C	C					
3.2.B	Estimate Initial Storage Capacity Requirements		A	C	R	C		C	C			
3.2.C	Project Multi-Year Growth Trends	I	A	R	C	R		C	C			
3.2.D	Identify Redundancy, Backup, and Compliance Storage Needs		A	C	R						C	C
3.2.E	Estimate Storage-Related Budget Requirements	A	R	C	C				C	C		

RACI Matrix Legend: R: Responsible A: Accountable C: Consulted I: Informed

Based on this Activity, the Entity can



Estimate Storage Capacity Requirements



Estimate Storage Related Budget Requirements

3.3 Target Data Storage and Infrastructure Architecture

This activity focuses on designing the future-state data storage infrastructure that will enable the Entity to meet evolving business needs, and technology shifts,. The target architecture should reflect flexibility, scalability, performance, and resilience while aligning with the Cloud First Policy. It incorporates inputs from the baseline assessment (3.1) and forecasting (3.2) to define a future-ready design that supports digital transformation, disaster recovery, high availability, and compliance with security and data governance standards.

Process – Target Data Storage and Infrastructure Architecture




Figure 6: Target Data Storage and Infrastructure Architecture Process Map

Legend: Process Step Decision Process Start Process End Document Off-Page Reference

3.3.A
(1/2)

Define Design Objectives and Architectural Principles: Set a clear foundation for the design of the target storage architecture by identifying strategic goals, performance and compliance needs, and the guiding principles that will inform all design decisions. This step ensures that the architecture is fit-for-purpose, future-proof, and aligned with national regulations. The Data Architect leads this activity, supported by the Data Architecture and Analytics Section Head, and Data Custodian.

- i. **Define business and technical objectives for the target architecture:** Articulate what the architecture must achieve to support strategic transformation and operational continuity.
 - a. **Identify core business goals driving the new architecture:**
 - **Examples:** enable digital public services, support smart infrastructure, expand analytics capabilities, or reduce infrastructure total cost of ownership
 - b. **Translate goals into technical outcomes for storage:**
 - **Examples:** achieve scalable cloud-native storage, improve RPO/RTO for critical systems, enable cross-site data replication, or reduce latency for real-time data systems
 - c. **Prioritize objectives based on forecasting and baseline insights:**
 - Use pain points and risks flagged in 3.1.E and growth patterns in 3.2.C to shape what the architecture must accommodate
- ii. **Define architectural design principles and guardrails:** Establish the technical and regulations-aligned rules that will govern architecture decisions.
 - a. **Specify key design principles to guide the architecture:**
 - **Examples:** cloud-first by default, modular and scalable design, minimize vendor lock-in, enable automation, support data classification and access control
 - b. **Include national and entity-specific guardrails:**
 - Align with the Qatar Cloud First Policy, data residency and security requirements, and the National Data Standards
 - Include internal thresholds such as minimum redundancy, encryption standards, or cost ceilings per TB
- iii. **Identify scope and boundaries of the target architecture:** Define what is included (and excluded) to avoid overreach or ambiguity.
 - a. **Determine which systems and environments are covered:**

3.3.A
Cont.
(2/2)

- Prioritize those marked as virtualizable or cloud-eligible in 3.1.D
- Exclude legacy systems with no planned modernization unless they impose architectural constraints

b. Clarify scope of environments (e.g., production, DR, backup):

- Decide if test/dev or sandbox environments are included in target architecture
- Define if this architecture covers centralized storage only or also distributed/edge nodes

c. Set temporal and lifecycle boundaries for the design:

- Define whether the architecture represents a 3-year, 5-year, or hybrid planning horizon
- Clarify whether it addresses only steady state or includes transitional architectures

3.3.B
(1/2)

Select Storage Models, Platforms, and Deployment Environments: Define the appropriate mix of storage technologies, platforms, and deployment options to support the Entity's forecasted storage needs and architectural objectives. This step ensures that technology choices are scalable, secure, and optimized for performance and cost across all relevant use cases. The Data Architect leads this step, with contributions from the IT Team, Data Engineers, and Data Custodians.

i. Identify the most suitable storage models for the target state: Choose storage technologies based on data types, performance needs, and business usage patterns.**a. Map storage models to business requirements:**

- Use block storage for transactional databases
- Use file storage for collaborative platforms or content management
- Use object storage for unstructured, archive, or cloud-native workloads

b. Consider architectural traits such as elasticity, durability, and access frequency:

- Apply hot/warm/cold tiering strategies to optimize performance and cost
- Use high IOPS storage for critical workloads and cost-efficient tiers for infrequently accessed data

c. Document each model:

- Include mapping between data types, systems, and recommended storage models
- Note technical constraints or regulatory considerations that affect model selection

ii. Select platform types and technologies aligned to future needs: Decide on underlying platforms and tools that will deliver the storage models selected**a. Evaluate platform options across on-premise and cloud:**

3.3.B
Cont.
(2/2)

- Examples: SAN/NAS, HDFS, object stores (e.g., S3, Azure Blob), distributed file systems
- Factor in manageability, scalability, fault tolerance, and integration with applications

b. Consider adoption of cloud-native services and abstraction layers:

- Examples: backup-as-a-service (BaaS), database-as-a-service (DBaaS), cloud-native snapshotting and encryption
- Explore use of Kubernetes Persistent Volumes, storage APIs, or orchestration tools

c. Reference Qatar-endorsed cloud platforms and tools:

- Prioritize platforms listed as approved CSPs in the Cloud First Policy
- Ensure compatibility with existing cloud credits or central support frameworks

iii. Define the deployment architecture across environments: Decide how storage will be physically or virtually hosted across production, backup, and DR tiers.

a. Select deployment models appropriate to each use case:

- On-premise for systems with high data sensitivity or latency constraints
- Public cloud for scalable, elastic workloads or archived content
- Hybrid deployments where system interoperability or policy requires dual presence

b. Align with Cloud First Policy priorities and data classification rules:

- Data classified as C4 is exempt from this policy; entities are directed to follow the relevant provisions in the NIA Policy and any other relevant laws and regulation

c. Capture deployment environments in the architecture blueprint or system registry:

- Tag each system or domain with its intended storage location (e.g., Azure blob, on-prem SAN, cold archive in S3 Glacier)
- Include fallback options or deployment dependencies (e.g., DR site hosting, cloud region availability)

3.3.C
(1/2)

Design Conceptual and Logical Storage Architecture: Develop the target-state conceptual and logical architecture for the Entity's storage landscape. This includes abstract representations of how data storage is structured, accessed, protected, and scaled across platforms and environments. The conceptual view focuses on high-level storage domains and functions, while the logical view details system relationships, capacity tiers, integration flows, and resiliency designs. This step is led by the Data Architect, in collaboration with the IT Team, and Data Custodian.

- i. **Develop a conceptual view of the target storage landscape:** Provide a high-level visual and narrative structure of storage domains and their purpose.
 - a. **Define primary storage domains in the architecture:**
 - Operational storage for live data used by business systems
 - Analytical storage for reporting, BI, and ML workloads
 - Backup and disaster recovery zones
 - Archival and compliance storage areas
 - b. **Outline how data flows across domains and environments:**
 - Show directionality and dependencies between core systems and storage zones
 - Represent high-level integrations with data ingestion, transformation, and user access layers
 - c. **Ensure conceptual design aligns with business and architectural principles:**
 - Confirm coverage of all data classes and usage scenarios defined in 3.3.A
 - Verify alignment with the Cloud First Policy, and the Data Architecture and Modeling Domain Standards
- ii. **Design the logical storage architecture:** Translate the conceptual model into more detailed, technology-neutral design structures.
 - a. **Define logical groupings of storage systems and tiers:**
 - Break down storage into hot/warm/cold layers based on performance and access patterns
 - Show replication layers, caching tiers, and compression or deduplication zones
 - b. **Map logical components to forecasted capacity and system classifications:**
 - Use input from 3.1.D (portability classification) and 3.2.C (growth trends) to guide sizing and distribution
 - Ensure logical zones accommodate forecasted volumes through Year 5
 - c. **Illustrate system-to-storage relationships and architectural dependencies:**
 - Include interfaces with key systems (e.g., ERP, CRM, data lake)
 - Document connections to cloud platforms, DR sites, and virtualization layers

3.3.C
Cont.
(2/2)

iii. **Integrate security, resiliency, and scalability considerations:** Embed operational and policy-level safeguards directly into the design.

a. **Incorporate data protection and access control layers into the logical view:**

- Show encryption, tokenization, or segmentation by data classification
- Include logical access zones with identity or role-based access overlays

b. **Reflect high availability and disaster recovery mechanisms in the design:**

- Include redundancy models (e.g., replication, failover, active-active clustering)
- Designate logical storage units to primary and secondary data centers or cloud regions

c. **Build in scalability mechanisms at the architectural level:**

- Support horizontal expansion via cloud-native services or modular on-prem components
- Design for elastic provisioning, multi-tenancy, and workload separation

3.3.D
(1/2)

Align with National Cloud First Policy and Technology Standards: Ensure that the target data storage architecture is aligned with the Qatar Cloud First Policy, endorsed cloud platforms, and national regulatory and security requirements. This step ensures that infrastructure decisions reflect national digital transformation priorities, and promote efficiency. It is led by the Data Architect and Data Architecture and Analytics Section Head, in coordination with the Information Security Team, and the Data Compliance Specialist.

i. **Review Cloud First Policy alignment requirements:** Validate that the architecture follows Qatar's strategic direction for cloud adoption.

a. **Confirm cloud-first evaluation across all storage-related decisions:**

- Validate that all in-scope systems have been assessed for cloud suitability using the classification from 3.1.D
- Prioritize cloud deployment for systems that are classified as cloud-eligible or virtualize-able

b. **Ensure use of endorsed Cloud Service Providers (CSPs):**

- Reference the list of CSPs approved by MCIT and confirm that proposed platforms are covered
- Use endorsed IaaS, PaaS, or SaaS offerings where applicable for storage services

c. **Reference the Cloud Policy Framework and guidance in design validation:**

- Include links to cloud governance models, cloud adoption reports, and exit strategy requirements in the architecture documentation

3.3.D
Cont.
(2/2)

- ii. **Embed national technology and security standards into the design:** Ensure compliance with national frameworks including data privacy, information assurance, and cybersecurity policies.
 - a. **Incorporate National Information Assurance (NIA) Policy storage mandates:**
 - Ensure the architecture respects data classification, storage location, encryption, and retention rules
 - Enforce data residency for critical datasets
 - b. **Apply NCSA Cloud Security Policy controls for cloud-based storage:**
 - Address encryption, access controls, incident response, and logging in cloud storage designs
 - Ensure compliance with multitenancy safeguards, vulnerability management, and interoperability protocols
 - c. **Validate alignment with the Personal Data Privacy Protection Law:**
 - Where personal data is stored, ensure compliance with privacy controls, retention limitations, and subject rights
 - Reference these controls in the architecture's constraints or risk documentation
- iii. **Align with national interoperability and architectural principles:** Ensure compatibility and future scalability across government platforms.
 - a. **Align storage architecture with national data standards:**
 - Use standardized metadata formats, data exchange protocols, and API-based storage integration where applicable
 - Ensure that the architecture supports future integration with shared platforms such as national data hubs or analytics environments
 - b. **Leverage centralized tools and shared services where available:**
 - Evaluate reuse of shared infrastructure for backup, disaster recovery, or long-term archival storage

3.3.E

Validate and Document Target Architecture: Consolidate the final conceptual and logical storage architecture design into a structured document and validate it through governance checkpoints. This step ensures all stakeholders agree on the future state before moving into implementation planning. It also guarantees that compliance, scalability, and cloud alignment considerations are properly captured and governed. The Data Architect leads this step, working closely with the Data Architecture and Analytics Section Head, IT Team, and the Data Committee.

i. **Consolidate architecture artifacts and supporting documentation:** Prepare a structured and complete design package for review and sign-off.

a. **Compile all architectural components into a single design document:**

- Include the conceptual architecture (storage domains, flows, dependencies)
- Include the logical architecture (tiers, capacities, platforms, integration points, safeguards)

b. **Attach all supporting assumptions, capacity projections, and compliance references:**

- Forecasted volumes from 3.2
- Technology selections and platform rationale from 3.3.B
- Alignment to Policy and Standards from 3.3.D

c. **Organize documentation for clarity and stakeholder relevance:**

- Include executive summary, visuals/diagrams, and a narrative of key design decisions
- Highlight pain points addressed, business value achieved, and strategic alignment

ii. **Facilitate internal and external architecture reviews:** Validate the architecture through formal governance procedures.

a. **Present the design to the Data Architecture and Analytics Section Head for internal review:**

- Ensure alignment with data architecture, security, and infrastructure standards
- Collect feedback, action items, and required refinements

b. **Submit the architecture to the Data Committee and DGO for governance endorsement:**

- Document feedback and track required changes for compliance with DSO.1.2.3

iii. **Finalize and store the architecture in a controlled repository:** Ensure traceability, version control, and access for downstream planning and audits.

a. **Apply formal versioning and approval metadata to the final design document:**

- Include version number, approval date, names of approvers, and change history
- Use naming conventions aligned with your entity's IT governance standards

RACI Matrix - Target Data Storage Infrastructure Architecture

Indicative

Sr No.	Activity	Data Committee	Data Management & Security Director	Data Architecture and Analytics Section Head	Data Compliance Specialist	Data Architect	Data Custodian	Data Engineer	IT Team	Information Security Team
3.3.A	Define Design Objectives and Architectural Principles		I	A		R	C			
3.3.B	Select Storage Models, Platforms, and Deployment Environments		I	A		R	C	C	C	
3.3.C	Design Conceptual and Logical Storage Architecture		I	A		R	C		C	
3.3.D	Align with National Cloud First Policy and Technology Standards		A	R	R	R				C
3.3.E	Validate and Document Target Architecture	C	A	C		R			C	

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Design Conceptual and Logical Storage Architecture



Validate and Document Target Architecture

3.4 Data Storage & Operations Roadmap

The Entity should establish a Data Storage & Operations Roadmap, that will be used and monitored to implement and activate the Data Storage & Operations initiatives across the Entity. The roadmap shall include, at a minimum:

- Key initiatives, milestones, and timelines needed to transition from the current to the target data storage infrastructure architecture
- The tools that need to be implemented to support Data Storage & Operations
- Resource allocation outlining the budget, personnel, and technology needed for each initiative

Process – Data Storage & Operations Roadmap




Figure 7: Data Storage & Operations Roadmap Process Map

Legend: Process Step Decision Process Start Process End Document Off-Page Reference

3.4.A

Outline Data Storage & Operations Initiatives: Identify and document a comprehensive list of initiatives to improve the Entity's data storage and operations capabilities. This should include both foundational enhancements and forward-looking transformation projects. Examples include optimizing backup and recovery procedures, implementing tiered storage models, migrating legacy systems to cloud-based platforms, and consolidating fragmented storage infrastructure. Ensure each initiative aligns with the Entity's broader data strategy and supports system reliability, cost efficiency, business continuity, and compliance with national data and cloud policies and standards.

3.4.B

Prioritize Initiative Short List: Use the long list of data storage and operations initiatives to develop a prioritized shortlist based on business impact, strategic alignment, and implementation feasibility. Focus on initiatives that enhance storage efficiency, improve backup and recovery reliability, increase system availability, or support cloud readiness. Highlight quick wins—such as optimizing existing backup schedules or decommissioning underutilized storage—that can deliver immediate value with minimal investment. Prioritization ensures that resources are directed to high-value efforts while maintaining momentum through early successes.

3.4.C

Initiate Initiative Card Development and Outline Initiative Activities: Develop initiative cards for each prioritized initiative outlined in the short list, capturing essential activities and steps required to guide implementation. For each initiative, outline the key activities by breaking them down into specific steps and sub-activities that need to be completed in sequence (where applicable). Include descriptions of each step to provide clear guidance on execution, along with any necessary details on tools, techniques, or methodologies that may be required.

Asset Guidance

An **Initiative Card Template** provides a structured format to document key information related to prioritized initiatives, such as initiative overview, key activities, sponsor, owner, stakeholders, timeline, key milestones, budget, technology requirements, dependencies, assumptions, risks, and milestones. The initiative cards provide a structured overview of proposed projects, allowing stakeholders to quickly understand each initiative's purpose, expected impact, and resource requirements.

Illustrative

OVERVIEW	SPONSOR	STAKEHOLDERS
xxx	xxx	xxx
OWNER		
xxx		
KEY ACTIVITIES		
1. Key Activity i. Sub-activity X - Detailed Steps ii. Sub-activity Y - Detailed Steps 2. Key Activity i. Sub-activity X - Detailed Steps ii. Sub-activity Y - Detailed Steps		
TIMELINE		BUDGET (IN QAR)
xxx		QAR xxx
EXPECTED BENEFITS		
• xxx		
KEY MILESTONES		
A. xxx B. xxx C. xxx		
TECHNOLOGY REQUIREMENTS		
• xxx		

OVERVIEW	SPONSOR	STAKEHOLDERS																																																																																																									
xxx	xxx	xxx																																																																																																									
OWNER																																																																																																											
xxxx																																																																																																											
KEY ACTIVITIES		TIMELINE																																																																																																									
<table border="1"><thead><tr><th>#</th><th>Milestones & Key Activities</th></tr></thead><tbody><tr><td>A.</td><td>Milestone</td></tr><tr><td>1.</td><td>Key Activity</td></tr><tr><td>B.</td><td>Milestone</td></tr><tr><td>2.</td><td>Key Activity</td></tr><tr><td>C.</td><td>Milestone</td></tr><tr><td>3.</td><td>Key Activity</td></tr></tbody></table>		#	Milestones & Key Activities	A.	Milestone	1.	Key Activity	B.	Milestone	2.	Key Activity	C.	Milestone	3.	Key Activity	<table border="1"><thead><tr><th>#</th><th>M1</th><th>M2</th><th>M3</th><th>M4</th><th>M5</th><th>M6</th><th>M7</th><th>M8</th><th>M9</th><th>M10</th><th>M11</th><th>M12</th></tr></thead><tbody><tr><td>A.</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>1.</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>B.</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>2.</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>C.</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr><tr><td>3.</td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td><td></td></tr></tbody></table>	#	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12	A.													1.													B.													2.													C.													3.												
#	Milestones & Key Activities																																																																																																										
A.	Milestone																																																																																																										
1.	Key Activity																																																																																																										
B.	Milestone																																																																																																										
2.	Key Activity																																																																																																										
C.	Milestone																																																																																																										
3.	Key Activity																																																																																																										
#	M1	M2	M3	M4	M5	M6	M7	M8	M9	M10	M11	M12																																																																																															
A.																																																																																																											
1.																																																																																																											
B.																																																																																																											
2.																																																																																																											
C.																																																																																																											
3.																																																																																																											
DEPENDENCIES		ASSUMPTIONS & RISKS																																																																																																									
• xxx		• xxx																																																																																																									

3.4.D

Outline Technology Requirements: Identify and document the technology requirements necessary to support each prioritized data storage and operations initiative. This includes outlining the required tools, software, and systems needed to ensure successful implementation and adoption across the Entity.

3.4.E

Document Assumptions: Record assumptions regarding resource availability, stakeholder engagement, and initiative complexity to ensure transparency and guide decision-making.

3.4.F

Analyze Potential Risks: Clearly identify and assess risks associated with document and content management initiatives. This includes:

- Technical risks such as incompatibility with legacy infrastructure, challenges integrating with existing databases or cloud environments, or limitations in performance monitoring and storage virtualization tools
- Operational risks such as disruptions to backup schedules during transition, misconfiguration of storage tiers, or insufficient IT staff readiness to manage new platforms
- Financial risks such as cost overruns due to underestimating cloud storage consumption, delayed procurement of infrastructure, or unanticipated licensing and support fees

3.4.G

Assign Initiatives Ownership: The Data Architecture and Analytics Section Head will be the owner of all document and content management initiatives. This designation ensures that the initiatives have a dedicated leader who can oversee progress and resolve roadblocks that may arise.

3.4.H

Identify Key Stakeholders: Identify individuals and departments with a vested interest in data storage and operations initiatives, including business units, IT infrastructure teams, data custodians, information security specialists, legal advisors, and compliance officers. These stakeholders play a vital role in defining storage requirements, data classification rules, performance expectations, and platform selection criteria. Engaging the right stakeholders early ensures cross-functional alignment, accelerates decision-making, and promotes adoption of the data storage and operations roadmap across the Entity.

3.4.I

Identify Interdependencies: Analyze how data storage and operations initiatives may impact or be influenced by existing or planned activities across the Entity. This includes:

- Reviewing how storage optimization, cloud migration, backup redesign, or disaster recovery efforts intersect with data governance, digital transformation, cybersecurity, or regulatory compliance programs
- Identifying overlapping teams, tools, or infrastructure components involved in multiple initiatives to prevent duplication of effort or configuration conflicts
- Ensuring coordination between business units, IT operations, data management, and information security teams to align priorities, timelines, and resource allocation

3.4.J

Develop Detailed Implementation Timeline: Break down each activity in an initiative into actionable steps with start and end dates, ensuring the timeline is realistic and achievable. Each step should have a clear description of the required task, dependencies on other activities, and allocated resources. Ensure that the timeline accounts for potential delays or risks identified earlier, so adjustments can be made without derailing the project.

3.4.K

Align Milestones with Timeline: Set specific checkpoints or milestones throughout the implementation to assess progress. These milestones allow the evaluation of whether the initiative is proceeding as planned and provide opportunities for course corrections.

Asset Guidance

The **Data Management Domain Plan Template** helps the Entity capture an overview of each identified initiative, detail its sub-activities, owners, stakeholders, and activity timeframe.

Illustrative

Data Management Initiative Plan Template									
The following table is a template capturing a brief overview of the domain relevant initiatives									
Field	Description								
Initiative ID	A unique identifier for each initiative to facilitate tracking and reference								
Initiative Title	the name of the initiative								
Initiative Description	A short description of the initiative, outlining its objectives and expected outcomes								
Cost	Estimated financial resources required for the initiative, including both direct and indirect costs								
Start Date	The date when the initiative is scheduled to begin								
End Date	The date by which the initiative is expected to be completed								
Interdependencies	Any other identified initiatives or projects that this initiative depends on for successful execution								
Assumptions	Conditions assumed to be true for the initiative's success (e.g., resource availability)								
Risks	Potential risks that could affect the initiative's success								

Initiative ID	Initiative Title	Initiative Description	Cost	Start Date	End Date	Interdependencies	Assumptions	Risks

Initiative ID	Initiative Title	Sub-activities	Owner	Stakeholder(s)	Status	Activity Timeframes		Roadmap			
						Start Date	Due Date	0	1	2	3
								1 Jan	8 Jan	15 Jan	###
DMSG-001					In Progress	1 Jul	30 Aug				
DMSG-002					Completed	1 Jul	31 Jan				
DMSG-003					Not Started	30 Mar	31 Jul				

3.4.L

Estimate Costs Associated with Each Initiative: Consider expenses related to tools, technology, personnel, and training required for effective implementation. Break down costs based on the needs of each task, such as acquiring software, hiring additional staff, or providing necessary training for existing teams. This step ensures the budget is well-planned, realistic, and includes all critical aspects to enable smooth implementation of the initiatives.

3.4.M

Develop Roadmap: Create an overall roadmap that outlines timelines, phases, and key milestones for implementing the prioritized document and content management initiatives. Sequence the initiatives based on dependencies, resource availability, and alignment with organizational goals. The roadmap should also include projected timelines for each initiative, allowing stakeholders to monitor progress and allocate resources effectively over time.

3.4.N

Gather Stakeholders' Feedback: Present the draft plan to all relevant parties, inviting input and suggestions to enhance its effectiveness. This feedback process helps surface any potential issues or overlooked areas, improving the plan's effectiveness and ensuring all voices are heard. Engaging stakeholders early also builds buy-in and commitment to the plan's success.

3.4.O

Incorporate Stakeholders' Feedback: Actively gather and incorporate constructive feedback from key stakeholders to refine the plan and ensure it aligns with their needs and expectations. This collaborative approach helps to address concerns early, fine-tune the details, and strengthen the overall effectiveness of the plan.

3.4.P

Obtain Approval of Data Storage & Operations Plan: After incorporating all feedback, the Data Management and Security Director will be presented with the finalized Data Storage & Operations Plan for review and approval. This ensures alignment with the Entity's priorities and secures the necessary support for successful implementation. Once approved, the plan can be formally adopted and published.

3.4.Q

Share Data Storage & Operations Plan with DGO: Submit the finalized Data Storage & Operations Plan to the National Data Governance Office (DGO) for review. The DGO will review the plan and provide recommendations, allowing the Entity to ensure their plan aligns with the broader national data program agenda and best practices.

RACI Matrix – Data Storage & Operations Roadmap

Indicative

Sr No.	Activity	DGO	Data Management and Security Director	Data Architecture and Analytics Section Head	Data Architect	Data Custodian	Data Strategy Expert
3.4.A	Outline Data Storage & Operations Initiatives		A	C	R	C	C
3.4.B	Prioritize Initiative Short List		A	C	R	C	C
3.4.C	Initiate Initiative Card Development and Outline Initiative Activities			A	R		C
3.4.D	Outline Technology Requirements			A	R	C	C
3.4.E	Document Assumptions			A	R	C	C
3.4.F	Analyze Potential Risks			A	R		C
3.4.G	Assign Initiatives Ownership		A	R	I		
3.4.H	Identify Key Stakeholders		A	R	C	C	
3.4.I	Identify Interdependencies			A	R		C
3.4.J	Develop Detailed Implementation Timeline			A	R		C
3.4.K	Align Milestones with Timeline			A	R		C
3.4.L	Estimate Costs Associated with Each Initiative			A	R	C	C
3.4.M	Develop Roadmap			A	R		C
3.4.N	Gather Stakeholders' Feedback			A	R		
3.4.O	Incorporate Stakeholders' Feedback			A	R		C
3.4.P	Obtain Approval of Data Storage & Operations Plan		A	R			
3.4.Q	Share Data Storage & Operations Plan with DGO	C/I	A	R			

RACI Matrix Legend: R: Responsible A: Accountable C: Consulted I: Informed

Based on this Activity, the Entity can



Outline Data Storage & Operations Initiatives



Develop and Review Data Storage & Operations Plan

4. Data Storage and Operations Implementation



4. Data Storage and Operations Implementation


What is Data Storage and Operations Implementation and why does the Entity need it?

The Data Storage & Operations Implementation process provides a structured approach to deploying secure, resilient, and regulations-compliant data infrastructure and operational procedures. It ensures that the designs and plans developed during the planning phase are executed consistently, with attention to reliability, security, and service continuity.

This process enables the Entity to put in place robust access controls, enforce disciplined backup strategies, and establish tested recovery protocols that minimize downtime and support business resilience. It connects infrastructure modernization with day-to-day operational readiness by translating standards into practical, executable procedures.

Through activities such as access provisioning, backup configuration, disaster recovery planning, and roadmap activation, the Entity can improve system reliability, enforce compliance, and reduce operational risk. These implementation efforts drive lasting impact by operationalizing the national data standards, enhancing continuity of service, and supporting secure digital transformation.

The remaining sections of this chapter will detail the below activities that guide the Entity through executing the Data Storage and Operations Implementation process.



Guiding Principles for Data Storage and Operations Implementation

1. **Secure by Design:** Integrate role-based access control, data classification alignment, and encryption practices into every aspect of storage and operations implementation.
2. **Ensure Recoverability:** Treat backup, restoration, and disaster recovery as core operational functions—not emergency procedures—and validate them regularly.
3. **Activate What's Planned:** Maintain fidelity to the designs and strategies defined during planning and architecture, ensuring clear traceability between roadmap intent and execution.
4. **Test Before Trust:** Validate all backup and recovery procedures through controlled testing and simulation before declaring them operational.
5. **Govern Through Visibility:** Track implementation progress transparently through milestones, reviews, and reporting to drive accountability and sustain momentum.

4.1 Database Access Control

Database Access Control ensures that only authorized personnel have access to data assets, based on clearly defined roles and responsibilities. It plays a central role in protecting sensitive information, ensuring privacy compliance, and minimizing the risk of data breaches or unauthorized data manipulation. This process enables the Entity to enforce consistent, auditable access policies that align with business roles, system criticality, and data classification levels.

Process – Database Access Control




Figure 8: Database Access Control Process Map

Legend: Process Step Decision Process Start Process End Document Off-Page Reference

4.1.A
(1/2)

Define Role-Based Access Control (RBAC) Framework: Establish a standardized RBAC framework that defines roles, access levels, and responsibilities across the Entity's database environments. This foundational step ensures database access is consistent, secure, and aligned with data sensitivity, job functions, and compliance mandates. The Information Security and Risk Management Section Head leads this step, with collaboration from the Data Custodian, Data Steward, and IT Team.

- i. **Identify and classify database roles across the Entity:** Define the roles that require access to databases and clarify their function-specific privileges.
 - a. **Categorize users by functional roles and access needs:**
 - Examples include Application Users, Data Analysts, Developers, System Integrators, and Auditors
 - Align each role with operational responsibility, technical expertise, and data classification level
 - b. **Determine access types per role:**
 - Read-only, write/edit, admin, export, configuration, or schema management
 - Include privileges such as access to logs, backups, stored procedures, or metadata
 - c. **Group roles into logical tiers (illustrative):**
 - **Tier 1:** General users with read-only access
 - **Tier 2:** Operational staff with write/update privileges
 - **Tier 3:** Admins with elevated or full access
- ii. **Define access control principles and assignment rules:** Set the foundational RBAC rules that govern how access is granted, limited, or escalated.
 - a. **Apply least privilege and need-to-know principles:**
 - Ensure users have the minimum access required to perform their duties
 - Prohibit blanket or default access to entire schemas or databases
 - b. **Align access with data classification and criticality:**
 - Limit access to Secret and Top Secret data to vetted and approved roles only
 - Define specific controls for access to sensitive tables, encrypted fields, or audit logs
 - c. **Define role ownership and approval responsibilities:**
 - Assign accountability for each role to a system or business owner
 - Establish approval workflows for role creation, modification, or removal

4.1.A
Cont.
(2/2)

iii. **Document the RBAC matrix and governance structure:** Capture the framework in a structured, reviewable, and maintainable format.

a. **Develop an RBAC matrix mapping roles to database access rights:**

- Use a table format showing each role and its approved privileges across relevant environments (e.g., dev, test, prod)
- Include role descriptions and example use cases

b. **Assign review responsibility and update frequency:**

- Define how often the RBAC matrix must be reviewed (e.g., semi-annually or during system onboarding)
- Assign review responsibility to Data Custodians and Information Security teams

c. **Store the RBAC framework in a secure, version-controlled repository:**

- Ensure the matrix is accessible to auditors, security officers, and designated admins
- Track historical versions for audit traceability



Key Message / Information

Database access should be granted exclusively through authorized applications to prevent direct access to the underlying database storage.

Allowing direct access bypasses key security controls such as authentication, authorization, input validation, and logging — increasing the risk of data breaches, unauthorized modifications, or accidental data loss. Accessing data through an application ensures that all interactions are governed by business logic, role-based permissions, and audit trails, providing a controlled and secure layer between users and sensitive data.

4.1.B
(1/2)

Implement Access Provisioning and Revocation Procedures: Establish standardized, auditable procedures for granting, modifying, and revoking database access. This ensures that access is tightly controlled throughout the user lifecycle and that changes are consistently tracked and aligned with the RBAC framework. The Information Security and Risk Management Section Head leads this activity in coordination with Data Custodians, IT Team, and the Data Compliance Specialist.

i. **Define procedures for access request and approval:** Ensure that database access is provisioned through structured and authorized workflows.

a. **Establish a formal request process for new access:**

- Require submission of a standardized access request form or IT ticket
- Capture user role, justification, system, and data sensitivity level

4.1.C
(1/2)

Integrate Access Control with Identity and Audit Systems: Establish integration between database access control mechanisms and the Entity's identity management, audit logging, and monitoring systems. This ensures real-time enforcement, centralized oversight, and traceable auditability of database access activities. The Information Security and Risk Management Section Head leads this step in collaboration with Data Custodian, IT Team, and Internal Audit.

- i. **Connect database access with enterprise identity management systems:** Ensure that user roles and permissions are centrally governed and traceable.
 - a. **Integrate database access with the Entity's Identity Access Management (IAM) or directory services:**
 - Examples include Active Directory (AD), LDAP, or Azure AD
 - Use group-based provisioning where possible to enforce consistency
 - b. **Enforce identity lifecycle controls through IAM tools:**
 - Ensure access is automatically adjusted when users change roles or leave the organization
 - Link provisioning and deprovisioning workflows with HR systems or organizational changes
 - c. **Maintain a centralized user-role mapping registry:**
 - Store the RBAC assignments for all database-connected users in a single access inventory
 - Use this as the source of truth during access reviews and audits
- ii. **Enable continuous audit logging and access monitoring:** Ensure that all access events are captured, monitored, and reviewable for compliance.
 - a. **Configure DBMS to generate detailed audit logs:**
 - Track events such as logins, privilege escalations, schema changes, data exports, or failed login attempts
 - Enable audit logs for both user and system accounts, especially for admin actions
 - b. **Forward logs to a centralized logging and monitoring platform:**
 - Integrate with SIEM systems (e.g., Splunk, IBM QRadar, Microsoft Sentinel) for real-time monitoring
 - Set alerts for anomalous behavior, privilege abuse, or access policy violations
 - c. **Review audit logs regularly in collaboration with Security and Compliance:**
 - Schedule monthly or quarterly log reviews for high-risk databases
 - Use logs during incident investigations, data breach response, or compliance audits

4.1.C
Cont.
(2/2)

iii. **Establish governance and escalation mechanisms:** Ensure ongoing accountability, issue resolution, and compliance enforcement.

a. **Define access governance roles and responsibilities:**

- Assign ownership of each system's access model to Data Custodians or System Owners
- Clarify roles for access reviewers, incident responders, and audit liaisons

b. **Develop escalation procedures for access-related incidents or violations:**

- Include pathways for handling unauthorized access, overdue revocations, or unapproved privilege changes
- Ensure escalation paths reach Security, Legal, and Executive Management as needed

c. **Use access control metrics to inform governance reporting:**

- Track metrics such as number of access requests granted/denied, privileged account usage, and number of revocation delays
- Include these metrics in security dashboards or governance board updates

RACI Matrix – Database Access Control

Indicative

Sr No.	Activity	Data Management & Security Director	Information Security and Risk Management Section Head	Data Compliance Specialist	Internal Audit	Data Custodian	Data Steward	IT Team
4.1.A	Define Role-Based Access Control (RBAC) Framework	A	R			C	C	C
4.1.B	Implement Access Provisioning and Revocation Procedures	A	R	C		C		C
4.1.C	Integrate Access Control with Identity and Audit Systems	A	R		C	R		C

RACI Matrix Legend: R: Responsible A: Accountable C: Consulted I: Informed

Based on this Activity, the Entity can



Define RBAC Framework



Integrate Access Control with Audit Systems

4.2 Data Backup & Recovery

Data Backup & Recovery ensures the ongoing availability, integrity, and recoverability of the Entity's critical data assets. It provides the technical foundation for operational resilience, disaster recovery, and long-term continuity of services. This activity translates storage and recovery policies into daily operational procedures, guaranteeing that critical systems and databases are regularly backed up, verified for integrity, and recoverable within defined timeframes.

4.2.A (1/2)

Define Backup Strategy and Frequency per System: Establish a structured and risk-aligned backup strategy that defines the frequency and type of backups for each information system based on business criticality, data classification, and recovery objectives. This ensures consistent protection of critical data and supports rapid recovery in the event of failure or data loss. The Data Custodian leads this activity, in collaboration with the Information Security team, System Owners, and IT Team.

- i. **Classify systems and datasets by backup priority:** Align backup strategy with business value, operational risk, and regulatory requirements.
 - a. **Categorize systems based on business criticality and data sensitivity:**
 - Mission-critical systems (e.g., ERP, financial, national registries)
 - Operational systems (e.g., HR, CRM, public service platforms)
 - Low-risk or archival systems (e.g., sandbox, logs, historical datasets)
 - b. **Map system categories to data classification levels:**
 - Ensure Secret and Top Secret systems are prioritized for frequent and secure backup
 - Consider legal retention requirements and breach exposure risks
 - c. **Define priority tiers for backup planning:**
 - **Tier 1:** Mission-critical, high-frequency backup
 - **Tier 2:** Moderate criticality, regular backup
 - **Tier 3:** Low-risk, infrequent or archival backup
- ii. **Define appropriate backup frequency and retention periods:** Tailor backup intervals to the system's operational profile and recovery needs.
 - a. **Set daily, weekly, or monthly backup frequencies based on system classification (illustrative):**
 - Daily or near real-time for Tier 1 systems
 - Weekly for Tier 2 systems
 - Monthly or quarterly for archival and Tier 3 systems
 - b. **Establish retention durations based on regulatory, operational, and legal requirements:**
 - Example: retain daily backups for 30 days, weekly for 3 months, monthly for 1 year
 - Comply with retention mandates for regulated datasets (e.g., financial, personal, medical)

4.2.A
Cont.
(2/2)

iii. **Align backup strategy with Recovery Objectives (RTO/RPO):** Ensure backup timing and design meet the Entity's recovery targets.

- a. **Review and confirm Recovery Point Objectives (RPO) per system (Detailed in 4.3):**
 - Define how much data loss is acceptable (e.g., 0 minutes, 4 hours, 1 day)
 - Tie backup frequency to these thresholds
- b. **Confirm Recovery Time Objectives (RTO) for restoring each system (Detailed in 4.3):**
 - Define how quickly the system must be back online after failure
 - Use RTOs to inform backup format, location, and resource allocation
- c. **Validate RTO and RPO alignment with business continuity and SLA documentation (Activity 5.1):**
 - Ensure documented targets are feasible and backed by tested infrastructure
 - Update targets as needed based on business or policy changes

4.2.B
(1/2)

Configure Backup Scope and Storage Architecture: Define what data must be included in backup operations and where that data should be stored to ensure resilience, performance, and compliance. This step ensures that backups are comprehensive, aligned with system-specific needs, and stored in environments that support recovery, retention, and security objectives. The Data Custodian leads this activity in collaboration with the Information Security team, IT Team, and System Owners.

- i. **Define the scope of backup content for each system:** Determine which data components need to be protected and how comprehensively.
 - a. **Identify critical data elements to include in each system's backup scope:**
 - Databases, transaction logs, configuration files, metadata, and system parameters
 - Include both structured and unstructured data where applicable
 - b. **Tailor the scope based on system type and business function:**
 - **For databases:** include full data tables and transaction logs
 - **For application systems:** include runtime configurations and supporting directories
 - **For file-based systems:** refer to National Document and Content Management Standards and Guidelines
 - c. **Document the backup scope in system-specific backup profiles:**
 - Use a standardized format that defines what to back up, who owns the data, and when it should be reviewed
 - Store backup profiles in a centralized repository for operational and audit reference

- ii. **Select appropriate backup storage media and location:** Ensure storage architecture supports durability, scalability, and recovery needs.
 - a. **Choose storage types based on system tier and access requirements:**
 - SAN/NAS for high-speed recovery
 - Object storage for long-term retention
 - Tape or cloud-based cold storage for infrequent access or archival backups
 - b. **Implement geographically separate storage for resiliency:**
 - Store primary backups on a RAID-configured SAN or NAS
 - Use offsite or cloud storage for secondary copies to ensure disaster recovery readiness
 - Avoid storing operational data and backup copies on the same media
 - c. **Ensure backups comply with encryption, immutability, and data segregation policies:**
 - Encrypt backups at rest and in transit for Secret and Top Secret datasets
 - Apply write-once-read-many (WORM) or immutability settings for compliance backups
 - Segment backup storage environments based on data classification
- iii. **Establish storage capacity and performance plans for backups:** Ensure infrastructure can accommodate growing backup volumes without service degradation.
 - a. **Estimate backup storage volume requirements:**
 - Calculate expected volume growth based on backup frequency, retention, and system size
 - Factor in deduplication, compression, and replication techniques to optimize space usage
 - b. **Monitor storage performance and availability:**
 - Ensure backup jobs complete within acceptable timeframes
 - Set alerts for capacity thresholds or storage failure risks
 - c. **Align backup storage architecture with overall data storage requirements:**
 - Integrate with the architecture designed in Activity 3.3
 - Coordinate with IT Infrastructure and Cloud Governance teams for hybrid or cloud-based implementations

4.2.C
(1/2)

Validate Backup Integrity and Recovery Capabilities: Establish procedures to regularly verify that backups are complete, restorable, and capable of meeting the Entity's business continuity and disaster recovery requirements. This step provides assurance that backup processes are not only implemented, but effective in real-world recovery scenarios. The Data Custodian leads this activity in collaboration with the IT Team, and Information Security Team.

- i. **Conduct routine validation of backup integrity:** Ensure that backup files are complete, uncorrupted, and consistent with production data.
 - a. **Establish a schedule for periodic backup file validation:**
 - Weekly or monthly integrity checks for Tier 1 systems
 - Quarterly checks for lower-tier systems or archival data
 - b. **Perform checksum or hash-based comparisons of backup files:**
 - Use tools to verify that backup files have not been altered or corrupted
 - Confirm the consistency of incremental and full backup sets
 - c. **Document validation results in a Backup Integrity Log:**
 - Record date, system, method used, validation outcome, and responsible reviewer
 - Store logs in a centralized repository for internal audit and compliance reviews
- ii. **Perform recovery testing in controlled environments:** Simulate real recovery scenarios to verify that systems can be restored fully and correctly.
 - a. **Restore backups into non-production (test or staging) environments:**
 - Test system state accuracy, data completeness, and application operability
 - Use different recovery points (e.g., daily vs. weekly backup) to test rollback accuracy
 - b. **Evaluate system behavior post-restoration:**
 - Confirm database consistency, application connectivity, and permissions
 - Identify any missing dependencies or configurations
 - c. **Schedule recovery drills for mission-critical systems:**
 - Conduct annual or bi-annual full recovery exercises for Tier 1 platforms
 - Involve both technical and business continuity teams in simulation planning and execution
- iii. **Track recovery performance and compliance metrics:** Ensure that recovery times and data recovery points align with defined objectives.

4.2.C
Cont.
(2/2)**a. Measure Recovery Time Actuals (RTA) and Recovery Point Actuals (RPA):**

- Compare these metrics against target RTOs and RPOs
- Log deviations and determine root causes for delayed recovery

b. Report backup and recovery success metrics to governance stakeholders:

- Include recovery validation status in monthly or quarterly IT operations reports
- Escalate failures, delays, or repeated inconsistencies to the Information Security team and Data Governance leads

c. Update backup and recovery procedures based on test findings:

- Refine retention policies, schedules, or tools if recovery issues are observed
- Incorporate corrective actions into system-specific backup plans

RACI Matrix – Data Backup & Recovery

Sr No.	Activity	Data Management & Security Director	Information Security and Risk Management Section Head	Data Custodian	Information Security Team	IT Team	System Owners
4.2.A	Define Backup Strategy and Frequency per System	I	A	R	C	C	C
4.2.B	Configure Backup Scope and Storage Architecture		A	R	C	C	C
4.2.C	Validate Backup Integrity and Recovery Capabilities	I	A	R	C	C	

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Define Backup & Recovery Frequency



Validate Capabilities

4.3 Disaster Recovery & Business Continuity

Disaster Recovery & Business Continuity ensures that the Entity can rapidly recover from unplanned disruptions, system failures, or data loss events while maintaining continuity of critical operations. This process involves the development, validation, and maintenance of recovery procedures for key systems and datasets, aligned with business impact, recovery time targets, and national security regulations.

4.3.A
(1/2)

Define System Priorities and Recovery Objectives: Establish the foundation for disaster recovery planning by classifying systems according to their business criticality and defining Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO). This step ensures recovery planning is risk-based, aligned with organizational impact tolerance, and tailored to the operational role of each system. The Information Security and Risk Management Section Head leads this activity, with support from System Owners, and Data Custodians.

- i. **Classify systems based on operational and business impact:** Group systems by their importance to mission-critical services, revenue, or regulatory obligations.
 - a. **Identify all active systems and applications across the Entity:**
 - Include databases, APIs, core platforms, SaaS solutions, and citizen-facing services
 - Use inputs from prior assessments in Section 3.1 and 3.3
 - b. **Assign each system a priority level (High, Medium, Low):**
 - **High:** Direct impact on essential services or public operations
 - **Medium:** Supportive systems with internal operational impact
 - **Low:** Non-critical or easily restorable services
 - c. **Record system priority in the Disaster Recovery Playbook:**
 - Populate the “Priority Level” field for each system using the defined criteria
 - Ensure this aligns with the Entity’s broader business continuity classifications

Asset Guidance

The **Disaster Recovery Playbook Template** can be leveraged by the Entities to create a business continuity plan and minimize downtime.

Illustrative

Disaster Recovery Playbook					
This template will help you to create a business continuity plan and minimize downtime by providing a structured recovery framework.					
Field	Description	Example			
System/Application	Name of the system or application	Customer Database			
Priority Level	High, Medium, Low	High			
RPO (Recovery Point Objective)	Maximum acceptable data loss (e.g., 15 mins, 1 hour)	15 minutes			
RTO (Recovery Time Objective)	Maximum acceptable downtime (e.g., 1 hour, 4 hours)	1 hour			
Backup Location	Location of backups (e.g., SAN, AWS S3, Off-site)	AWS S3 - Geo-Redundant			
Recovery Steps	Step-by-step instructions for recovery	1. Access backup in AWS S3. 2. Restore database to test environment. 3. Validate data integrity.			
Dependencies	Other systems or resources required for recovery	Application Server, Network Configuration			
Assigned Personnel	Name and contact information of the responsible individual(s)	John Doe, johndoe@example.com			
Escalation Contacts	Contact details for escalation (e.g., DR Manager, IT Director)	Jane Smith, DR Manager, janesmith@example.com			
Validation Checklist	Checklist for verifying the system after recovery	1. Verify database restoration			
Last Updated	Date the playbook was last reviewed or updated	25-Aug-24			

System/Application	Priority Level	RPO (Recovery Point Objective)	RTO (Recovery Time Objective)	Backup Location	Recovery Steps

Dependencies	Assigned Personnel	Escalation Contacts	Validation Checklist	Last Updated

4.3.A
Cont.
(2/2)

- ii. **Define Recovery Point Objectives (RPO) per system:** Determine acceptable data loss in the event of system failure or disaster.
 - a. **Consult with business and IT stakeholders to set realistic RPO targets:**
 - Examples: 0 minutes for financial systems, 15 minutes for transactional systems, 24 hours for archival systems
 - b. **Align RPO with backup schedules and system volatility:**
 - Ensure that backup frequency and retention strategies (from 4.2) support the defined RPOs
 - Adjust RPO upward for systems with minimal daily changes, and downward for real-time platforms
 - c. **Populate the “RPO” field in the Disaster Recovery Playbook for each system:**
 - Use standardized formats (e.g., 5 min, 1 hour, 1 day) for consistency and traceability
- iii. **Define Recovery Time Objectives (RTO) per system:** Determine how quickly each system must be restored after an incident.
 - a. **Establish RTOs based on business continuity thresholds and SLA (Activity 5.1) obligations:**
 - Examples: 1 hour for public services, 4 hours for internal systems, 48 hours for low-impact systems
 - b. **Consider infrastructure dependencies and recovery complexity:**
 - Account for time needed to provision environments, restore backups, test functionality, and reconfigure integrations
 - c. **Populate the “RTO” field in the Disaster Recovery Playbook:**
 - Ensure RTO values are validated by both IT and business teams to reflect actual impact tolerance

4.3.B
(1/2)

Develop and Document System-Level Recovery Procedures: Translate recovery priorities and objectives into clear, actionable playbooks for each critical system. This step ensures that in the event of a disruption, personnel can execute predefined recovery steps quickly and effectively, reducing downtime and ensuring compliance with RTO/RPO targets. The Information Security and Risk Management Section Head leads this activity in coordination with System Owners, Data Custodians, and the IT Team.

- i. **Document backup locations and recovery responsibilities:** Ensure recovery plans specify where data is stored and who is accountable for restoration.
 - a. **Specify backup storage location for each system:**
 - Examples: on-prem RAID array, SAN, AWS S3, Azure Blob, offsite vault
 - Indicate whether backups are geo-redundant or region-specific
 - b. **Assign named personnel responsible for executing the recovery:**
 - Include data engineers, system administrators, and infrastructure leads
 - Ensure contact details are up to date for time-sensitive execution
 - c. **Populate the “Backup Location” and “Assigned Personnel” fields in the Disaster Recovery Playbook:**
 - Link to the most recent backup schedule or validation logs if needed
- ii. **Define and document step-by-step recovery instructions:** Provide clear, system-specific recovery tasks that are executable under pressure.
 - a. **Develop recovery steps tailored to each system’s architecture and environment:**
 - Examples: Restore VM, reattach storage volume, restore database, restart services, validate application connections
 - b. **Include environment-specific variations or dependencies:**
 - Note whether recovery differs for cloud vs. on-prem or production vs. DR environment
 - Identify sequences that require external approvals or API integrations
 - c. **Populate the “Recovery Steps” and “Dependencies” fields in the Playbook:**
 - Use numbered, step-by-step format with bullet points for clarity
 - Ensure logical order and include reference links or scripts where applicable
- iii. **Define escalation and validation protocols:** Ensure issues are quickly escalated and system integrity is verified after recovery.

4.3.B
Cont.
(2/2)

- a. **Identify escalation contacts and chain of command for recovery failures:**
 - Include IT Director, Legal or Executive stakeholders
 - Provide names, roles, and multiple contact channels
- b. **Include a validation checklist to confirm successful system recovery:**
 - Examples: test user login, verify data integrity, confirm job processing, check audit logs
 - Tailor checklist to the business function and system complexity
- c. **Populate the “Escalation Contacts” and “Validation Checklist” fields in the Playbook:**
 - Ensure responsible parties are aware of their validation duties
 - Version-control the checklist and tie it to the most recent backup validation results

4.3.C
(1/2)

Validate, Test, and Maintain the Recovery Plan: Regularly test the effectiveness of disaster recovery playbooks and ensure that documented procedures remain accurate and up to date. This step provides confidence that the Entity can recover its systems and data within the required timeframes, and supports auditability, compliance, and operational readiness. The Information Security and Risk Management Section Head leads this activity, with collaboration from System Owners, IT Team, and Internal Audit.

- i. **Conduct periodic recovery testing and simulation exercises:** Validate recovery procedures in practice through scheduled testing.
 - a. **Perform full and partial disaster recovery simulations for critical systems:**
 - Examples include restoring databases, launching cloud failover environments, or executing full DR switchovers
 - Test under controlled conditions with minimal production disruption
 - b. **Track recovery time and data accuracy during tests:**
 - Measure actual RTO/RPO performance against documented targets
 - Identify any gaps, missing steps, or resource constraints
 - c. **Log test outcomes and lessons learned:**
 - Maintain detailed records of test dates, outcomes, success rates, and responsible participants
 - Use this data to update procedures, address weaknesses, and inform future planning

4.3.C
Cont.
(2/2)

- ii. **Review and update the Disaster Recovery Playbook routinely:** Ensure that all system-level recovery documentation remains current and usable.
 - a. **Set review frequencies based on system priority:**
 - High-priority systems: quarterly or bi-annual review
 - Medium and low-priority systems: annual review
 - b. **Update fields such as RTO/RPO, assigned personnel, backup locations, and dependencies:**
 - Ensure changes in infrastructure, vendors, or contacts are promptly reflected
 - Cross-check with backup strategy updates (from 4.2) and architectural changes (from 3.3)
 - c. **Track revision history and maintain version control:**
 - Update the “Last Updated” field in the Playbook after each revision
 - Keep archived versions for audit traceability
- iii. **Ensure governance, accountability, and continual improvement:** Establish oversight mechanisms and continuous refinement processes.
 - a. **Assign responsibility for plan ownership and oversight:**
 - Designate DR Coordinators or System Owners as accountable for each playbook entry
 - Establish escalation to IT governance or crisis management teams for unresolved issues
 - b. **Involve internal or external auditors in periodic plan reviews:**
 - Include DR testing and playbook reviews in annual audit scope
 - Address any findings or non-conformities through corrective action plans
 - c. **Integrate lessons learned into Entity-wide resilience planning:**
 - Use DR test results to enhance infrastructure design, backup architecture, or user training
 - Share insights across departments to elevate disaster preparedness maturity

RACI Matrix – Disaster Recovery & Business Continuity

Indicative

Sr No.	Activity	Data Management & Security Director	Information Security and Risk Management Section Head	Data Custodian	Internal Audit	IT Team	System Owners
4.3.A	Define System Priorities and Recovery Objectives	A	R	C			C
4.3.B	Develop and Document System-Level Recovery Procedures	A	R	C		C	C
4.3.C	Validate, Test, and Maintain the Recovery Plan	A	R		C	C	C

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Define System Priorities and Recovery Objectives



Validate Recovery Plan

4.4 Data Storage & Operations Roadmap Activation

With the Data Storage & Operations initiatives fully defined earlier—including detailed activities, timelines, technology needs, risks, dependencies, and budget—this activity ensures the roadmap is activated through disciplined execution, tracked via structured milestones, and updated in coordination with Entity and national governance bodies.

4.4.A

Initiate Execution and Governance Routines: Formally launch the execution of the Data Storage & Operations roadmap by activating governance routines, assigning initiative oversight roles, and initiating regular reporting cycles. This step ensures that the implementation of prioritized initiatives progresses in a structured and accountable manner, without duplicating what is already defined in the initiative cards. The Data Architecture and Analytics Section Head leads this step in collaboration with the Data Architect, Data Custodian, Initiative Owners, and IT Team.

- i. **Formally launch execution across all approved initiatives:** Trigger implementation based on previously developed initiative cards.
 - a. **Confirm budget availability, team assignments, and tool readiness:**
 - Reconfirm that each initiative in the roadmap is approved for execution
 - Reconfirm resource allocations for internal staff, vendors, and third-party partners
 - Ensure required platforms, environments, or licenses are in place
 - b. **Communicate initiative scope and execution start across teams:**
 - Notify stakeholders of approved priorities, delivery windows, and responsibilities
 - Use internal briefings, kick-off meetings, or project charters to establish alignment
- ii. **Activate governance routines for execution oversight:** Establish a rhythm for reporting, coordination, and issue management.
 - a. **Launch a recurring initiative governance forum:**
 - Schedule monthly or bi-weekly roadmap review meetings with initiative owners
 - Use these sessions to review progress, unblock delays, and share risks across initiatives
 - b. **Define reporting cadence and documentation requirements:**
 - Standardize progress update formats (e.g., % complete, blockers, next steps)
 - Align reporting with milestone tracking and executive dashboards
 - c. **Monitor alignment with strategic goals and compliance expectations:**
 - Ensure initiatives remain aligned to national regulations (e.g., Cloud First, NIA, NCSA)
 - Use roadmap reviews to highlight changes that may require governance escalation

A

Monitor Initiative Execution and Track Key Milestones: Track progress of Data Storage & Operations initiatives using predefined milestones and structured reporting. This step ensures that initiatives are progressing as planned, allows early identification of delays or risks, and provides visibility to stakeholders. The Data Architecture and Analytics Section Head leads this step in collaboration with the Data Architect, Initiative Owner, and Data Custodian.

- i. **Use initiative cards and dashboards for progress monitoring:** Leverage the initiative cards as the central source for execution tracking.
 - a. **Monitor key activities and completion status for each initiative:**
 - Use the “Key Activities” section to track task-level progress
 - Flag delays or incomplete tasks that may impact downstream milestones
 - b. **Update initiative timelines based on actual progress:**
 - Reflect changes directly in the timeline section of each card
 - Highlight revised start/end dates in PMO dashboards
 - c. **Consolidate progress reports into a central tracking view:**
 - Use a shared dashboard or spreadsheet to display % completion, risk level, and owner status across all initiatives
 - Make updates accessible to stakeholders for transparency
- ii. **Track milestone delivery and escalate delays:** Monitor milestone completion and manage issues proactively.
 - a. **Review initiative milestones at regular intervals:**
 - Validate achievement of milestones as defined in initiative cards
 - Compare actual timing against planned delivery windows
 - b. **Flag missed or at-risk milestones for review:**
 - Require initiative owners to document causes, impact, and proposed mitigation
 - Escalate high-impact delays to the Data Management and Security Director
 - c. **Use milestone reviews to assess roadmap alignment and adjust priorities:**
 - Revisit initiative sequencing, dependencies, or resource allocation based on milestone performance
 - Document updates in roadmap planning tools or initiative logs
- iii. **Communicate progress to governance and executive stakeholders:** Maintain visibility and accountability throughout execution.

4.4.B
Cont.
(2/2)**a. Prepare periodic roadmap status reports:**

- Include summary of completed, ongoing, and delayed initiatives
- Provide commentary on progress, risks, and cross-cutting issues

b. Present updates during governance review forums:

- Share dashboards or executive briefings with the Data Committee or senior leadership
- Use feedback to refine execution plans and remove bottlenecks

c. Ensure traceability between initiative-level tracking and roadmap-level oversight:

- Link initiative progress updates to strategic goals and planned outcomes
- Use milestone tracking to inform reporting to the DGO

4.4.C
(1/2)

Review, Update, and Submit Roadmap for Governance Oversight: Maintain the relevance and accuracy of the Data Storage & Operations roadmap by reviewing it regularly, incorporating execution feedback, and submitting updates to internal and external oversight bodies. This step promotes transparency, alignment, and continuous improvement throughout implementation. The Data Architecture and Analytics Section Head leads this activity in collaboration with the Data Architect, and Initiative Owners.

i. Conduct periodic roadmap reviews and incorporate feedback: Ensure the roadmap reflects up-to-date progress, changes, and lessons learned.**a. Schedule roadmap reviews aligned with milestone checkpoints:**

- Conduct reviews quarterly or after major implementation phases
- Use these sessions to validate completed work, delays, and realigned delivery dates

b. Update initiative cards and roadmap documents based on actual execution:

- Adjust timelines, dependencies, resource allocations, and risk statuses
- Document new insights, blockers, or scope changes identified during execution

c. Capture and incorporate stakeholder feedback into roadmap refinements:

- Engage initiative leads, technical teams, and compliance stakeholders
- Ensure changes are validated through governance forums before formal submission

ii. Align roadmap changes with strategic and compliance expectations: Ensure adjustments still support the Entity's objectives and compliance to the national data standards.

A

a. Revalidate strategic alignment of the roadmap:

- Confirm ongoing relevance of initiatives to cloud-first, business continuity, and data governance priorities
- Adjust or reprioritize initiatives if national guidance or internal policies change

b. Reassess compliance with backup, recovery, and access control mandates:

- Confirm that roadmap changes do not introduce compliance gaps
- Engage legal, information security, and internal audit teams if material changes arise

c. Maintain documentation traceability across versions:

- Apply version control to roadmap and initiative card documents
- Archive superseded versions for audit and historical reference

iii. Submit finalized roadmap to governance stakeholders: Report roadmap status and evolution to both internal executives and external oversight.**a. Present the updated roadmap to the Data Management and Security Director for endorsement:**

- Include a summary of progress, changes, and resource impacts
- Ensure leadership has visibility into challenges, next steps, and delivery outlook

b. Share the approved roadmap with the National Data Governance Office (DGO):

- Submit via the preferred channel for roadmap reporting
- Address any DGO recommendations or feedback in future roadmap iterations

c. Use the roadmap submission to reinforce accountability and alignment:

- Signal transparency and commitment to the national data program
- Demonstrate the Entity's maturity in delivering structured, regulations-aligned storage and operations capabilities

RACI Matrix – Data Storage & Operations Roadmap Activation

Indicative

Sr No.	Activity	DGO	Data Management & Security Director	Data Architecture and Analytics Section Head	Data Custodian	Initiative Owner	IT Team	Data Architect
4.4.A	Initiate Execution and Governance Routines		A	R	C	C	C	C
4.4.B	Monitor Initiative Execution and Track Key Milestones		A	R	C	C		C
4.4.C	Review, Update, and Submit Roadmap for Governance Oversight	I	A	R		C		C

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Monitor Data Storage & Operations Initiatives



Review, Update, and Submit Roadmap

5. Data Storage and Operations Performance and Change Management



5. Data Storage and Operations Performance and Change Management


What is Data Storage and Operations Performance and Change Management and why does the Entity need it?

The Data Storage & Operations Performance & Change Management process provides a structured framework for continuously monitoring, maintaining, and improving the performance, security, and stability of the Entity's database environments. It ensures that storage systems and operational procedures remain aligned with service level expectations, regulatory obligations, and evolving technology standards.

This process enables the Entity to track and respond to performance issues in real time, maintain configuration integrity, enforce timely maintenance practices, and ensure that system changes are governed through structured documentation and reviews. By embedding automation, monitoring, and auditability into daily operations, the Entity reduces risks, improves reliability, and ensures ongoing alignment with national data and cloud policies.

Through activities such as SLA tracking, performance monitoring, configuration management, and patch governance, the Entity can detect issues early, maintain system health, and continuously align its infrastructure with strategic goals. These practices strengthen operational maturity, support audit readiness, and help ensure the long-term resilience of the Entity's data infrastructure.

The following activities guide the Entity through executing the Data Storage & Operations Performance & Change Management process.



Guiding Principles for Data Storage and Operations Performance and Change Management

1. **Monitor Continuously:** Establish real-time visibility into system health, usage, and SLA compliance to detect and resolve issues before they impact operations.
2. **Audit What You Configure:** Ensure all configuration settings are documented, versioned, and periodically reviewed to support compliance, troubleshooting, and change traceability.
3. **Patch with Purpose:** Apply maintenance updates and patches systematically, with proper testing, change control, and rollback plans to minimize risk.
4. **Track What You Own:** Maintain a centralized software asset inventory with complete licensing, support status, and update history to guide decision-making and audits.

5.1 Data Storage SLAs

Data Operations Service Level Agreements (SLAs) define measurable expectations for how database and storage systems should perform in terms of availability, responsiveness, backup, and recovery. This process ensures that operational expectations are formalized between IT service providers and data owners, enabling clarity, accountability, and continuous performance monitoring across the Entity's data infrastructure.

Process – Data Storage SLAs




Figure 9: Data Storage SLAs Process Map

Legend: Process Step Process Start Process End Document Off-Page Reference

Define SLA Metrics for Data Operations: Establish a clear and measurable set of SLA metrics that define performance expectations across key areas of data operations, including system availability, query execution, backup frequency, and recovery readiness. The Data Management and Governance Section Head and Data Governance Specialist lead this step in collaboration with the Data Custodian, IT Team, Data Owners, and the Information Security Team.

- i. **Identify systems and services that require SLA coverage:** Determine which systems require formal operational guarantees.
 - a. **Prioritize critical business systems and platforms for SLA definition:**
 - Examples: ERP, citizen services, analytics platforms, finance databases
 - Use the system priority levels defined in Section 4.3 to guide SLA tiering
 - b. **Engage system owners and business stakeholders to validate coverage:**
 - Validate that all systems with performance, availability, or recovery expectations are included
 - Exclude non-production and sandbox environments unless operationally required
 - c. **Classify systems by SLA tier to reflect service criticality:**
 - **Tier 1:** 24/7, mission-critical
 - **Tier 2:** Business hours or regional services
 - **Tier 3:** Low-usage or internal systems
- ii. **Populate the SLA Template for each in-scope system:** Use a consistent, structured format to define operational expectations.
 - a. **Fill out the SLA Template fields per system or platform:**

Asset Guidance

The **Data Storage and Operations SLA Template** can be leveraged by the Entity to capture SLA expectations.

Illustrative

SLA Template							
The SLA template can be leverage by the Entity to capture SLA expectations for Data Storage and Operations							
Field	Description					Example	
System Name	The name or identifier of the database or data platform under SLA					Citizen Services Database (CSD)	
SLA Tier	Categorization of criticality for the system (drives performance targets)					Tier 1 – Mission-Critical	
Data Owner	Person or unit responsible for the business operations of the system					Head of e-Services	
Technical Owner	Person or team responsible for system performance and compliance					Data Custodian	
Availability Window	Timeframe during which the system must be operational					24/7 or Sun-Thu, 8am–8pm	
Availability Target (%)	Uptime goal during the defined window					99.90%	
Max Execution Time (Critical TX)	Maximum allowable time for key transactions or queries					≤ 2 seconds for user authentication	
Backup Frequency	How often full or incremental backups are taken					Full daily + incremental every 6 hours	
Recovery Time Objective (RTO)	Max acceptable downtime before full recovery					2 hours	
Recovery Point Objective (RPO)	Max acceptable data loss from time of failure					15 minutes	
Monitoring Approach	How SLA metrics are tracked and reviewed					Automated – Integrated with Splunk SIEM	
Breach Notification Protocol	Process followed if SLA is breached, including who is notified and how					Email to System Owner + Escalation to IT Director	
Escalation Path	Defined path if SLA is not restored within threshold					Tiered escalation: DBA → Security → CIO	
Review Frequency	How often the SLA terms and performance are reviewed					Quarterly	
Notes / Exceptions	Any special conditions, exclusions, or temporary overrides					Excludes planned downtime (3rd Saturday, 2–4am)	

System Name	SLA Tier	Data Owner	Technical Owner	Availability Window	Availability Target (%)	Max Execution Time (Critical TX)	Backup Frequency

Recovery Time Objective (RTO)	Recovery Point Objective (RPO)	Monitoring Approach	Breach Notification Protocol	Escalation Path	Review Frequency	Notes / Exceptions

5.1.A
Cont.
(2/2)

b. Collaborate with technical and business stakeholders for each entry:

- Data Owners provide operational needs and performance expectations
- Data Custodians and the IT Team define technical feasibility and backup capabilities

c. Ensure all fields are completed consistently:

- Use standard units (e.g., % uptime, seconds, hours)
- Flag exceptions or special cases in the “Notes” section of the template

iii. Prepare templates for review and approval: Set the stage for formal SLA validation and operational implementation.

a. Store completed templates in a centralized, version-controlled location:

- Ensure accessibility to governance teams, auditors, and initiative owners

b. Use these templates as inputs for implementation (Step 5.1.B):

- Reference completed templates in formal SLA review sessions
- Ensure they serve as the source of truth for monitoring and reporting

5.1.B
(1/2)

Formalize and Implement SLA Agreements: Review, validate, and approve the completed SLA Templates for each system, ensuring formal agreement between data owners and IT service teams. This step operationalizes SLA commitments by clarifying responsibilities, embedding enforcement mechanisms, and ensuring compliance with national performance and availability standards. The Data Management and Governance Section Head and Data Governance Specialist lead this activity in collaboration with Data Custodians, Data Owners, IT Team, and the Information Security and Risk Management Section Head.

A

i. Facilitate stakeholder review and sign-off of SLA templates: Ensure agreement between business and technical owners on service expectations.

a. Review completed SLA Templates in a joint business–IT session:

- Validate availability windows, execution thresholds, backup frequencies, and recovery objectives
- Discuss implications of breach scenarios and escalation responsibilities

b. Assign formal sign-off from both the Data Owner and Technical Owner:

- Signatures may be digital (via workflow tool) or documented through formal email approval
- Record approvers’ names, roles, and approval dates in the template

c. Store approved templates in a central SLA repository:

- Use a version-controlled, secure repository accessible to Governance, IT, and Audit teams
- Reference these templates in policy audits or performance reviews

ii. **Integrate SLA parameters into operational and technical systems:** Ensure agreed metrics are actively tracked and enforced.

a. **Configure monitoring tools to track SLA metrics:**

- Examples: use Splunk, Prometheus, SolarWinds, or native DBMS tools to monitor uptime, execution time, and backup success
- Align log retention and metric granularity with audit and RPO/RTO needs

b. **Implement alerting and breach detection mechanisms:**

- Set threshold-based triggers for SLA violations
- Route alerts automatically to Data Owner, Data Custodian, and IT leadership per escalation path

c. **Coordinate with IT Ops and Security to embed SLA into daily operations:**

- Ensure backup jobs, access monitoring, and performance logs align with the SLA requirements
- Confirm alignment with cloud provider SLAs for hybrid systems

iii. **Prepare for governance and performance tracking:** Ensure SLA performance can be measured, reported, and improved over time.

a. **Define responsibility for SLA performance monitoring and compliance tracking:**

- Assign ownership to Data Custodians or IT Team
- Clarify roles for regular SLA reporting and breach remediation

b. **Establish a regular SLA review and refresh cycle (illustrative):**

- **Tier 1 systems:** quarterly
- **Tier 2–3 systems:** bi-annually or annually
- Include reviews in governance board or audit committee agendas

c. **Link SLAs to performance dashboards and operational KPIs:**

- Track metrics like % of SLA breaches, average resolution time, and compliance by system
- Use reports to support operational maturity assessments and risk reviews

Monitor SLA Performance and Trigger Responses: Implement automated monitoring of SLA metrics and ensure that breach alerts and escalation protocols are clearly defined and actively used. This step supports transparency, accountability, and service reliability by ensuring that all deviations from agreed service levels are detected, documented, and responded to in a timely manner. The IT Team lead this activity, in coordination with Information Security, Data Owners, and the Data Management and Governance Section Head.

i. **Set up automated SLA monitoring systems:** Ensure SLA metrics are tracked continuously and without manual intervention.

a. **Integrate database and infrastructure systems with monitoring platforms:**

- Use tools like Prometheus, Splunk, SolarWinds, or native DBMS monitors to track uptime, performance, and backup success
- Link SLA metrics (e.g., uptime %, max query time, backup status) to dashboards

b. **Configure system thresholds and compliance indicators:**

- Set performance thresholds for each SLA metric defined in the template
- Enable SLA dashboards to visually highlight systems at risk of breaching SLAs

c. **Enable automated breach detection:**

- System flags SLA violations based on defined thresholds
- Store event logs for audit and trend analysis

ii. **Trigger notifications and initiate escalation workflows:** Ensure all SLA breaches are immediately known and actioned.

a. **Configure automatic email or alert triggers:**

- Notify Data Custodians, system owners, and data governance leads when thresholds are breached
- Route high-severity issues to service desk or incident response teams

b. **Define and document escalation paths for unresolved breaches:**

- **Escalation tiers (illustrative):** Data Custodian → IT Operations Manager → Information Security → Executive Stakeholder
- Include response time expectations and responsibilities at each level

c. **Capture SLA breach data in centralized tracking system:**

- Use ITSM platforms (e.g., ServiceNow, Jira) to log SLA breaches, responses, and resolution timelines
- Tag incidents by system, root cause, resolution, and impact level

5.1.C
Cont.
(2/2)

iii. **Use SLA monitoring data for governance and service improvement:** Ensure SLA compliance is visible and informs future decisions.

a. **Generate regular SLA compliance reports:**

- Include % compliance by system, number of breaches, and mean resolution time
- Share reports with governance boards, risk committees, or performance councils

b. **Identify trends and recurring issues:**

- Analyze patterns of SLA failure across systems, time windows, or teams
- Recommend systemic improvements (e.g., infrastructure scaling, workload balancing, access optimization)

RACI Matrix – Data Storage SLAs

Indicative

Sr No.	Activity	Data Management & Security Director	Data Management and Governance Section Head	Information Security and Risk Management Section Head	Data Governance Specialist	Data Custodian	Information Security Team	IT Team	Data Owner
5.1.A	Define SLA Metrics for Data Operations	A	R		R	C	C	C	
5.1.B	Formalize and Implement SLA Agreements:	A	R	C	R	C	C	C	C
5.1.C	Monitor SLA Performance and Trigger Responses	I	A					R	C

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Define SLA Metrics for Data Operations



Monitor SLA Performance

5.2 Database Performance Monitoring

Database Performance Monitoring ensures that the Entity continuously observes and evaluates the health, responsiveness, and efficiency of its database environments. It enables the detection of performance bottlenecks, slow-running queries, system saturation, and degraded user experiences before they escalate into service disruptions or SLA violations.

Process – Database Performance Monitoring




Figure 10: Database Performance Monitoring Process Map

Legend: Process Step Decision Process Start Process End Document Off-Page Reference

Define Key Database Performance Metrics: Establish a set of metrics that provide visibility into the operational health, efficiency, and responsiveness of the Entity's database systems. These metrics serve as the foundation for proactive monitoring, SLA compliance tracking, and long-term capacity planning. The Data Custodian leads this activity in collaboration with the IT Team, Data Architects, and Information Security.

- i. **Identify core performance dimensions to be monitored:** Establish broad categories of performance relevant to daily operations and stability.
 - a. **Define capacity-related metrics:**
 - Storage used (GB/TB), storage growth rate, number of database objects, number of storage containers
 - Monitor volume of committed vs. uncommitted blocks or pages
 - b. **Define performance-related metrics:**
 - Query execution time, average transaction duration, number of active sessions
 - IOPS (input/output operations per second), throughput, cache hit ratios
 - c. **Define availability and uptime metrics:**
 - Uptime %, failover events, connection latency, system errors or lockups
 - Percentage of time system is operational within defined SLA windows
- ii. **Establish metric definitions and measurement methods:** Ensure consistency and accuracy across all performance metrics.
 - a. **Define each Metric:**
 - Define metric name, unit of measure, calculation method, frequency of collection, and responsible owner
 - Example: "Query Execution Time – measured in seconds, calculated as average time to complete SELECT operations during business hours"
 - b. **Link metrics to system or environment scope:**
 - Specify whether metrics are collected for production, staging, dev/test, or DR environments
 - Tailor collection depth by system criticality or SLA tier
 - c. **Align metrics with compliance, SLA, and capacity planning needs:**
 - Ensure that selected metrics support reporting for 5.1 (SLAs), 5.4 (Configuration Management), and roadmap decisions

5.2.B
(1/2)

A

Configure Monitoring Tools and Dashboards: Set up technical tools and visualization interfaces to collect, monitor, and display real-time performance metrics across all database environments. This step enables proactive detection of issues, supports SLA enforcement, and provides visibility to technical teams and governance stakeholders. The Data Custodian and IT Team lead this activity, in collaboration with the Data Visualization Officer, Business Data Analyst, and the Data Management and Governance Section Head.

- i. **Select and configure monitoring tools across database platforms:** Ensure performance monitoring is active, consistent, and integrated across environments
 - a. **Choose appropriate monitoring tools based on database technology:**
 - Examples: SQL Server Management Studio (SSMS), Oracle Enterprise Manager, MySQL Workbench, Prometheus, SolarWinds, Zabbix
 - For hybrid or cloud systems, consider native monitoring from AWS CloudWatch, Azure Monitor, or GCP Operations
 - b. **Configure agents or collectors to track the metrics defined in 5.2.A:**
 - Metrics may include CPU usage, IOPS, query time, connection failures, or uptime
 - Ensure alignment with SLA targets and compliance policies
 - c. **Set metric thresholds and data collection frequency:**
 - Example: capture query execution times every 15 seconds, log uptime every 5 minutes
 - Define thresholds for warning and critical alerts per metric category
- ii. **Build centralized dashboards to visualize database performance:** Provide visibility and actionable insights for operations and governance.
 - a. **Develop real-time dashboards for technical teams:**
 - Show active connections, failed queries, performance bottlenecks, and backup job status
 - Enable filtering by environment, system, or SLA tier
 - b. **Create executive-level views for governance and reporting:**
 - Visualize compliance status (% of systems meeting SLA), trend graphs for query time or uptime, and monthly summaries
- iii. **Integrate alerts and logs with service management workflows:** Ensure the system can trigger response actions and support audit traceability.
 - a. **Configure alert thresholds to trigger notifications and escalations:**
 - Route alerts based on severity (e.g., CPU > 90%, query time > 3 seconds)
 - Notify Data Custodians, service desk, or escalation teams based on defined roles

5.2.B
Cont.
(2/2)

b. Forward alerts and logs to centralized SIEM or ITSM platforms:


- Examples: Splunk, ServiceNow, Jira, Microsoft Sentinel
- Support integration with security and compliance monitoring

c. Maintain logs and alert history for governance use:

- Store records for internal audits, SLA tracking, and system tuning reviews (5.2.C)
- Ensure retention policies comply with national regulations

Asset Guidance

A Data Storage and Operations Dashboards provides real-time insights into the performance, availability, and efficiency of storage systems within an entity. It displays key metrics such as storage utilization, system uptime, query response times, and backup success rates, enabling stakeholders to monitor operational performance briefly. *Illustrative*

5.2.C
(1/2)

Analyze Trends and Detect Performance Issues: Leverage performance monitoring data to identify negative trends, diagnose root causes, and trigger actions that maintain or improve database efficiency and reliability. This step supports continuous improvement by helping teams spot issues early, resolve bottlenecks, and optimize resource allocation. The Data Custodian leads this activity in collaboration with the IT Team, and the Data Management and Governance Section Head.

i. Establish routines for performance data review and analysis: Ensure performance monitoring is actionable, not just reactive.

a. Schedule regular analysis of monitoring dashboards and logs:

- Daily reviews for Tier 1 systems, weekly or monthly for lower-tier systems
- Focus on uptime, query speed, system load, and backup success

b. Compare current performance against SLA thresholds and historical baselines:

- Identify deviations from expected norms and SLA metrics

- Track performance over time to detect slow degradation
- a. **Create summary reports and trend snapshots:**
 - Monthly or quarterly insights that highlight patterns and anomalies
 - Include visual charts to show changes in performance over time
- ii. **Detect early warning signs of potential performance issues:** Use data patterns and alerts to uncover emerging risks.
 - a. **Monitor for recurring anomalies or threshold breaches:**
 - Examples: repeated spikes in query latency, backup job failures, I/O saturation during peak hours
 - Flag systems with more than two consecutive threshold breaches for deeper investigation
 - b. **Identify underperforming queries or inefficient data access patterns:**
 - Use query profiling tools or explain plans to identify slow or unoptimized SQL
 - Correlate performance issues with specific user actions or application events
 - c. **Link performance issues to infrastructure or configuration root causes:**
 - Diagnose storage limitations, CPU or memory bottlenecks, and misconfigured indexes
 - Document findings and propose corrective actions
- iii. **Recommend improvements and initiate tuning or escalation:** Act on findings to optimize performance and maintain system health.
 - a. **Propose tuning actions for underperforming systems:**
 - Examples: reindex tables, refactor queries, upgrade storage tier, scale compute resources
 - Prioritize based on business impact and recurrence
 - b. **Escalate systemic or high-risk issues through governance or IT leadership:**
 - Route unresolved or recurring issues to the data architecture and analytics section head, or security teams as needed
 - Trigger change requests if tuning requires downtime or configuration changes

RACI Matrix – Database Performance Monitoring

Indicative

Sr No.	Activity	Data Management & Security Director	Data Management and Governance Section Head	Data Architecture and Analytics Section Head	Data Architect	Data Custodian	Information Security Team	IT Team	Data Visualization Officer	Business Data Analyst
5.2.A	Define Key Database Performance Metrics	I	I	A	C	R	C			
5.2.B	Configure Monitoring Tools and Dashboards	I	C	A		R		R	C	C
5.2.C	Analyze Trends and Detect Performance Issues		C	A		R		C		

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Define Database Performance Metrics



Analyze Trends and Detect Issues

5.3 Database Configuration Management

Database Configuration Management ensures that all configuration parameters across development, testing, and production environments are documented, standardized, and maintained in a controlled and auditable manner. This process reduces inconsistencies, improves performance stability, and ensures compliance with internal policies and national standards.

5.3.A
(1/2)

A

Establish Standardized Configuration Parameters: Define and document a core set of configuration parameters that must be consistently applied across database environments to support performance, security, and compliance. The Data Custodian leads this activity in collaboration with the IT Team, Data Architects, Information Security, and the Data Architecture and Analytics Section Head.

- i. **Identify key configuration parameters for standardization:** Focus on settings that affect performance, scalability, security, or compliance.
 - a. **Select commonly used parameters across the Entity's database platforms:**
 - Examples: Max Connections, Query Timeout, Memory Allocation, Auto-Commit, Logging Settings, Backup Compression
 - b. **Include parameters that vary across environments or carry operational risk:**
 - Identify parameters that often differ between Dev/Test/Prod and may lead to inconsistencies
 - Prioritize those tied to stability, access control, or resource limits
 - c. **Validate parameter relevance with application teams:**
 - Ensure chosen parameters reflect real-world configurations and are actively managed
 - Check against past incidents or performance issues linked to misconfiguration
- ii. **Populate the Database Configuration Template for each parameter:** Use the standardized template to ensure structured documentation and traceability.
 - a. **Fill out each field in the template:**

Asset Guidance

The **Database Configuration Template** can be leveraged by the Entity to document, manage, and standardize database configurations across environments. **Illustrative**

Database Configuration Template					
The primary purpose of this template is to provide a structured and centralized way to document, manage, and standardize database configurations across environments.					
Field	Description			Example	
Configuration Parameter	Name of the specific database configuration setting being standardized.			Max Connections	
Default Value	The standard value for the configuration parameter applied across environments.			200	
Description	A brief explanation of the parameter's purpose or role in database operations.			Maximum number of simultaneous database connections.	
Environment	Specifies which environments (e.g., Development, Testing, Production) the configuration applies to.			Dev, Test, Prod	
Compliance/Policy Reference	Links the parameter to a regulatory requirement, organizational policy, or best practice guideline.			CIS Database Benchmark	
Change History	Tracks updates to the parameter, including the reason for changes and the approver's details.			Updated from 150 to 200 on 31-Dec-2024 by John Doe.	
Configuration Parameter	Default Value	Description	Environment	Compliance/Policy Reference	Change History

5.3.A
Cont.
(2/2)

- b. Use one row per configuration entry and one template per environment grouping if needed:

- Separate templates can be maintained for each DB engine (e.g., Oracle, MySQL, PostgreSQL) if configurations vary significantly

- c. Maintain a centralized, version-controlled repository for completed templates:

- Ensure accessibility to Data Custodians, Audit, and governance teams
- Align naming conventions and versioning with IT change management processes

- iii. Confirm baseline standards for rollout and policy alignment: Set expectations for adoption and enforcement across systems.

- a. Review and approve configuration parameters through a technical review board:

- Include representatives from infrastructure, security, and data architecture teams
- Validate proposed defaults against performance benchmarks and vendor guidance

- b. Confirm which parameters are required, recommended, or optional:

- Categorize each setting for clarity and prioritization in enforcement

5.3.B
(1/2)

Maintain Configuration Records and Audit Trail: Implement a structured process to track, store, and maintain database configuration records over time. This step ensures all changes are transparent, traceable, and aligned with policy, and that the Entity has a reliable audit trail to support governance, troubleshooting, and compliance. The Data Custodian leads this activity in collaboration with the IT Team, Audit Team, and the Information Security and Risk Management Section Head.

- i. Centralize storage of completed configuration templates: Ensure documentation is secure, version-controlled, and easily accessible.

- a. Store completed configuration templates in a centralized repository:

- Use platforms such as SharePoint, Git, internal document control systems, or a CMDB
- Organize by database type, system, or environment for ease of access

- b. Apply file naming conventions and versioning rules:

- Example: DB_Config_MySQL_Prod_v1.2_2024-08-25
- Include version number, environment, DBMS type, and date

- c. Assign access permissions based on roles and responsibilities:

- Ensure read/write access is restricted to Data Custodians and authorized reviewers
- Allow read-only access for audit, compliance, and governance teams

A

5.3.B
Cont.
(2/2)

- ii. **Maintain change history and documentation trail:** Track what changes were made, why, by whom, and when.
 - a. **Populate the “Change History” field in the configuration template for each update:**
 - Record the previous value, updated value, date of change, person responsible, and reason for the change
 - b. **Link configuration changes to supporting tickets or approvals:**
 - Reference ITSM or change management IDs for traceability (e.g., ServiceNow, Jira)
 - Store related meeting notes or review board decisions where applicable
 - c. **Archive superseded versions of configuration records:**
 - Maintain a chronological history of configuration changes
 - Ensure historical records are retained in accordance with data retention and audit policies
- iii. **Ensure alignment with audit, compliance, and operational reviews:** Use the configuration records as an auditable source of truth.
 - a. **Provide configuration templates during internal or external audits:**
 - Demonstrate adherence to configuration baselines and change management processes
 - Address audit queries with complete and verified documentation
 - b. **Schedule periodic reviews of active configuration templates:**
 - Quarterly or semi-annual reviews for production systems
 - Verify actual configurations match documented baselines in the template
 - c. **Use templates as input to performance optimization, compliance checks, and future tuning:**
 - Cross-reference with performance metrics (Section 5.2)
 - Ensure consistency with SLAs (Section 5.1) and system requirements

5.3.C
(1/2)

A

Review and Align Configurations with Policy and Environment Requirements: Conduct periodic reviews of database configuration parameters to ensure they remain aligned with security policies, performance expectations, and environment-specific operational needs. This step focuses on evaluating the ongoing appropriateness of configuration baselines, identifying misalignments, and ensuring that updates reflect changing technical, regulatory, or business contexts. The Data Custodian leads this activity in collaboration with the Information Security team, Data Architects, and Audit team.

- i. **Schedule and conduct periodic configuration reviews:** Review existing configurations at regular intervals to assess relevance and applicability.
 - a. **Define review frequency based on system criticality and risk profile (Illustrative):**
 - Tier 1 / Production: Quarterly
 - Tier 2 / Staging & DR: Bi-annually
 - Tier 3 / Dev/Test: Annually
 - b. **Reassess active settings in live systems against documented templates:**
 - Use automation or scripts to extract current values and compare them to the approved configuration records
 - Identify drift or undocumented adjustments
 - c. **Flag parameters that are outdated, deprecated, or inconsistent:**
 - Highlight misaligned or unsupported values, including hardcoded legacy defaults or post-incident changes
- ii. **Validate alignment with current policies and performance expectations:** Ensure configurations reflect internal standards, external regulations, and operational context.
 - a. **Reevaluate configurations against compliance benchmarks:**
 - Confirm settings meet CIS benchmarks, national security policies (e.g., NIA, NCSA), or vendor guidance
 - Include encryption, logging, authentication, and memory allocation parameters
 - b. **Assess whether current settings meet operational needs and SLA commitments:**
 - Align configurations with SLA metrics (e.g., response time, uptime, IOPS)
 - Identify settings that are limiting performance or causing instability
- iii. **Update configuration baselines and communicate changes:** Apply approved changes and ensure documentation remains current.
 - a. **Submit change recommendations through formal review channels:**
 - Use established governance processes for approving parameter changes
 - Provide justifications, risk assessments, and rollback procedures

5.3.C
Cont.
(2/2)**b. Update configuration templates and change history once approved:**

- Ensure version number, change date, and approver are recorded
- Archive previous versions for traceability

c. Notify stakeholders and update operational systems:

- Inform application teams, DevOps, and incident response teams of new baseline settings
- Ensure changes are applied consistently across all environments

RACI Matrix – Database Configuration Management*Indicative*

Sr No.	Activity	Information Security and Risk Management Section Head	Data Architecture and Analytics Section Head	Data Architect	Data Custodian	IT Team	Audit Team	Information Security Team
5.3.A	Establish Standardized Configuration Parameters		A/C	C	R	C		
5.3.B	Maintain Configuration Records and Audit Trail	A/C			R	C	C	
5.3.C	Review and Align Configurations with Policy and Environment Requirements	A/C		C	R		C	C

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Establish Configuration Parameters

5.4 Database Maintenance and Updates

Database Maintenance and Updates ensure that the Entity's database environments remain secure, performant, and compliant with vendor recommendations and regulatory obligations. This process involves applying vendor patches, performing regular updates, and tracking licensing and asset costs—all while minimizing disruption to business services.

5.4.A
(1/2)

Apply Vendor Maintenance Patches Across Environments: Ensure that vendor-issued patches and minor updates are applied regularly and consistently across development, test, and production environments. This step reduces vulnerability to security threats, maintains platform stability, and ensures ongoing vendor support. The Data Custodian leads this activity in collaboration with the IT Team, and Information Security Team.

i. **Maintain visibility into vendor patch release cycles:** Monitor and plan around the patch schedules of supported database platforms.

a. **Subscribe to vendor release notifications and advisories:**

- Use Oracle, Microsoft, PostgreSQL, or cloud provider channels to stay informed
- Track CVEs (Common Vulnerabilities and Exposures) and patch dependencies

b. **Maintain a rolling patch calendar:**

- Record expected release dates and product-specific timelines (e.g., quarterly patch cycles)
- Share with IT Operations, Security, and Governance teams

c. **Link patch notifications to your Software Asset Inventory:**

- Identify affected systems based on DBMS version and environment
- Use the "Next Scheduled Patch" and "Last Patch Applied" fields to track patch status

ii. **Apply patches across dev, test, and production environments:** Execute patching in a controlled, staged manner.

a. **Define a standardized patch deployment sequence:**

- Start with Development → Test → UAT → Production
- Allow adequate validation time between stages

b. **Validate patch behavior and system compatibility before production rollout:**

- Test for functional integrity, performance impacts, and security behavior
- Involve application owners if schema or driver changes are introduced

c. **Schedule production patching within approved maintenance windows:**

- Minimize disruption by aligning with planned downtime periods
- Communicate planned changes to system users and business owners in advance

A

5.4.A
Cont.
(2/2)

iii. **Document patch application and retain change history:** Ensure traceability, accountability, and audit readiness.

a. **Log all patching activity per environment:**

- Capture date, version applied, systems affected, and patch reference ID
- Use a central patch log or link to ITSM ticketing system (e.g., ServiceNow)

b. **Record updates in the Software Asset Inventory Template:**

- Update “Last Patch Applied” and “Next Scheduled Patch” fields
- Link to change approval records or governance board decisions if applicable

c. **Maintain patch documentation for internal audits and compliance reviews:**

- Include rollback plans, test results, and risk mitigation notes
- Store in a version-controlled repository accessible to security and audit teams

Asset Guidance

The **Software Asset Inventory Template** can be leveraged by the Entities for tracking database software licenses, versions, and support contracts.

Illustrative

Software Asset Inventory							
The Software Asset Inventory can be leveraged by the Entity for the tracking of database software licenses, versions, and support contracts							
Field	Description				Example		
Database System Name	Name or identifier of the database or DBMS in use				Core Financial Database		
DBMS Platform	The type of database software or vendor technology used				Oracle 19c Enterprise Edition		
Version Number	Installed version of the DBMS software				19.13.0.0.0		
Environment	The deployment context of the database instance				Production / Dev / Test / DR		
License Type	Type of software license (e.g., commercial, open-source, subscription)				Commercial – Per Core License		
License Status	Whether the license is active, expired, or under evaluation				Active		
License Owner / Responsible Team	Entity or team responsible for maintaining the license				IT Operations – DBA Team		
License Expiry Date	Date on which the license will expire or require renewal				31-Dec-25		
Annual Support Cost (QAR)	Yearly support fee charged by the vendor (if applicable)				QAR 120,000		
Vendor Support Status	Whether the current version is under active vendor support				Yes – Supported until Dec 2026		
Server / Hosting Platform	Hosting location or platform for the database				On-Prem – Virtual Server OR Azure SQL DB		
Server Lease Cost (if applicable)	Annual cost for server hosting (physical or cloud)				QAR 25,000 (Cloud Instance – Azure)		
Last Patch Applied	Date of the most recent patch or maintenance update applied				1-Feb-24		
Next Scheduled Patch	Date for the next scheduled patch or update				15-Jul-24		
Notes / Exceptions	Any caveats, known risks, or change control notes				Upgrade planned Q1 2025; legacy code dependency		

Database System Name	DBMS Platform	Version Number	Environment	License Type	License Status	License Owner / Responsible Team	License Expiry Date

Annual Support Cost (QAR)	Vendor Support Status	Server / Hosting Platform	Server Lease Cost (if applicable)	Last Patch Applied	Next Scheduled Patch	Notes / Exceptions

Maintain Version Control and Software Asset Inventory: Track database versions, licenses, and associated costs in a centralized, structured, and auditable format. This step ensures the Entity complies with licensing agreements, maintains vendor support eligibility, and enables informed decision-making for renewals, upgrades, and budgeting. The Data Custodian leads this activity in collaboration with IT Team, Finance Team, and Audit team.

- i. **Populate and centralize the Software Asset Inventory:** Ensure comprehensive visibility into software assets across all environments.
 - a. **Use the Software Asset Inventory Template to document asset details:**
 - Fields include system name, DBMS platform, version, license type, vendor support status, and cost-related information
 - Include all environments (Prod, Test, Dev, DR) and associated hosting platforms
 - b. **Store the completed inventory in a centralized, secure repository:**
 - Examples: SharePoint, enterprise CMDB, internal asset management system
 - Apply version control to track changes over time
 - c. **Assign data ownership and access permissions:**
 - Restrict edit access to Data Custodians and IT asset managers
 - Enable read-only access for Finance, Procurement, and Governance stakeholders
- ii. **Track licensing agreements, renewal timelines, and support status:** Ensure ongoing compliance and readiness for audits or contract changes.
 - a. **Record license types and associated terms per system:**
 - Examples: per-core, enterprise edition, open-source with support, SaaS license
 - Track key terms such as number of users, CPU limits, and renewal frequency
 - b. **Maintain visibility into license renewal and support expiry dates:**
 - Populate and monitor the “License Expiry Date” and “Vendor Support Status” fields
 - Set alerts 60–90 days in advance of renewal deadlines
 - c. **Collaborate with Finance and Procurement for license budgeting:**
 - Use annual support costs and lease costs to forecast IT spending
 - Ensure compliance with procurement policies for renewals or expansions
- iii. **Align version control with security, performance, and vendor guidance:** Ensure all version-related decisions are documented and justified.

5.4.B
Cont.
(2/2)**a. Log all version upgrades or downgrades in the inventory:**

- Record date, old version, new version, and reason for change
- Link to patch records, SLA changes, or compliance drivers

b. Track systems running unsupported or nearing end-of-life versions:

- Flag such systems for remediation planning in Section 5.4.C
- Validate against vendor lifecycle roadmaps

c. Use the asset inventory to inform roadmap, risk management, and audit reviews:

- Ensure updates feed into architecture planning, governance reporting, and security assessments
- Maintain evidence for internal and external audits of software compliance

5.4.C
(1/2)

Review and Justify Deferment of Non-Applied Updates: Establish a structured process to document, evaluate, and formally justify cases where database patches, updates, or upgrades are intentionally deferred. This step ensures transparency, risk awareness, and alignment with governance expectations when the Entity chooses not to implement vendor-recommended changes. The Data Custodian leads this activity in collaboration with the Information Security team, System Owners, Data Architect, and the Data Architecture and Analytics Section Head.

i. Identify and document deferred updates and patches: Create visibility into systems where updates were not applied.**a. Maintain a running list of deferred patches and version upgrades:**

- Use logs, release bulletins, or vendor advisory trackers to flag updates not applied
- Track affected systems, patch/version details, and release dates

b. Record each deferment in the Software Asset Inventory Template:

- Add a note to the “Notes / Exceptions” field explaining the pending or skipped update
- Include references to the original patch bulletin or vendor release notes

c. Highlight deferred updates that involve known vulnerabilities or critical functionality:

- Tag high-risk or security-impacting deferments for additional review

ii. Justify deferments through technical and business rationale: Ensure deferrals are not arbitrary and are risk-aware.**a. Provide a formal rationale for each deferred update:**

A

5.4.C
Cont.
(2/2)

- Examples: dependency on unsupported code, risk of breaking integrations, pending system replacement, compatibility concerns
- Document input from Data Custodians, application owners, and security officers

b. Conduct risk assessments where applicable:

- Assess impact on performance, availability, security, and compliance
- Determine whether mitigation controls (e.g., WAF rules, access limits) can reduce exposure

c. Submit justifications to the appropriate governance body for review:

- Examples: Data Architecture and Analytics Section Head
- Include a recommended review timeframe (e.g., revisit in 3–6 months)

iii. Review deferments periodically and update governance records: Ensure that deferred updates are not forgotten or left unaddressed.**a. Set a review cadence for all deferred updates (illustrative):**

- Critical systems: Monthly or Quarterly
- All others: Bi-annually or aligned with patch cycle

b. Update the Software Asset Inventory and patch schedule accordingly:

- If a deferred patch is later applied, update the “Last Patch Applied” field
- If deferment is extended or revoked, update the rationale and next review date

c. Report deferred update status in governance and audit reports:

- Include in quarterly IT risk summaries or annual compliance audits
- Use deferment trends to inform future investment or modernization efforts

RACI Matrix – Database Maintenance and Updates

Indicative

Sr No.	Activity	Information Security and Risk Management Section Head	Data Architecture and Analytics Section Head	Data Architect	Data Custodian	IT Team	Audit Team	Information Security Team	Finance Team	System Owners
5.4.A	Apply Vendor Maintenance Patches Across Environments	I	A		R	C		C		
5.4.B	Maintain Version Control and Software Asset Inventory		A		R	C	C		C	
5.4.C	Review and Justify Deferment of Non-Applied Updates		A/C	C	R			C		C

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can




Apply Maintenance Patches

5.5 QDTI

The Qatar Data Trust Index (QDTI) is designed to measure and enhance data integrity across national entities by promoting adherence to governance standards and ensuring the accuracy, reliability, and responsible management of data. This initiative aims to build trust and support a robust digital economy. The overall QDTI score is derived from two dimensions: the Data Quality Score (DQS) and the Data Compliance Score (DCS), which together combine to determine the Entity's QDTI score.

This section outlines the steps needed to gather the necessary information for calculating the DCS.



***Note:** The sum of the weights across all dimensions should add up to 100%

****Note:** This score is calculated by computing the weighted average score of all the sub-dimensions under each respective Dimension (scores are on a scale of 0 to 100)

Figure 11: Illustrative QDTI Scoring Methodology for an Entity

5.5.A

Initiate Data Storage and Operations Compliance Assessment: The Data Management and Governance Section Head assigns the Data Compliance Specialist to oversee the assessment process. The specialist is responsible for managing the evaluation of the Entity's adherence to the National Data Storage and Operations Standards.



Key Message / Information

The Data Storage and Operations Domain falls under the Data Enablement Dimension in the National Data Management Framework and, consequently, within the QDTI Dimension.


5.5.B

Conduct Data Storage and Operations Compliance Assessment: Initiate the self-assessment to evaluate adherence to the National Data Storage and Operations Standards. This assessment will ensure compliance and accountability while identifying areas for improvement.

Asset Guidance

The **Qatar Data Trust Index Methodology Guide** can be leveraged to guide the Entity through the steps required to conduct the self-assessment.

Illustrative



Asset Guidance

The **QDTI Assessment Template** can be used to document compliance with the National Data Standards. This template is intended for use prior to the DGO launching the self-assessment on its platform to calculate the overall QDTI Score.

Illustrative

Domain	Control	Control Description
DQM	DQM.1.1.1 Data Quality Policy	As part of the Data Quality Policy control, the Entity shall create and maintain an enterprise-wide Data Quality Management Policy...
DQM	DQM.1.1.2 Data Quality Policy	As part of the Data Quality Policy control, the Entity shall publish and distribute a Data Quality Management Policy...
DQM	DQM.1.1.3 Data Quality Planning	As part of the Data Quality Planning control, the Entity shall perform initial data quality profiling for the identified Critical Data Elements (CDEs)...
DQM	DQM.1.1.4 Data Quality Planning	As part of the Data Quality Planning control, the Entity shall develop a data quality plan aligned with its data strategy and profiling results...

Specification	Control Specification	Capability Level	Compliance
DQM.1.1.1.1 Data Quality Policy Creation	The Entity shall create a data quality policy in alignment with the NDQAF published by the NPC...	Foundational	Compliant
DQM.1.1.2.1 Data Quality Policy Publication	The Entity shall publish a data quality policy....	Foundational	Compliant
DQM.1.1.3.1 Initial Data Quality Profiling	The Entity shall identify CDEs and perform initial data quality profiling for the identified CDEs....	Foundational	Compliant
DQM.1.1.4.1 Data Quality Plan Design	Based on the initial data profiling results, the Entity shall develop a data quality plan in alignment with the Entity's data strategy...	Foundational	Compliant



Key Message / Information

The DGO will launch the self-assessment on its platform and will provide toolkits and host workshops to guide entities through the self-assessment process during the first year.

5.5.C

Submit Data Storage and Operations Score to DGO: Submit the Data Storage and Operations Compliance Score to the DGO for evaluation.



Key Message / Information

The DGO will calculate the Entity's overall Qatar Data Trust Index (QDTI) score only after the Entity has completed both the Data Quality Score (DQS) assessment and the required self-assessments for all other data management domains to determine the Data Compliance Score (DCS). Once all assessments are completed, the DGO will aggregate the compliance scores across domains to finalize the DCS, which will then be combined with the DQS to determine the Entity's overall QDTI score.

RACI Matrix - QDTI

Indicative

Sr No.	Activity	DGO	Data Management & Security Director	Data Management and Governance Section Head	Data Compliance Specialist	Data Governance Specialist
5.5.A	Initiate Data Storage and Operations Compliance Assessment		I	A	R	
5.5.B	Conduct Data Storage and Operations Assessment		I	A	R	C
5.5.C	Submit Data Storage and Operations Score to DGO	I	I	A	R	I

RACI Matrix Legend: **R:** Responsible **A:** Accountable **C:** Consulted **I:** Informed

Based on this Activity, the Entity can



Calculate Data Storage and Operations Compliance Score

6. Appendix



6.1. Templates and Frameworks

Asset No.	Asset Name	Description	File Name
1	Data Storage & Retention Policy Template	Template containing a sample table of contents the Entity can refer to when drafting its policy.	NPC_DGO_QDKC_DSO_Data Storage and Retention Policy Template
2	Data Management Policy Communication Email Template	Template containing a sample email the Entity can use to announce the publishing of the new policy.	NPC_DGO_QDKC_General_Data Management Policy Communication Email Template
3	Baseline Storage Assessment Report	Template is designed to document the current state of an entity's data storage infrastructure.	NPC_DGO_QDKC_DSO_Base line Storage Assessment Report
4	Storage and Usage Requirements Forecasting Template	Template is designed to provide a structured approach for estimating future storage capacity needs.	NPC_DGO_QDKC_DSO_Stor age and Usage Requirements Forecasting Template
5	Initiative Card Template	Template used to outline and track the details of a specific initiative.	NPC_DGO_QDKC_General_Ini tiative Card Template
6	Data Management Domain Plan Template	Template that helps the Entity capture an overview of identified initiatives, detail their sub-activities, owners, stakeholders, and activity timeframe.	NPC_DGO_QDKC_General_Data Management Domain Plan Template
7	Disaster Recovery Playbook Template	Template that helps the Entity create a business continuity plan and minimize downtime by providing a structured recovery framework.	NPC_DGO_QDKC_DSO_Disa ster Recovery Playbook Template
8	Data Storage and Operations SLA Template	Template can be leverage by the Entity to capture SLA expectations for Data Storage and Operations	NPC_DGO_QDKC_DSO_SLA Template
9	Database Configuration Template	Template is to provide a structured and centralized way to document, manage, and standardize database configurations across environments.	NPC_DGO_QDKC_DSO_Data base Configuration Templates

Asset No.	Asset Name	Description	File Name
10	Software Asset Inventory	Template that can be leveraged by Entities for tracking database software licenses, versions, and support contracts.	NPC_DGO_QDKC_DSO_Software Asset Inventory
11	Qatar Data Trust Index Methodology Guide	The Qatar Data Trust Index Methodology Guide can be leveraged to guide the Entity through the compliance self-assessment.	22012025_NPC_QDTI Methodology Guide_V1.0
12	QDTI Assessment Template	The QDTI Assessment Template can be used to document compliance with the National Data Standards.	23012025_NPC_QDTI Assessment Template_V1.0



المجلس الوطني للتخطيط
National Planning Council

دولة قطر • State of Qatar