DIGITAL TRANSFORMATION

Enterprise Data Backup, Disaster Recovery and Cyber Resilience Frameworks for Secure Digital Transformation in Saudi Arabia under Vision 2030

Nishad Mohammad Abdul

Published November 2026

Abstract

Secure digital transformation depends on an enterprise's ability to recover trustworthy information and critical services after cyberattack, operational failure or infrastructure disruption. In Saudi Arabia, integration among government platforms, private enterprises and cloud services makes conventional backup success rates an inadequate indicator of resilience. This structured narrative review synthesizes scholarly research and authoritative guidance published between 2020 and 2025 concerning backup isolation, recovery engineering, ransomware response, data integrity, service dependencies and national governance. The review distinguishes routine disaster recovery from adversarial cyber recovery, where an attacker may compromise identity, management interfaces and replicated datasets before disruption becomes visible. Findings are organized into six interconnected capabilities: business impact classification, protected recovery copies, trustworthy recovery control planes, clean room restoration, service level validation and auditable governance. A Saudi oriented Enterprise Cyber Recovery Assurance Framework is proposed to connect these capabilities with recovery time and recovery point objectives, operational evidence and applicable national controls. The analysis indicates that resilience is not obtained by purchasing redundant storage alone; it emerges from independently protected recovery paths and repeated demonstrations that essential services can resume with verified data. The paper identifies implementation challenges, limitations of existing evidence and research priorities for measuring recovery readiness across complex Saudi digital ecosystems.

1 Introduction

Saudi Arabia's Vision 2030 has accelerated the modernization of public services, enterprise platforms, financial operations and digitally connected infrastructure. Digital transformation, however, creates dependence on interconnected information assets. A single compromise involving identity services, administrative credentials, application databases or storage infrastructure can impair business processes simultaneously. The Vision 2030 Annual Report 2024 describes continued investment in digital capability, while national digital government guidance explicitly links service reliability with risk management and continuity planning (Saudi Vision 2030, 2025; Digital Government Authority, 2025). Accordingly, secure transformation requires recovery that is designed, governed and demonstrated before a damaging incident occurs.

Backup and disaster recovery are related but distinct capabilities. A backup creates recoverable data copies; disaster recovery coordinates the restoration of infrastructure, applications, connectivity and operations after disruption. Cyber resilience extends beyond both concepts by assuming intelligent adversaries can attack recovery systems themselves. Ransomware may encrypt production workloads, exfiltrate sensitive records, delete backup catalogs, corrupt replicas and disable administrative access. McIntosh et al. (2024) show why a narrow focus on file encryption no longer captures the

contemporary threat. NIST storage guidance likewise treats data protection, isolation and restoration assurance as connected security engineering concerns (NIST, 2020a).

Traditional recovery assumptions are fragile in estates combining older enterprise applications, virtualized infrastructure, cloud native services, third party platforms and outsourced support. Replication can duplicate corruption, while a successful storage restore can still leave an application unusable because identity providers, encryption keys or upstream services are missing. Reviews of multi cloud architecture identify portability and provider heterogeneity as persistent operational constraints (Tomarchio et al., 2020; Alonso et al., 2023). These concerns imply that recovering a service, rather than merely retrieving files, should become the unit of enterprise resilience measurement.

This review addresses a specific gap between security control compliance and practical recoverability. Frameworks prescribe safeguards, yet organizations still need ways to demonstrate that restoration procedures survive compromised credentials, uncertain data cleanliness and dependencies that cross organizational boundaries. The purpose here is neither to claim a tested Saudi deployment nor to calculate national incident rates. Instead, the study develops a literature grounded operating model through which Saudi organizations can prioritize recovery investments, assess control effectiveness and support continuity of critical digital services under Vision 2030.

2 Conceptual background and literature

Business continuity defines the organizational capability to maintain priority activities at acceptable levels during disruption. Disaster recovery provides the technology procedures needed to restore those activities. Information and communication technology readiness for business continuity emphasizes the alignment of technical recovery with organizational requirements (ISO/IEC, 2025). An enterprise may therefore meet a documented infrastructure recovery time while failing its actual business requirement if records are incomplete, a transaction queue is inconsistent or users cannot authenticate. Technical metrics require interpretation in relation to business outcomes.

Recovery time objective, or RTO, specifies a target duration for restoring an activity after interruption; recovery point objective, or RPO, expresses the maximum tolerable interval of data loss measured backward from disruption. These targets are planning requirements rather than guarantees. Service owners should establish them using business impact analysis, legal obligations and operational dependencies, then test whether the recovery environment can meet them. The NIST Cybersecurity Framework 2.0 places recovery within a broader cycle that also includes governance, identification, protection, detection and response (NIST, 2024). Such integration avoids treating recovery as an isolated storage team responsibility.

Three distinctions are central to the literature. First, high availability minimizes interruption through redundancy and rapid failover, but it may preserve malicious changes if replicas trust the same data stream. Second, backup retention provides historical versions, but only when access policies and deletion protection prevent an attacker from destroying all usable copies. Third, cyber recovery demands establishing a trustworthy execution environment before recovering selected data. The systems engineering view of cyber resilience stresses anticipation, absorption, recovery and adaptation despite adverse conditions (NIST, 2021). These capabilities address different failure mechanisms and should not be represented as interchangeable products.

Research on cloud orchestration and resource provisioning demonstrates why multi provider recovery is difficult. Providers differ in snapshot semantics, metadata, identity models, networking and automation interfaces (Tomarchio et al., 2020; Li et al., 2022). Studies of hybrid cloud adoption additionally identify concerns about quality of service and security governance (Khan et al., 2021). An alternate region or cloud account is thus not automatically an independent recovery domain. Its independence depends on which credentials, policies, encryption keys, management tools and network paths remain shared with production.

Distributed applications deepen this challenge. Berardi et al. (2022) catalogue recurring microservice security concerns, while Soldani and Brogi (2023) explain how anomalies and failures propagate across interacting services. Rezaei Nasab et al. (2023) document practical security activities in microservice development, and Jayalath et al. (2024) identify vulnerability patterns arising from architectural complexity. Together, these studies indicate that backup planning must include configuration states, secrets, service discovery, deployment manifests and dependency relationships, not only application databases. The literature supports dependency aware recovery but contains less direct evidence concerning verified end to end restoration in Saudi enterprises.

3 Research aim, objectives and questions

The aim of this review is to synthesize current evidence into a practical framework for protecting enterprise recovery assets and restoring trusted digital services in Saudi Arabia under Vision 2030. Unlike studies focused primarily on cloud placement or infrastructure utilization, the investigation examines how organizations can prove recoverability under both non adversarial failure and deliberate cyber compromise. Its contribution is an integrated assurance model connecting technical recovery design, business priorities and regulatory evidence.

Four objectives guide the analysis. The first is to distinguish backup integrity, disaster recovery performance and cyber recovery trustworthiness. The second is to evaluate architectural mechanisms for protecting recovery data, administrative authority and restoration environments. The third is to interpret relevant Saudi continuity, data and cybersecurity requirements alongside international standards. The fourth is to derive measurable readiness indicators and an implementation sequence suitable for complex enterprise services. These objectives lead to three questions: Which controls create independently recoverable enterprise assets? How should organizations establish trust before resuming critical services? Which governance and testing practices can demonstrate that recovery arrangements support Saudi digital transformation?

4 Review methodology

A structured narrative review with transparent thematic synthesis was selected because the topic spans storage engineering, adversarial security, continuity management, cloud operations and Saudi policy. The goal was to compare mechanisms and synthesize design principles, not to estimate a pooled intervention effect. Relevant material was identified from peer reviewed journal publications and official sources dated from January 2020 through December 2025. Core search combinations included enterprise backup, ransomware recovery, immutable storage, disaster recovery, cyber resilience, clean room restoration, multi cloud continuity, identity protection, data integrity and Saudi digital government controls.

Source selection prioritized publications that explained technical mechanisms, governance obligations or generalizable operational challenges. Review articles, empirical studies and vendor neutral standards were favored over promotional product literature. Saudi regulatory and policy materials were selected from the National Cybersecurity Authority, Digital Government Authority, Communications, Space and Technology Commission and Vision 2030 reporting. International benchmarks included NIST, ISO/IEC, CISA and ENISA. Earlier foundational literature was excluded from the final reference list to satisfy the 2020–2025 date constraint, although established concepts such as RTO and RPO are described in their documented form.

The review used a qualitative evidence matrix recording source type, publication period, primary contribution, limitations and relevance to enterprise recovery. Included documents were grouped into six analytical themes: service criticality; backup protection; identity and storage isolation; incident containment and restoration trust; recovery verification; and Saudi governance. Claims were compared across academic and institutional sources to separate broadly supported principles from design judgments proposed by this paper. No unreported database hit counts, independent reviewer agreement scores or experimental tests are asserted. The resulting framework is explicitly conceptual and requires validation against operational data.

This methodology has strengths and constraints. It captures connections that specialized storage or cybersecurity reviews often treat separately, but it cannot support causal estimates or claims that one commercial technology has universally superior recovery performance. Government guidance defines obligations and recommended practice, not measured outcomes. Peer reviewed studies offer analytical depth but may employ datasets, architectures or threat assumptions unlike those of Saudi institutions. These differences are retained in the discussion rather than masked by an artificial precision in evidence synthesis.

5 Findings: Enterprise backup protection

The first finding is that backup design should begin with service classification rather than storage capacity. Business owners need an inventory of essential services, data sets, recovery dependencies and minimum acceptable functionality. Financial ledgers, patient systems, government portals, industrial monitoring and general collaboration platforms have different consequences of lost transactions or delayed restoration. Criticality should determine frequency of copies, retention periods, integrity checking, geographic placement and the amount of reserved restoration capacity. NIST control guidance provides a broad basis for contingency planning and access governance, while ISO/IEC 27001:2022 establishes a management system context for systematic risk decisions (NIST, 2020b; ISO/IEC, 2022a).

The second finding concerns the independence of recovery copies. A resilient program may combine frequent operational snapshots, versioned backups and a separately administered protected copy. The familiar principle of maintaining multiple copies across distinct media or domains is valuable, yet numerical copy counts do not by themselves establish security. Backups controlled by the same privileged identity, reachable through the same management network or dependent on one compromised key management system can fail together. NIST storage guidance recommends isolation, access control and restoration assurance, while the CISA ransomware guide emphasizes offline, encrypted backups and routine testing (NIST, 2020a; CISA et al., 2023).

Immutability can make designated backup objects resistant to premature modification or deletion during a defined retention period. Nevertheless, an immutable repository can still contain

maliciously altered data if compromise predates the protected copy. Likewise, an offline archive may be unusable when credentials, software versions, decryption keys or restore procedures are unavailable. Consequently, immutability is one defense within an integrated design. Its effectiveness depends on enforceable retention, independent administrators, secure key custody, inventory completeness, monitored replication and a verified path from stored data to functioning services. Organizations should document these dependencies before treating an archive as a recovery guarantee.

The third finding is that backup operations themselves require monitoring. Unusual changes in backup schedules, sudden increases in data entropy, mass deletions, failed jobs, shortened retention and administrative role assignments can indicate attack activity or control deterioration. Alerts must reach monitoring channels not wholly controlled by the same operators whose accounts might be compromised. ISO/IEC 27002:2022 addresses information security controls, while NIST zero trust guidance supports continuing verification of access rather than implicit trust based on network location (ISO/IEC, 2022b; NIST, 2020c). Restoration readiness reporting should include coverage gaps and copy usability, not merely completed job counts.

6 Findings: Disaster recovery and service dependencies

An effective disaster recovery plan begins with an application dependency map. Database recovery can fail if upstream identity, DNS, network routing, application secrets, enterprise integration services or vendor licenses remain unavailable. For a revenue platform, correct restoration order might require authentication and network segmentation before message queues, databases and application servers are brought online. Recovery teams should document both technical dependencies and business sequencing, including manual workarounds. Alonso et al. (2023) show that distributed multi cloud applications face interoperability and runtime management barriers that are consequential during restoration, not only during normal service deployment.

Different service tiers justify different approaches. Synchronous or near synchronous replication can limit data loss but introduces network and consistency dependencies. Asynchronous copies support distance between failure domains at the cost of potentially larger RPOs. Periodic full and incremental backups offer historical recovery points yet may increase restoration time when long chains must be reconstructed. The architecture should balance these trade offs against business consequences and attack scenarios. Li et al. (2022) describe resource selection complexity across cloud services, reinforcing the need to account for restoration compute, network throughput and licensing rather than evaluating the backup repository alone.

Hybrid recovery may use on premise resources, approved cloud environments or a dedicated recovery account, depending on classification and contractual requirements. Recovery design must also consider large scale data movement, egress charges, system compatibility and the ability to rebuild infrastructure from verified configuration. Infrastructure as code and documented application manifests can reduce configuration drift, but automation should not blindly redeploy compromised images or secrets. Service owners should distinguish a functioning technical endpoint from successful business transactions. A restored portal that displays a login page while its core workflow remains unavailable should not be declared recovered.

Recovery orchestration must preserve accountability. Automation can execute preapproved steps, capture timestamps and reduce operator error, but decisions about destructive restoration, evidence preservation and acceptance of data loss require defined authority. Runbooks should identify escalation paths when clean copies are uncertain or when the nominal RTO conflicts with forensic

needs. NIST incident response guidance encourages embedding response coordination within broader cybersecurity risk management (NIST, 2025a). This alignment matters because the incident commander, business owner, security lead and technology recovery team may otherwise pursue incompatible priorities during the same outage.

7 Findings: Cyber recovery under active compromise

Cyber recovery begins with the possibility that a threat actor has already obtained privileged access to production and backup systems. Before reconnecting recovery assets, responders should determine whether the suspected attack remains active, identify compromised identities and assess the confidence of candidate recovery points. The ENISA threat landscape highlights ransomware and threats to availability and data as consequential concerns (ENISA, 2024). McIntosh et al. (2024) further show that recovering encrypted files does not resolve data exfiltration, persistence or continuing extortion risk. Restoration therefore requires a combined cybersecurity and business decision, not simply execution of a storage job.

An isolated recovery environment, sometimes described operationally as a clean room, is a deliberately controlled environment for assessing and rebuilding services without granting untrusted production systems direct authority over the process. It should use separately protected administrative identities, monitored access, tightly controlled data movement, validated images and a documented approval sequence. Network isolation may be physical, logical or organizational; the essential property is that a compromise affecting production cannot automatically revoke, delete or corrupt protected recovery resources. Zero trust architecture provides a useful principle of explicit authorization and least privilege, while NIST guidance on cloud native zero trust extends those ideas across distributed application components (NIST, 2020c; NIST, 2023).

Clean restoration does not mean every file must be declared malicious or every application rebuilt manually. It means that chosen restore points and rebuilt services are subjected to checks matched to the threat. Those checks may examine malware indicators, unexpected privileged accounts, integrity baselines, configuration drift, key rotation and application transactions. The standards based zero trust implementation work published by NIST in 2025 offers examples of integrating identity, policy enforcement and visibility (NIST, 2025b). Such controls support restoration assurance, but their suitability depends on operational complexity and on the organization's ability to maintain the required independent trust anchors.

Data compromise creates additional dilemmas. A clean backup may predate a large volume of legitimate transactions, whereas a recent backup may contain attacker persistence or corrupted records. Recovery leaders must select a point that balances integrity, recoverable transactions and acceptable business loss, then document any reconciliation needed. This decision cannot be reduced to the newest available timestamp. Cyber recovery plans should define a process for approving exceptional RPO deviations, replaying verified transaction logs and notifying affected stakeholders. No architecture can guarantee zero loss under all adversarial scenarios, but transparent criteria can prevent hurried restoration from introducing further damage.

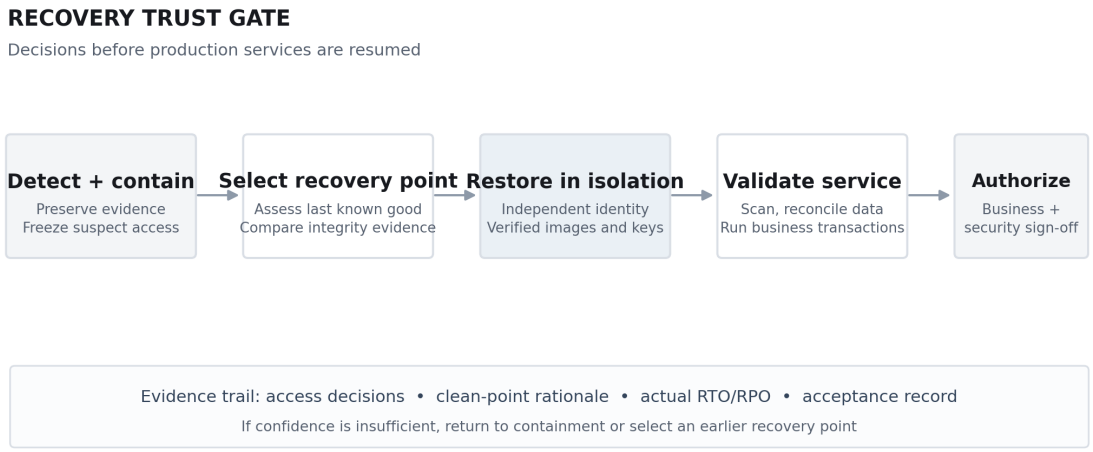


Figure 1. Trust-aware cyber recovery decision path (original conceptual illustration).

8 Findings: Assurance through testing and measurement

Backups are only an input to recovery; demonstrated restoration is the outcome. Job monitoring should be supplemented by isolated restore exercises, application level validation and scenarios that remove assumed dependencies. A technically successful restore should be checked for usable records, valid user permissions, functioning integrations and acceptable service performance. The CISA ransomware guide and Saudi digital government continuity controls both emphasize planned preparedness and effectiveness testing (CISA et al., 2023; Digital Government Authority, 2025). Evidence from exercises can reveal missing keys, expired certificates or insufficient standby capacity before an actual disruption.

An assurance dashboard should combine four measurement families. Coverage indicators establish whether every high priority service and required configuration element has a protected recovery path. Integrity indicators assess checksums, restore point screening and protection against unauthorized deletion. Performance indicators compare achieved recovery duration and actual data loss with approved targets. Governance indicators measure test completion, unresolved exceptions, access review outcomes and the currency of dependencies and runbooks. None is sufficient alone. A high percentage of completed backup jobs can coexist with an inability to restore the institution's most important service.

Testing can progress from low disruption evidence checks toward more demanding exercises. At the first level, teams verify inventory, retention locks, encryption key availability and restoration documentation. At the second, they restore selected datasets and validate integrity. At the third, they recover a complete application chain in an isolated environment while deliberately withholding production credentials or connectivity. At the fourth, a coordinated exercise tests leadership decisions, communications, cross provider actions and the resumption of a prioritized business process. NIST cyber resiliency engineering supports iterative assessment of these organizational capabilities (NIST, 2021).

Recovery results should report elapsed time from incident declaration or authorized exercise start through business acceptance, with milestones for containment, infrastructure readiness, data restoration and transaction validation. This decomposition permits realistic comparison and helps

locate bottlenecks. Time targets must not reward unsafe shortcuts: a quick restoration that reintroduces compromised credentials may reduce measured downtime while worsening overall harm. Soldani and Brogi (2023) underline the importance of root cause analysis in distributed services, suggesting that visibility into failure propagation should accompany speed focused indicators. Table 1 converts these principles into measurable evidence.

Table 1. Suggested enterprise cyber recovery assurance indicators (conceptual).

Assurance measure	Required evidence	Decision supported
Protected coverage	Mapped critical services and protected datasets	Which activities lack a usable recovery path?
Recovery point integrity	Restore point checks and approved clean point rationale	Which copies can be trusted?
Isolation effectiveness	Distinct identities, keys and management access	Can production compromise affect recovery?
Actual RTO/RPO	Timed full chain exercises and reconciled data loss	Are business recovery targets achievable?
Service acceptance	Verified workflows, integrations and sign off	Is the recovered service genuinely usable?
Governance closure	Exercise findings, owners, dates and exceptions	Which residual risks require management action?

9 Saudi Arabian governance and Vision 2030 alignment

Saudi requirements should be interpreted through the actual sector, entity classification, service arrangement and data involved. The National Cybersecurity Authority's Essential Cybersecurity Controls establish a national baseline for in scope entities; its Cloud Cybersecurity Controls address obligations in cloud relationships, and Data Cybersecurity Controls focus on protection of data (National Cybersecurity Authority, 2024a, 2024b, 2022). These instruments are complementary but not interchangeable. A compliance mapping should specify the applicable version, responsible party, required evidence and exceptions, rather than presenting every Saudi enterprise as subject to identical obligations.

Government entities also operate within Digital Government Authority guidance that connects organizational risk management, continuity and technology choices. The government cloud adoption guideline discusses migration and deployment considerations, while the 2025 risk and continuity controls call for structured responsibilities and ongoing preparedness (Digital Government Authority, 2023, 2025). Cloud service provider eligibility and contractual obligations must be considered alongside the Communications, Space and Technology Commission's service provisioning regulations (CST, 2023). Together these sources suggest that recovery location, data classification, service contracts, subcontractors and exit arrangements should be resolved before an incident forces urgent decisions.

A practical governance model separates three accountabilities. Business executives own service criticality and acceptable operational disruption. Technology teams own recoverable architectures, backup procedures and tested restoration dependencies. Security and compliance functions assess access segregation, evidence quality, threat exposure and regulatory fit. A cross functional continuity forum should resolve conflicts and oversee exceptions. This governance structure is consistent with

ISO/IEC 27031:2025 and with Saudi guidance, but each organization must tailor roles to its actual structure. No single committee or technical product substitutes for documented decision rights and management attention.

Vision 2030 alignment should be expressed through demonstrable service outcomes rather than general references to innovation. Examples include protecting access to digital government services, maintaining trusted financial records, restoring connected logistics workflows and preserving essential enterprise transactions. The Vision 2030 reporting context supports the importance of reliable digital infrastructure, but it does not prove that a particular recovery design has been deployed or certified nationally (Saudi Vision 2030, 2025). This review therefore proposes a Saudi oriented implementation framework, not an assertion of nationwide compliance. Organizations must verify legal requirements, including data location and sector rules, with their responsible authorities.

10 Proposed Enterprise Cyber Recovery Assurance Framework

The proposed Enterprise Cyber Recovery Assurance Framework comprises five decision layers supported by continuous evidence. The governance layer establishes service ownership, applicable requirements, risk appetite and approval authority. The criticality layer translates business impact analysis into recoverable service groupings, dependencies, target restoration order and acceptable data loss. The protection layer manages isolated backup copies, retention controls, key custody and identity segregation. The recovery layer provides a controlled environment for containment, clean restoration and business validation. The assurance layer collects exercise results, unresolved weaknesses and management decisions. Figure 2 displays this architecture as a connected operating model.

Implementation should proceed through a sequence of controlled decisions. Organizations first inventory business services and classify essential data. They next identify shared failure domains across production, backup consoles, identity providers and administrative networks. They then protect at least one logically or physically independent recovery path appropriate to the risk. Subsequent work establishes documented restore procedures and test environments. Finally, recovery exercises produce signed evidence and prioritized remediation. This sequence intentionally begins with dependencies and governance rather than commercial platform selection. It is compatible with ISO/IEC 27001 management system thinking and NIST's comprehensive security control approach (ISO/IEC, 2022a; NIST, 2020b).

The framework distinguishes normal state availability from adverse state trust. During ordinary disruption, authorized automation can accelerate failover or recovery. During suspected compromise, the same mechanisms may need to be paused until identity trust and data cleanliness are reassessed. Figure 1 illustrates this trust aware decision path. Organizations should preapprove conditions for invoking cyber recovery, preserving forensic evidence and authorizing restored services. Rehearsing these decisions reduces dependence on improvised judgment when operational pressure is highest. The design does not assume one storage vendor, one cloud provider or a uniform recovery architecture for all applications.

Table 2 compares the mechanisms used for routine backup, disaster recovery and adversarial cyber recovery. Its purpose is not to rank one method as universally superior: the three capabilities address complementary risks and require a layered strategy. Cyber recovery generally demands stronger independence and more cautious validation, while conventional operational recovery may prioritize speed after hardware failure. Resource constraints remain important, particularly for smaller

organizations. Tiered protection allows critical services to receive enhanced isolation and testing without imposing the same cost and complexity on every enterprise application.

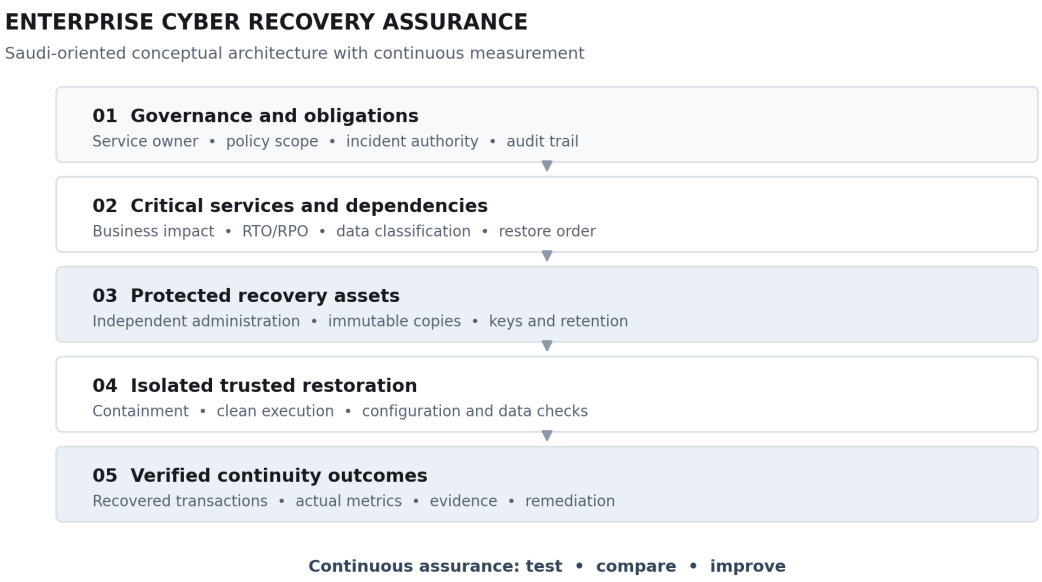


Figure 2. Saudi-oriented Enterprise Cyber Recovery Assurance Framework (original conceptual illustration).

Table 2. Differences between conventional operational recovery and cyber recovery.

Capability	Operational disruption	Adversarial compromise
Recovery goal	Resume processing after failure	Resume processing with independently established trust
Recovery copy	Recent backup or replica	Protected historical copy with verified integrity
Administrative access	Authorized standard operators	Separated identities; privileged access revalidated
Restore environment	Standby site or alternative capacity	Isolated clean environment with controlled data ingress
Acceptance test	Service connectivity and basic function	Malware checks, data reconciliation and business transactions
Key failure	Capacity or dependency unavailable	Recovery assets, credentials or restored data remain compromised

11 Discussion, limitations and future research

The synthesis reveals a recurring distinction between possession of recovery technology and possession of recovery capability. Replication, immutable archives and redundant infrastructure are necessary in some contexts, but each protects only selected parts of the failure chain. A mature program asks whether the organization can identify a trustworthy recovery point, retrieve the necessary configuration, operate an independent control plane, reconstruct service dependencies and prove business acceptance. Studies of multi cloud security and application operations suggest that complexity grows when organizations distribute technologies without corresponding operational coherence (Alyas et al., 2022; Alonso et al., 2023).

One important tension is between protection and restoration speed. Deeply isolated recovery data may be less accessible to an attacker but require additional time and coordination for retrieval. Frequent replication can produce low nominal data loss but rapidly propagate corruption. Automated orchestration improves consistency yet may amplify mistakes when trust conditions are not verified. The appropriate solution therefore depends on criticality, threat exposure, regulatory context and tested human capability. Recovery metrics should make these trade offs visible rather than compressing them into a single readiness percentage. Decision makers need evidence about uncertainty and residual risk as well as performance.

The review is limited by its narrative design, the heterogeneity of included sources and the absence of primary organizational fieldwork. There is no claim of a comprehensive bibliometric search, pooled effect size or experimentally verified effectiveness of the proposed framework. Policy materials describe normative expectations, while research papers may evaluate narrowly defined systems. The literature also provides limited public, comparable evidence on Saudi enterprise recovery performance. The framework should accordingly be tested through interviews, anonymized exercise records and cross sector case studies before conclusions are drawn about feasibility, cost or achievable RTOs.

Future research could compare how Saudi government, finance, healthcare, industrial and logistics organizations measure trustworthy restoration. Controlled exercises could examine the effect of shared identity compromise on backup usability, the relationship between immutability and verified data cleanliness, and the practical cost of maintaining isolated recovery environments. Further work should evaluate how cloud portability, automated dependency reconstruction and policy based evidence capture influence end to end recovery. Quantitative studies need transparent service definitions and comparable disruption scenarios; otherwise, published recovery durations will remain difficult to interpret. Such studies could help convert the present conceptual contribution into an empirically validated assessment instrument.

12 Conclusion

Enterprise resilience is demonstrated when an organization can restore trusted services under the conditions that actually threaten them. This review shows why successful backup jobs, redundant storage and conventional disaster recovery plans are insufficient indicators of cyber recovery readiness. Effective arrangements combine business led classification, independently protected copies, separated administrative authority, validated restore points, dependency aware orchestration and repeated exercises that end with business acceptance. Saudi cybersecurity and digital government guidance provides important governance context, while international standards and research clarify the underlying technical and organizational mechanisms.

The proposed Enterprise Cyber Recovery Assurance Framework connects these elements through governance, criticality, protection, recovery and evidence. Its principal contribution is a practical shift from counting recovery assets toward proving recoverable services, with explicit attention to adversarial compromise. For organizations contributing to Vision 2030, the central objective is not maximum backup volume or maximum cloud adoption; it is continuity of secure, dependable digital operations that can be measured and improved. Further empirical evaluation is needed before the framework can be treated as a validated national model.

References

[1] Alyas T, Alissa K, Alqahtani M, Faiz T, Alsaif SA, N. and Naqvi HH (2022). Multi-Cloud Integration Security Framework Using Honeypots. *Mobile Information Systems*, 2022, 2600712. https://doi.org/10.1155/2022/2600712

[2] Alonso J, Orue-Echevarria L, Casola V, Torre AI, Huarte M, E. and Lobo JL (2023). Understanding the challenges and novel architectural models of multi-cloud-native applications: a systematic literature review. *Journal of Cloud Computing*, 12, 6. https://doi.org/10.1186/s13677-022-00367-6

[3] Berardi D, Giallorenzo S, Mauro J, Melis A, F. and Prandini M (2022). Microservice security: a systematic literature review. *PeerJ Computer Science*, 8, e779. https://doi.org/10.7717/peerj-cs.779

[4] Communications, Space and Technology Commission (CST) (2023). Cloud Computing Services Provisioning Regulations, Version 4 [Internet]. *2023*. Available from: https://www.cst.gov.sa/en/regulations-and-licenses/regulations/cloud-computing-services-provisioning-regulations

[5] Cybersecurity and Infrastructure Security Agency, Federal Bureau of Investigation, National Security Agency and Multi-State Information Sharing and Analysis Center (CISA et al.) (2023). #StopRansomware Guide [Internet]. 2023. Available from: https://www.ic3.gov/CSA/2023/231019.pdf

[6] Digital Government Authority (2023). Guideline for Cloud Computing Adoption by Government Agencies, DGA-1-2-5-212, Version 1.0 [Internet]. *2023*. Available from: https://dga.gov.sa/sites/default/files/2023-09/Guideline%20for%20Cloud%20Computing%20Adoption%20by%20Government%20Agencies%20%E2%80%93%20V1.0_1.pdf

[7] Digital Government Authority (2025). Controls of Risk Management and Business Continuity for Digital Government, DGA-1-2-5-107, Version 4.0 [Internet]. *2025*. Available from: https://dga.gov.sa/sites/default/files/2025-09/Controls%20of%20Risk%20Management%20and%20Business%20Continuity%20for%20Digital%20Government-V4.0.pdf

[8] European Union Agency for Cybersecurity (ENISA) (2024). ENISA Threat Landscape 2024 [Internet]. *2024*. Available from: https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024

[9] International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) (2022a). ISO/IEC 27001:2022: Information security, cybersecurity and privacy protection: Information security management systems: Requirements [Internet]. *2022a*. Available from: https://www.iso.org/standard/27001

[10] International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) (2022b). ISO/IEC 27002:2022: Information security, cybersecurity and privacy protection: Information security controls [Internet]. *2022b*. Available from: https://www.iso.org/standard/75652.html

[11] International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) (2025). ISO/IEC 27031:2025: Cybersecurity: Information and communication technology readiness for business continuity [Internet]. *2025*. Available from: https://www.iso.org/standard/27031

[12] Jayalath RK, Ahmad H, Goel D, M. S. and Ullah F (2024). Microservice Vulnerability Analysis: A Literature Review With Empirical Insights. *IEEE Access*, 12, 155168–155204. https://doi.org/10.1109/ACCESS.2024.3481374

[13] Khan SU, Khan HU, N. and Khan RA (2021). Challenges and Their Practices in Adoption of Hybrid Cloud Computing: An Analytical Hierarchy Approach. *Security and Communication Networks*, 2021, 1024139. https://doi.org/10.1155/2021/1024139

[14] Li X, L. and Liu S (2022). A survey of resource provisioning problem in cloud brokers. *Journal of Network and Computer Applications*, 203, 103384. https://doi.org/10.1016/j.jnca.2022.103384

[15] McIntosh T, Sušnjak T, Liu T, Xu D, Watters P, Liu D, et al (2024). Ransomware Reloaded: Re-examining Its Trend, Research and Mitigation in the Era of Data Exfiltration. *ACM Computing Surveys*, 57(1). https://doi.org/10.1145/3691340

[16] National Cybersecurity Authority (2022). Data Cybersecurity Controls (DCC 1:2022) [Internet]. *2022*. Available from: https://nca.gov.sa/en/regulatory-documents/controls-list/dcc/

[17] National Cybersecurity Authority (2024a). Essential Cybersecurity Controls (ECC 2:2024) [Internet]. *2024a*. Available from: https://nca.gov.sa/en/regulatory-documents/controls-list/ecc/

[18] National Cybersecurity Authority (2024b). Cloud Cybersecurity Controls (CCC 2:2024) [Internet]. *2024b*. Available from: https://nca.gov.sa/en/regulatory-documents/controls-list/ccc/

[19] National Institute of Standards and Technology (NIST) (2020a). Security Guidelines for Storage Infrastructure, SP 800–209. https://doi.org/10.6028/NIST.SP.800-209

[20] National Institute of Standards and Technology (NIST) (2020b). Security and Privacy Controls for Information Systems and Organizations, SP 800-53 Revision 5. https://doi.org/10.6028/NIST.SP.800-53r5

[21] National Institute of Standards and Technology (NIST) (2020c). Zero Trust Architecture, SP 800–207. https://doi.org/10.6028/NIST.SP.800-207

[22] National Institute of Standards and Technology (NIST) (2021). Developing Cyber-Resilient Systems: A Systems Security Engineering Approach, SP 800-160 Volume 2 Revision 1. https://doi.org/10.6028/NIST.SP.800-160v2r1

[23] National Institute of Standards and Technology (NIST) (2023). A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Cloud Environments, SP 800-207A. https://doi.org/10.6028/NIST.SP.800-207A

[24] National Institute of Standards and Technology (NIST) (2024). The NIST Cybersecurity Framework (CSF) 2.0, CSWP 29. https://doi.org/10.6028/NIST.CSWP.29

[25] National Institute of Standards and Technology (NIST) (2025a). Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile, SP 800-61 Revision 3. https://doi.org/10.6028/NIST.SP.800-61r3

[26] National Institute of Standards and Technology (NIST) (2025b). Implementing a Zero Trust Architecture, SP 1800–35. https://doi.org/10.6028/NIST.SP.1800-35

[27] Rezaei Nasab A, Shahin M, Hoseyni Raviz SA, Liang P, A. and Lenarduzzi V (2023). An empirical study of security practices for microservices systems. *Journal of Systems and Software*, 198, 111563. https://doi.org/10.1016/j.jss.2022.111563

[28] Saudi Vision 2030 (2025). Vision 2030 Annual Report 2024 [Internet]. *2025*. Available from: https://www.vision2030.gov.sa/media/25042025ev1/En-Annual%20Report-Vision2030-2024.pdf

[29] J. and Brogi A (2023). Anomaly Detection and Failure Root Cause Analysis in (Micro)Service-Based Cloud Applications: A Survey. *ACM Computing Surveys*, 55(3), 59. https://doi.org/10.1145/3501297

[30] Tomarchio O, D. and Di Modica G (2020). Cloud resource orchestration in the multi-cloud landscape: a systematic review of existing frameworks. *Journal of Cloud Computing*, 9, 49. https://doi.org/10.1186/s13677-020-00194-7