

Multi-Cloud Database Architecture for Saudi Enterprises: Improving Scalability, Data Sovereignty and Digital Resilience under Vision 2030.

Abstract:

Saudi enterprises are expanding cloud adoption to support digital services, analytics, artificial intelligence, enterprise applications, and data-intensive operations. As this transformation accelerates, database architecture becomes a strategic design problem rather than a simple infrastructure choice. A single-cloud deployment can provide elasticity and managed services, but it can also concentrate operational dependency, complicate exit strategies, and create tension between global cloud capabilities and national requirements for cybersecurity, data governance, privacy, and business continuity. Multi-cloud database architecture offers an alternative by distributing selected database capabilities across two or more cloud environments while retaining centralized governance and clear workload ownership. However, multi-cloud does not automatically improve resilience or sovereignty; poorly designed deployments can increase latency, consistency problems, security complexity, operational cost, and compliance risk. This paper develops a structured conceptual framework for multi-cloud database architecture for Saudi enterprises under Vision 2030. The study synthesizes recent research on multi-cloud systems, cloud-native databases, interoperability, security, data sovereignty, distributed data processing, and high availability, together with Saudi regulatory and digital-government sources. It proposes a six-domain architecture covering workload classification, sovereign data placement, database portability, cross-cloud resilience, unified security and observability, and lifecycle governance. The framework distinguishes active-active, active-passive, segmented, and portable-recovery patterns and links them to recovery objectives, consistency requirements, data sensitivity, and regulatory constraints. The paper also proposes a phased implementation roadmap and governance metrics for enterprise adoption. The central argument is that effective multi-cloud database strategy in Saudi Arabia should be policy-driven, workload-specific, and resilience-oriented: sensitive data placement must be determined by classification and lawful transfer conditions; portability must be designed at schema, data, tooling, and operational layers; and resilience must be validated through measurable recovery, failover, and cyber-recovery exercises. The resulting model supports Vision 2030 objectives by combining scalable cloud adoption with stronger digital resilience, data control, and reduced concentration risk.

Keywords: multi-cloud; cloud database; data sovereignty; digital resilience; database architecture; cloud portability; high availability; disaster recovery; data governance; Saudi Vision 2030.

1. Introduction

Cloud computing has become a foundational component of enterprise digital transformation. Databases that previously operated in fixed on-premises environments are increasingly delivered through infrastructure-as-a-service, platform-as-a-service, database-as-a-service, managed distributed databases, and cloud-native data platforms. These models can accelerate provisioning, improve elasticity, reduce infrastructure-management overhead, and provide access to advanced security, analytics, and automation capabilities. At the same time, moving mission-critical data platforms into the cloud changes the enterprise risk model. Availability becomes partly dependent on provider regions and services; data governance must account for external

processing; operational teams must manage new identity and configuration layers; and long-term architecture can become coupled to provider-specific database features.

Multi-cloud architecture addresses some of these concerns by using services from more than one cloud provider or cloud environment. The term is broader than simply maintaining accounts with several providers. A meaningful multi-cloud database architecture deliberately allocates data, database services, recovery copies, or supporting workloads across independent environments according to business and technical objectives. These objectives may include resilience, geographic diversity, bargaining power, cost management, access to specialized capabilities, regulatory alignment, data sovereignty, or avoidance of strategic lock-in. Research on multi-cloud systems consistently identifies flexibility and reduced dependency as potential benefits while also identifying interoperability, orchestration, security, consistency, and operational complexity as central challenges (Tomarchio, Calcaterra and Di Modica, 2020; Alonso et al., 2023).

The database layer deserves special treatment because data is stateful, persistent, highly interconnected, and difficult to move quickly at scale. Stateless application components can often be redeployed with infrastructure-as-code, containers, or standardized deployment tooling. Databases contain large volumes of durable state, transaction histories, schema dependencies, security models, stored procedures, replication configurations, and recovery requirements. Vendor lock-in can therefore become strongest at the data layer. Shaik and Kalyanasundaram (2024) show that database migration is constrained by data volume, bandwidth, dependencies, and acceptable downtime, while cloud-native database research highlights how new database architectures optimize elasticity and availability through techniques such as disaggregated storage and distributed execution (Dong et al., 2024). A multi-cloud strategy that ignores these characteristics may create an architecture that appears diversified but cannot actually fail over, migrate, or recover when required.

For Saudi Arabia, the topic has additional strategic relevance. The Kingdom's Cloud First direction encourages cloud adoption for government technology investments, while the Digital Government Authority (DGA) has embedded cloud architecture and migration planning within digital transformation standards. The Saudi Data & AI Authority (SDAIA) positions data as a national asset and emphasizes advanced digital infrastructure, reliability, resilience, cybersecurity, and government cloud capabilities. The National Cybersecurity Authority (NCA) has issued Cloud Cybersecurity Controls (CCC-2:2024) and Data Cybersecurity Controls, while the Communications, Space and Technology Commission (CST) regulates cloud service provisioning. The Personal Data Protection Law (PDPL) and its transfer regulation also make cross-border data movement a legal and governance consideration rather than a purely technical networking decision (MCIT, 2020; DGA, 2024; NCA, 2024; SDAIA, 2023a; CST, 2023).

This paper therefore asks: how should Saudi enterprises design multi-cloud database architecture so that scalability, data sovereignty, and digital resilience reinforce rather than undermine one another? The study develops a conceptual framework rather than presenting a fabricated empirical benchmark. It combines peer-reviewed research, standards, and Saudi regulatory sources to define architectural patterns, decision criteria, security controls, and an implementation roadmap. The framework is intended for heterogeneous enterprise estates that may include relational databases, cloud-managed databases, distributed SQL platforms, analytics stores, and legacy systems across public, private, and government cloud environments.

2. Literature Review and Conceptual Background

2.1 From single-cloud adoption to multi-cloud architecture

Cloud adoption initially emphasized migration from privately managed infrastructure to a single provider. Over time, enterprises adopted multiple providers through mergers, business-unit autonomy, specialized software services, geographic requirements, or deliberate diversification. Multi-cloud architecture formalizes this reality. Tomarchio, Calcaterra and Di Modica (2020) describe multi-cloud orchestration as a problem involving heterogeneous resources, portability, brokerage, policy, and lifecycle management. Alonso et al. (2023), in a systematic literature review of multi-cloud-native applications, identify unresolved challenges involving portability, runtime reconfiguration, state synchronization, data consistency, security monitoring, network integration, and non-functional optimization.

These findings are particularly important for database platforms. A database is not simply a compute workload that can be restarted elsewhere. It has persistent state, transaction ordering, recovery logs, data locality requirements, backup chains, encryption keys, user privileges, and application dependencies. The multi-cloud problem at the database layer therefore has at least four dimensions: service portability, data portability, operational portability, and policy portability. Service portability asks whether the database engine or equivalent capability can run in another environment. Data portability concerns whether the data can be moved efficiently and without unacceptable loss. Operational portability concerns whether backup, monitoring, patching, security, and recovery procedures remain executable. Policy portability concerns whether the same governance, access, and compliance requirements can be enforced consistently.

Ramalingam and Mohan (2021) show that portability and interoperability remain constrained by proprietary interfaces and incomplete standardization. This means that using two providers does not necessarily create a portable architecture. An enterprise can become locked into two separate proprietary stacks. Multi-cloud therefore requires deliberate abstraction only where abstraction contributes measurable business value. The objective is not to reduce every provider to a lowest common denominator. Instead, enterprises should separate portable core data capabilities from provider-specific accelerators and document the recovery or migration path for each dependency.

2.2 Cloud-native databases, scale, and the stateful challenge

Distributed data processing research also demonstrates that cross-environment scaling introduces measurable overhead. Ullah et al. (2024) found that hybrid-cloud data processing performance depends on data movement and the distribution of nodes between environments. For transaction databases, this effect can be more severe because synchronous replication across long-distance links increases commit latency and because network partitions force choices about availability and consistency. For this reason, a Saudi enterprise should not interpret multi-cloud scalability as the ability to spread every transaction across every provider. A more practical principle is workload-aware scaling: keep latency-sensitive transactional paths close to the authoritative data, while distributing read replicas, analytical workloads, backups, or recovery environments according to business objectives.

2.3 Data sovereignty as an architectural property

Data sovereignty is sometimes reduced to the question of where a server is physically located. In practice, sovereignty is multidimensional. It includes legal jurisdiction, control over encryption and access, data classification, administrative authority, transfer conditions, subcontractor relationships, operational visibility, and the ability to recover or move data without depending on a single external party. Galij, Pawlak and Grzyb (2024) argue that public-cloud data sovereignty requires both technical and governance mechanisms and show that sovereignty solutions can impose performance and economic trade-offs.

For Saudi enterprises, sovereignty must be interpreted through applicable Saudi requirements rather than generic international assumptions. SDAIA and the National Data Management Office (NDMO) establish national data-governance roles and emphasize protection of personal and sensitive data. The PDPL and the Regulation on Personal Data Transfer Outside the Kingdom allow transfers subject to legal conditions and safeguards; they do not support a simplistic assumption that every category of enterprise data must always remain in-country. At the same time, the NCA Cloud Cybersecurity Controls explicitly address cloud risks and data-localization requirements, and sector-specific or critical-system controls may impose additional restrictions (NCA, 2024; SDAIA, 2023a; SDAIA, 2023b).

Therefore, sovereignty should be embedded into the data-placement algorithm. Before selecting a region or provider, the enterprise must know the data class, data owner, regulatory basis, transfer restrictions, encryption requirements, recovery locations, and administrative-access model. Multi-cloud can strengthen sovereignty if it creates controlled choice and prevents critical data from becoming operationally captive. It can weaken sovereignty if uncontrolled replication causes sensitive data to cross boundaries or if centralized cloud-management tools expose metadata, keys, or backups outside approved environments.

2.4 Resilience, high availability, and concentration risk

Resilience is the ability to continue or restore services under failure, cyberattack, operational error, and external disruption. High availability (HA) usually addresses localized component failures through redundancy and automated failover, while disaster recovery (DR) addresses larger events through recovery in a separate environment. Xiong, Fowley and Pahl (2016) demonstrate a database-specific multi-cloud HADR pattern using synchronous and asynchronous replication to balance continuity and recovery objectives. Wang, Yang and Song (2020) similarly show that distributing a cloud database across multiple clouds can increase availability and reduce dependence on one service provider, although their SHAMC design also demonstrates that security and availability mechanisms can introduce computational overhead.

Multi-cloud is attractive because it can reduce concentration risk: a regional or provider-wide disruption does not automatically disable every copy of a critical database. However, architectural independence must be genuine. Two database instances in separate accounts but dependent on the same identity provider, domain-name service, network hub, key-management system, CI/CD pipeline, or operational team may share common failure modes. Digital resilience therefore requires dependency mapping, diversity at selected control points, tested failover, immutable or isolated recovery copies, and a clear distinction between high availability and cyber recovery.

2.5 Security and governance complexity

Security is one of the strongest arguments for and against multi-cloud. On one hand, diversified architecture can reduce single-provider dependency and create isolated recovery paths. On the other hand, every additional provider introduces new identity models, APIs, logging systems, security services, configuration rules, and skills requirements. Yeboah-Ofori et al. (2024) identify governance, access control, interoperability, configuration, and cross-cloud attack surfaces as significant challenges. Diningrat, Suhardi and Rahardjo (2025) similarly categorize multi-cloud security risks and emphasize the need for coordinated technical controls. Ali et al. (2025) highlight identity management, secure inter-cloud communication, regulatory compliance, and privacy-preserving controls as ongoing research priorities.

3. Methodology

This paper uses a structured narrative review and conceptual architecture method. It is not presented as a PRISMA systematic review, randomized experiment, or production benchmark. The goal is to synthesize technical and regulatory evidence into an implementable enterprise framework for Saudi Arabia. Sources were selected from five categories: multi-cloud architecture and orchestration; database and cloud-native data architecture; sovereignty, privacy, and data-protection research; resilience and database recovery; and Saudi official policies, regulations, and cybersecurity controls.

The synthesis followed four stages. First, literature findings were mapped to three primary objectives in the paper title: scalability, data sovereignty, and digital resilience. Second, recurring architectural risks were identified, including vendor lock-in, cross-cloud latency, data consistency, identity fragmentation, regulatory transfer risk, untested failover, shared dependencies, and operational complexity. Third, these risks were translated into design controls and workload-placement criteria. Fourth, a six-domain reference architecture and phased implementation roadmap were constructed.

The framework uses a workload-level decision model rather than prescribing one topology for the entire enterprise. For each database workload i and candidate environment j , a conceptual placement score can be expressed as:

$$P(i,j) = wsS(i,j) + wrR(i,j) + wpP(i,j) + wcC(i,j) + woO(i,j) - wlL(i,j) - wtT(i,j),$$

where S represents sovereignty and compliance fit, R resilience contribution, P performance fit, C cost fit, O operational capability, L latency penalty, and T portability or transition risk. The weights are not universal constants. They should be approved according to the workload's business criticality and regulatory context. The equation is used only to make decision logic explicit; it does not claim empirical measurement.

A second principle is that resilience must be evaluated with measurable service objectives. At a minimum, architectures should specify recovery time objective (RTO), recovery point objective (RPO), failover decision authority, replication mode, backup isolation, and a tested return-to-normal procedure. Data sovereignty must likewise be evidenced through classification, location, encryption ownership, transfer records where applicable, privileged-access controls, and provider/subprocessor governance.

4. Proposed Multi-Cloud Database Architecture for Saudi Enterprises

The proposed framework contains six domains: (1) workload and data classification, (2) sovereign placement and residency control, (3) portable database and integration architecture, (4) resilience and recovery engineering, (5) unified security and observability, and (6) lifecycle governance and FinOps. These domains are intentionally linked. A technically elegant replication design is unacceptable if it violates data-transfer conditions, while a compliant architecture is weak if recovery cannot be demonstrated.




Figure 1. Six-domain multi-cloud database architecture for Saudi enterprises. Source: author conceptualization based on reviewed literature and Saudi governance requirements.

4.1 Domain 1: workload and data classification

The first step is to classify workloads before selecting clouds. Each database should be assigned a business owner, technical owner, service tier, data classification, personal-data status, regulatory scope, latency requirement, maximum acceptable data loss, maximum outage time, expected growth profile, and dependency map. This inventory should distinguish system-of-record databases from caches, read replicas, analytical stores, development copies, and backup repositories.

Table 1 summarizes four practical patterns. The active-active pattern provides the highest immediate continuity potential but is the most difficult because distributed writes require conflict management, consensus, or database-native global transaction support. Active-passive is simpler and often more appropriate for mission-critical relational systems. Segmented multi-cloud places different domains or workloads on different providers, reducing enterprise concentration without cross-cloud transaction complexity. Portable recovery keeps production on one provider while maintaining tested data, infrastructure definitions, and recovery procedures for another; it provides strategic resilience without paying the full cost of continuously active duplication.

4.2 Domain 2: sovereign data placement and residency controls

A sovereign placement layer converts legal and policy requirements into enforceable technical rules. The decision should start with data classification and purpose. For personal data, cross-border replication must be evaluated under the PDPL and transfer regulation. For government or critical data, applicable NCA, DGA, NDMO, or sector-specific requirements must be considered. The design should record the lawful or policy basis for each primary copy, replica, backup, archive, and support-access path.

The architecture should also distinguish data residency from operational sovereignty. Data may physically reside in Saudi Arabia while the enterprise remains dependent on external control planes, proprietary tooling, or inaccessible encryption mechanisms. A stronger sovereignty model includes the ability to retrieve data in documented formats, maintain independent backups, control cryptographic keys or key policies, audit privileged operations, and execute a tested provider-exit or continuity plan.

4.3 Domain 3: portable database and integration architecture

Portability should be designed in layers. At the schema layer, teams should document proprietary data types, stored procedures, extensions, sequences, partitioning features, and security constructs. At the data layer, they should maintain tested export, change-data-capture, replication, and restore methods. At the application layer, database access should be encapsulated where possible so that connection routing, credentials, and read/write roles can change without redesigning the whole application. At the operational layer, backup, monitoring, patching, and incident procedures should be represented in automation and runbooks rather than retained only as provider-console knowledge.

Inter-cloud integration requires special caution. Synchronous transaction replication across providers can create latency and availability problems. For many Saudi enterprise workloads, asynchronous replication or change-data-capture is more practical for distant recovery targets. Data integration platforms should preserve transaction ordering where required, monitor replication lag, validate schema compatibility, and expose data-loss risk clearly. Where distributed reads are used, routing should consider locality, freshness, and data class.

4.4 Domain 4: resilience and recovery engineering

Digital resilience should be designed around failure scenarios rather than provider marketing claims. The enterprise should model at least five events: database-node failure, availability-zone failure, regional cloud failure, provider control-plane impairment, and cyber compromise. Each event may require a different recovery mechanism. High availability inside one region can address node or zone failures but does not solve provider concentration. Cross-region replicas may address regional failure but may still share provider identity and control planes. A second cloud can reduce these shared dependencies if the recovery path is operationally independent.

Resilience architecture should define criticality tiers. Tier 1 systems may require automated or orchestrated failover, near-real-time replication, frequent recovery exercises, and isolated immutable backups. Tier 2 systems may use warm standby with asynchronous replication. Tier 3 workloads may rely on infrastructure-as-code and backup restore. The objective is not to make every database active-active; it is to align cost and complexity with business impact.

Recovery testing must go beyond confirming that a replica exists. Exercises should measure actual RTO and RPO, verify application connectivity, validate identity and secrets, confirm DNS or routing changes, check data consistency, and demonstrate return to the primary environment. Cyber-recovery exercises should also test restoration from a point before malicious changes, with isolated credentials and clean infrastructure. A multi-cloud environment that has never been failed over is an architectural hypothesis, not proven resilience.

4.5 Domain 5: unified security and observability

Security governance should be standardized across clouds through common policy objectives. Identity should follow least privilege, strong authentication, privileged-role separation, and short-lived credentials where possible. Service identities should be preferred over embedded passwords. Encryption policies should define approved algorithms, key rotation, key administration, and the treatment of database backups and replication channels. Network connectivity between clouds should be authenticated, encrypted, segmented, monitored, and capacity-tested.

Cloud-specific logs should feed a normalized security and operations layer. This does not mean copying all sensitive telemetry to one external location; normalization can occur through approved regional collectors or federated analysis. The goal is a common view of database health, replication lag, backup status, privileged actions, configuration drift, vulnerability exposure, and security incidents. NCA's cloud and data cybersecurity controls make governance, protection, resilience, and data-lifecycle security central requirements for Saudi organizations in scope (NCA, 2022; NCA, 2024).

4.6 Domain 6: lifecycle governance, architecture review, and FinOps

Multi-cloud architecture must be governed throughout its lifecycle. An enterprise cloud architecture board should approve workload placement, acceptable proprietary dependencies, cross-border data flows, resilience tiers, and exception handling. The board should include database engineering, cloud architecture, cybersecurity, data governance, privacy/legal, business continuity, finance, and business ownership. This prevents multi-cloud from becoming a purely infrastructure-led program.

5. Architectural Patterns and Selection Criteria

A Saudi enterprise should select a pattern based on business criticality, transaction characteristics, data sensitivity, and recovery objectives. Table 1 compares the major patterns.

Table 1. Multi-cloud database architectural patterns and selection considerations.

Pattern	Primary use	Scalability / resilience value	Sovereignty considerations	Key cautions
Active-active multi-cloud	Simultaneous service across independent environments	Highest continuity potential; can distribute reads/writes	Data copies and transactions must remain within approved locations and transfer rules	Consistency, latency, conflict handling, cross-cloud network dependence, high cost
Active-passive multi-cloud	Primary database with secondary takeover environment	Strong DR and concentration-risk reduction with simpler write model	Secondary replicas/backups must meet the same classification and access controls	Replication lag, failover orchestration, idle/standby cost, return-to-primary complexity

Pattern	Primary use	Scalability / resilience value	Sovereignty considerations	Key cautions
Segmented multi-cloud	Different domains or applications assigned to different providers	Reduces enterprise-wide provider concentration; scales independently by domain	Each domain can use a placement policy matched to its data class	Cross-domain dependencies may recreate shared failure points; duplicated governance effort
Portable recovery	Production on one provider; tested build/restore path on another	Lower-cost strategic resilience and provider-exit readiness	Alternate recovery location must be approved before backups or data are transferred	Recovery is slower than always-on standby; portability must be tested with realistic data volumes
Hybrid sovereign core	Authoritative sensitive data in approved Saudi/private/government environment; scalable services elsewhere	Scales application/analytics while retaining controlled authoritative core	Strong alignment for sensitive data where cross-boundary exposure must be minimized	Integration latency, API/event design, derived-data governance, possible central-core bottleneck

Source: author synthesis informed by Xiong, Fowley and Pahl (2016), Alonso et al. (2023), Shaik and Kalyanasundaram (2024), and related reviewed literature.

Pattern selection should be evidence-led. Active-active deployment is justified only where concurrent service from independent environments creates a clear continuity or geographic benefit and the database technology can manage consistency across distance. Active-passive deployment is generally easier to govern for transactional systems because one environment remains authoritative while the second is engineered for takeover. Segmented multi-cloud reduces enterprise-wide concentration by placing independent domains on different providers, but it does not protect a particular workload unless that workload has its own recovery path. Portable recovery is a lower-cost alternative in which production uses one provider while the enterprise maintains tested backups, infrastructure definitions, build procedures, and recovery automation for another. A hybrid sovereign-core pattern can keep authoritative sensitive data in an approved Saudi environment while exposing controlled APIs, events, or derived datasets to scalable public-cloud services.

The chosen pattern should be recorded as part of the database service design together with the accepted consistency model, expected failover behavior, replication technology, data-location rules, dependency map, and provider-exit assumptions. Architecture review should also examine common-mode dependencies such as identity, network hubs, certificate services, key management, monitoring, and deployment pipelines. A nominally multi-cloud topology that shares these critical dependencies may provide less resilience than its diagram suggests.




Figure 2. Workload placement and multi-cloud pattern selection flow. Source: author conceptualization.

6. Data Sovereignty and Saudi Regulatory Alignment

Saudi cloud architecture operates within a rapidly maturing policy environment. The Cloud First Policy was established to accelerate cloud adoption and guide government entities toward cloud-based investment. DGA guidance and digital transformation standards require formal cloud adoption planning and cloud architecture capabilities. CST's Cloud Computing Service Provisioning Regulations version 4, effective from October 2023, define obligations for cloud service provisioning. NCA's updated Cloud Cybersecurity Controls (CCC-2:2024) establish minimum cloud cybersecurity expectations and explicitly reflect data-localization considerations. SDAIA and NDMO provide the national data-governance context, and the PDPL regulates personal-data processing and transfer (MCIT, 2020; CST, 2023; DGA, 2024; NCA, 2024; SDAIA, 2023a).

For enterprise architects, the practical lesson is to build a regulatory control matrix rather than rely on general statements such as "Saudi data must stay in Saudi Arabia." The matrix should ask: What is the data class? Does it contain personal data? Is the entity subject to sector-specific rules? Is the cloud service registered or otherwise acceptable for the intended use? Where are primary data, replicas, logs, backups, and support data stored? What entities can access them? Does transfer outside the Kingdom occur, and if so, under what basis and safeguards? What cybersecurity controls apply? Who approves exceptions?

Table 2. Data sovereignty and regulatory control matrix for multi-cloud databases.

Control area	Architecture question	Recommended control	Evidence / artefact
Data classification	What data is stored, replicated, logged or backed up?	Assign owner, classification, personal-data status, criticality and permitted locations	Data inventory; classification register; data-flow map
Residency / transfer	Will any copy or telemetry move outside an approved Saudi location?	Region allow-lists, documented transfer assessment where applicable, provider/location approval	Architecture diagram; provider region configuration; transfer record

Control area	Architecture question	Recommended control	Evidence / artefact
Encryption & keys	Who controls encryption and privileged access?	Encrypt at rest/in transit; defined key ownership, rotation and emergency access	KMS/HSM policy; key audit logs; privileged-access records
Backups & archives	Where are backups, snapshots and long-term archives retained?	Apply the same classification and location rules to recovery data	Backup policy; retention settings; restore-test evidence
Operational telemetry	Do SQL text, logs, plans or diagnostics expose sensitive data?	Filter, mask or regionalize telemetry; least-privilege access to observability data	Logging design; masking rules; SIEM access controls
Provider exit	Can data and metadata be recovered outside the primary provider?	Maintain documented export/restore methods and a tested alternate recovery path	Exit runbook; recovery exercise report; data-format inventory

Source: author synthesis based on Saudi PDPL/transfer requirements, NDMO governance, NCA controls, and DGA cloud guidance.

Table 2 converts these questions into architectural controls. A compliant design should be able to produce evidence for each answer. This evidence can include architecture diagrams, data-flow maps, provider contracts, region configurations, encryption settings, access logs, transfer assessments, backup policies, and tested recovery records.

Sovereignty should also include exit capability. If an enterprise cannot recover its database in an alternate environment, export data at operational scale, or reconstruct security and metadata, its practical control over the data is reduced even if the primary copy is locally hosted. Multi-cloud architecture can therefore be viewed as a sovereignty instrument when it creates credible technical alternatives and documented operational independence.

7. Scalability without Uncontrolled Complexity

Vertical and horizontal scaling within the primary database environment should remain the first option for many transactional workloads. Read-intensive services can use local read replicas. Analytical workloads can consume controlled replicated datasets through change-data-capture. Batch or AI workloads can be moved to the provider with the best computational capability if the data is approved for that environment. Data products can be published with classification metadata so that consuming platforms know what can be replicated.

Capacity planning should include network limits as well as database resources. Replication throughput must be sufficient to prevent RPO violations during peak periods. Recovery environments must be able to absorb the production workload after failover. Backup restoration must be timed with realistic data volumes. These measures are particularly important for large enterprise databases where the time required to move terabytes of data can exceed acceptable recovery windows.

8. Digital Resilience and Cyber-Recovery Model

Digital resilience under Vision 2030 requires continuity not only during hardware failures but also during cyber incidents and large-scale service disruption. The proposed model defines resilience as four capabilities: **resist**, **absorb**, **recover**, and **adapt**. Resist includes security hardening and redundancy. Absorb includes graceful degradation, read-only modes, queueing, and local failover. Recover includes alternate-cloud restoration or activation. Adapt includes post-incident learning and architecture improvement.

The multi-cloud database design should maintain at least one recovery path that does not depend entirely on the same administrative trust domain as production. For example, backup copies intended for ransomware recovery should not be deletable with the same credentials used to administer the production database. Secondary-cloud accounts should have separately controlled emergency access. Recovery automation should be versioned and tested but protected from unauthorized changes.

Table 3. Tiered digital resilience model for database workloads.

Tier	Business characteristic	Recovery architecture	Operational expectation	Validation
Tier 1 – Critical	Severe business/public-service impact from outage or data loss	Automated/orchestrated failover, near-real-time replica plus isolated recovery copy	Dedicated ownership, high observability, controlled emergency access	Frequent full-service failover and cyber-recovery exercises
Tier 2 – Important	Material impact but short controlled interruption is tolerable	Warm standby or asynchronous cross-cloud replica	Documented failover, capacity reserved or rapidly provisioned	Scheduled recovery exercises and restore verification
Tier 3 – Standard	Moderate impact; restore-based recovery acceptable	Infrastructure-as-code plus protected backups in approved alternate environment	Clear rebuild sequence and dependency documentation	Periodic restore test with application validation
Tier 4 – Non-production	Low business impact; data can be regenerated or sanitized	Rebuild from code and approved data sources	Cost-optimized; no unnecessary multi-cloud duplication	Configuration rebuild test and data-governance review

Source: author conceptualization; recovery principles informed by HADR literature and NCA resilience requirements.

Table 3 presents a tiered resilience model. Tier 1 services require stringent RTO/RPO and frequent exercises. Tier 2 services prioritize controlled warm recovery. Tier 3 services may use restore-based recovery. Tier 4 non-critical environments can be rebuilt from code and sanitized data. The actual numeric objectives should be defined by each enterprise's business impact analysis rather than copied from generic benchmarks.

9. Implementation Roadmap

A controlled transition to multi-cloud should be phased. Phase 1 establishes governance and inventory. Enterprises identify critical databases, data classes, current provider dependencies, recovery objectives, and applicable regulatory requirements. Existing backups and disaster-recovery plans should be tested before introducing more complexity.

Phase 2 builds foundational portability. Teams standardize schema documentation, backup formats, infrastructure-as-code, secret management, monitoring, and database build procedures. The objective is to ensure that critical systems can be reconstructed and data can be restored outside the primary operational path.

Phase 3 introduces targeted secondary-cloud recovery for selected critical systems. Replication or backup transfer is enabled according to sovereignty requirements. Failover routing, identity, and operational runbooks are tested. At this stage, the enterprise should avoid scaling the program faster than its ability to validate recovery.

Phase 4 introduces advanced patterns such as active-active, distributed read scaling, or multi-cloud analytics only where justified by measurable requirements. Centralized policy, security

telemetry, data-governance controls, and FinOps are expanded. Provider-specific features remain allowed but must be classified and governed.

Phase 5 institutionalizes continuous resilience testing. Recovery exercises, provider-exit simulations, configuration-drift checks, data-location audits, cost reviews, and cybersecurity assessments become recurring governance activities. Table 4 summarizes the roadmap and exit criteria.

Table 4. Phased multi-cloud database implementation roadmap.

Phase	Objective	Key deliverables	Exit criterion
1. Govern & inventory	Understand risk before adding cloud complexity	Database inventory, data classification, dependency map, RTO/RPO, regulatory matrix	Critical workloads have approved owners, classifications and recovery objectives
2. Build portability	Make reconstruction and data recovery repeatable	Schema/dependency documentation, backup standards, IaC, secrets, monitoring, restore automation	A selected workload can be rebuilt and restored outside its normal operating path
3. Add targeted recovery	Reduce concentration risk for highest-priority systems	Secondary-cloud account, replication/backup transfer, routing, runbooks, independent access	Controlled exercise meets approved RTO/RPO and security checks
4. Optimize patterns	Use advanced multi-cloud only where requirements justify it	Active-active/read scaling/analytics patterns, policy-as-code, unified observability, FinOps	Benefits and risks are measured; provider-specific dependencies are documented
5. Institutionalize resilience	Make recovery and sovereignty continuous controls	Recurring drills, exit simulation, configuration-drift checks, data-location audit, cost review	Evidence is produced continuously and exceptions are tracked to closure

Source: author conceptualization based on multi-cloud lifecycle research and Saudi cloud governance guidance.

Each phase should have a measurable exit condition rather than a calendar deadline. For example, a recovery environment should not be declared production-ready until a controlled exercise has demonstrated database restoration or failover, application reconnection, security validation, and evidence that the resulting data state is within the approved RPO. This staged approach limits the risk of creating a broad multi-cloud footprint before the organization has the operational capability to manage it.

10. Discussion

The proposed framework suggests that the strongest case for multi-cloud databases in Saudi Arabia is not simply "more clouds equal more reliability." The value comes from controlled architectural optionality. An enterprise that can classify data, place it deliberately, operate independent recovery, and move or restore critical state has greater strategic control than one that merely consumes services from several providers.

A particularly important finding from the literature is that stateful portability remains harder than stateless portability. Multi-cloud-native research continues to identify state synchronization and cloud-agnostic data design as open challenges (Alonso et al., 2023). Database vendor-lock-in research similarly shows that data size, bandwidth, and service dependencies can make migration difficult (Shaik and Kalyanasundaram, 2024). Saudi enterprises should therefore prioritize portability at the data and recovery layers even when production services remain provider-specific.

The Saudi regulatory environment also supports a mature interpretation of cloud governance. Cloud adoption is encouraged, but it exists alongside cybersecurity controls, data governance,

personal-data protection, and formal digital transformation standards. The resulting architectural objective is neither cloud avoidance nor unrestricted cloud expansion. It is governed cloud adoption that preserves national and organizational control while using cloud capabilities to improve service quality and innovation.

11. Limitations and Future Research

This paper is a conceptual review and framework. It does not benchmark specific commercial cloud providers, and it does not claim that one multi-cloud pattern is universally superior. Costs, latency, service availability, database features, and regulatory obligations vary by sector and workload. The framework should therefore be validated through organization-specific pilots and recovery exercises.

Future research could evaluate the framework empirically using representative Saudi enterprise workloads. Useful experiments would compare cross-cloud replication latency, recovery time, failover reliability, and operating cost for Oracle, PostgreSQL, SQL Server, and distributed SQL systems. Further work could also develop formal scoring models for sovereignty risk, provider concentration, and portability debt. Another research direction is AI-assisted multi-cloud operations, where machine learning predicts capacity, replication lag, failure risk, and cost while respecting data-location and security policies.

Confidential computing and privacy-enhancing technologies may also expand the range of workloads that can operate across cloud boundaries, but performance and operational complexity remain important considerations. Finally, future work should examine how sovereign cloud initiatives, evolving data-transfer rules, and Saudi digital infrastructure policies affect multi-cloud database design over time.

12. Conclusion

Saudi Arabia's digital transformation creates a strong need for database architectures that can scale without sacrificing control or resilience. Multi-cloud architecture can contribute to this goal, but only when designed as a governed stateful-data strategy rather than as simple provider diversification. The paper proposed a six-domain framework covering workload classification, sovereign placement, portability, resilience, unified security and observability, and lifecycle governance.

The framework emphasizes four principles. First, data placement must follow classification, lawful transfer conditions, cybersecurity requirements, and business purpose. Second, database portability must include data, schema, tooling, operations, and recovery—not merely the ability to create a virtual machine in another cloud. Third, resilience must be measured through tested RTO, RPO, failover, and cyber-recovery outcomes. Fourth, multi-cloud adoption must be selective: active-active designs should be reserved for workloads that justify their complexity, while active-passive, segmented, or portable-recovery models can provide substantial resilience with lower risk.

Under Vision 2030, the combination of cloud-first adoption, national data governance, cybersecurity regulation, and investment in digital infrastructure creates a strong foundation for this approach. A well-designed multi-cloud database strategy can help Saudi enterprises achieve scalable digital services, stronger data sovereignty, reduced concentration risk, and credible

continuity during technical or cyber disruption. The most valuable outcome is not the number of cloud providers used, but the enterprise's demonstrated ability to control, protect, recover, and evolve its critical data platforms.

References

- Alonso, J., Orue-Echevarria, L., Casola, V., Torre, A.I., Huarte, M., Osaba, E. et al. (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>.
- Tomarchio, O., Calcaterra, 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>.
- Ramalingam, C. and Mohan, P. (2021) 'Addressing Semantics Standards for Cloud Portability and Interoperability in Multi Cloud Environment', *Symmetry*, 13(2), 317. <https://doi.org/10.3390/sym13020317>.
- Galij, S., Pawlak, G. and Grzyb, S. (2024) 'Modeling Data Sovereignty in Public Cloud—A Comparison of Existing Solutions', *Applied Sciences*, 14(23), 10803. <https://doi.org/10.3390/app142310803>.
- Plazotta, M. and Klettke, M. (2024) 'Data Architectures in Cloud Environments', *Datenbank-Spektrum*, 24, pp. 243–247. <https://doi.org/10.1007/s13222-024-00490-5>.
- Dong, H., Zhang, C., Li, G. and Zhang, H. (2024) 'Cloud-Native Databases: A Survey', *IEEE Transactions on Knowledge and Data Engineering*, 36(12), pp. 7772–7791. <https://doi.org/10.1109/TKDE.2024.3397508>.
- Wang, L., Yang, Z. and Song, X. (2020) 'SHAMC: A Secure and highly available database system in multi-cloud environment', *Future Generation Computer Systems*, 105, pp. 873–883. <https://doi.org/10.1016/j.future.2017.07.011>.
- Shaik, V. and Kalyanasundaram, N. (2024) 'Cloud databases: A resilient and robust framework to dissolve vendor lock-in', *Software Impacts*, 21, 100680. <https://doi.org/10.1016/j.simpa.2024.100680>.
- Xiong, H., Fowley, F. and Pahl, C. (2016) 'A database-specific pattern for multi-cloud high availability and disaster recovery', in Celesti, A. and Leitner, P. (eds.) *Advances in Service-Oriented and Cloud Computing. Communications in Computer and Information Science*, 567, pp. 374–388. https://doi.org/10.1007/978-3-319-33313-7_29.
- Ullah, F., Dhingra, S., Xia, X. and Babar, M.A. (2024) 'Evaluation of distributed data processing frameworks in hybrid clouds', *Journal of Network and Computer Applications*, 224, 103837. <https://doi.org/10.1016/j.jnca.2024.103837>.
- Goswami, P., Faujdar, N., Debnath, S., Khan, A.K. and Singh, G. (2024) 'Investigation on storage level data integrity strategies in cloud computing: classification, security obstructions, challenges and vulnerability', *Journal of Cloud Computing*, 13, 45. <https://doi.org/10.1186/s13677-024-00605-z>.
- Chandramouli, R. and Hales, W. (2024) *A Data Protection Approach for Cloud-Native Applications*. NISTIR 8505. Gaithersburg, MD: National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.IR.8505>.
- Yeboah-Ofori, A., Jafar, A., Toluwalaju, T., Hilton, I., Oseni, W. and Musa, A. (2024) 'Data Security and Governance in Multi-Cloud Computing Environment', 2024 11th International Conference on Future Internet of Things and Cloud (FiCloud), Vienna, Austria. <https://doi.org/10.1109/FiCloud62933.2024.00040>.
- Diningrat, D.C., Suhardi and Rahardjo, B. (2025) 'Security Issues in Multi-Cloud: A Systematic Literature Review', *IEEE Access*, 13, pp. 70006–70017. <https://doi.org/10.1109/ACCESS.2025.3561352>.
- Ali, S., Talpur, D.B., Abro, A., Alshudukhi, K.S., Alwakid, G.N., Humayun, M., Bashir, F., Wadho, S.A. and Shah, A. (2025) 'Security and privacy in multi-cloud and hybrid cloud environments: Challenges, strategies, and future directions', *Computers & Security*, 157, 104599. <https://doi.org/10.1016/j.cose.2025.104599>.
- Ameur, Y., Bouzeffrane, S. and Thinh, L.V. (2023) 'Handling security issues by using homomorphic encryption in multi-cloud environment', *Procedia Computer Science*, 220, pp. 390–397. <https://doi.org/10.1016/j.procs.2023.03.050>.
- Nabawy, R.M., Hassanin, M., Ibrahim, M.H. and Kaseb, M.R. (2024) 'Mixed fragmentation technique for securing structured data using multi-cloud environment (MFT-SSD)', *Ad Hoc Networks*, 164, 103625. <https://doi.org/10.1016/j.adhoc.2024.103625>.
- Ministry of Communications and Information Technology (MCIT) (2020) *KSA Cloud First Policy*. Riyadh: MCIT. Available at: https://www.mcit.gov.sa/sites/default/files/ksa_cloud_first_policy_en.pdf (Accessed: 3 September 2026).

533 Digital Government Authority (DGA) (2023) Guideline for Cloud Computing Adoption by Government Agencies –
534 Version 1.0. Riyadh: DGA. Available at: <https://dga.gov.sa/> (Accessed: 3 September 2026).

535 Digital Government Authority (DGA) (2024a) Cloud Computing and its Role in Accelerating Digital
536 Transformation and its Sustainable Impact in the Government Sector – Version 1.0. Riyadh: DGA. Available at:
537 <https://dga.gov.sa/> (Accessed: 3 September 2026).

538 Digital Government Authority (DGA) (2024b) Digital Transformation Basic Standards – Version 3.0. Riyadh:
539 DGA. Available at: <https://dga.gov.sa/> (Accessed: 3 September 2026).

540 Communications, Space and Technology Commission (CST) (2023) Cloud Computing Service Provisioning
541 Regulations, Version 4. Riyadh: CST. Available at: <https://www.cst.gov.sa/> (Accessed: 3 September 2026).

542 National Cybersecurity Authority (NCA) (2022) Data Cybersecurity Controls (DCC-1:2022). Riyadh: NCA.
543 Available at: <https://nca.gov.sa/en/regulatory-documents/controls-list/dcc/> (Accessed: 3 September 2026).

544 National Cybersecurity Authority (NCA) (2024a) Cloud Cybersecurity Controls (CCC-2:2024). Riyadh: NCA.
545 Available at: <https://nca.gov.sa/en/regulatory-documents/controls-list/ccc/> (Accessed: 3 September 2026).

546 National Cybersecurity Authority (NCA) (2024b) Essential Cybersecurity Controls (ECC 2-2024). Riyadh: NCA.
547 Available at: <https://nca.gov.sa/en/regulatory-documents/controls-list/ecc/> (Accessed: 3 September 2026).

548 Saudi Data & AI Authority (SDAIA) (2023a) Personal Data Protection Law. Riyadh: SDAIA. Available at:
549 <https://sdaia.gov.sa/> (Accessed: 3 September 2026).

550 Saudi Data & AI Authority (SDAIA) (2023b) Regulation on Personal Data Transfer Outside the Kingdom. Riyadh:
551 SDAIA. Available at: <https://dgp.sdaia.gov.sa/> (Accessed: 3 September 2026).

552 Saudi Data & AI Authority (SDAIA) (n.d.-a) National Data Management Office (NDMO). Available at:
553 <https://sdaia.gov.sa/en/Sectors/Ndmo/Pages/default.aspx> (Accessed: 3 September 2026).

554 Saudi Data & AI Authority (SDAIA) (n.d.-b) National Strategy for Data & AI. Available at:
555 <https://sdaia.gov.sa/en/SDAIA/SdaiaStrategies/Pages/NationalStrategyForDataAndAI.aspx> (Accessed: 3
556 September 2026).

557 Saudi Data & AI Authority (SDAIA) (n.d.-c) National Information Center (NIC) and Government Cloud Services.
558 Available at: <https://sdaia.gov.sa/en/Sectors/Nic/Pages/default.aspx> (Accessed: 3 September 2026).