Global Academic Journal of Economics and Business

Available online at https://www.gajrc.com
DOI: https://doi.org/10.36348/gajeb.2026.v08i06.014

OPEN ACCESS



ISSN: 2706-9001 (P)
ISSN: 2707-2584 (O)

Original Research Article

Strengthening Cyber Resilience in Saudi Arabia's Cloud Infrastructure: An Integrated Framework for Threat Management, Security Governance, and Business Continuity

Ataur Rahman Atif Mohammed[1]*
[1]Independent Researcher

***Corresponding Author**
Ataur Rahman Atif Mohammed
Independent Researcher

Article History
Received: 09.09.2026
Accepted: 02.11.2026
Published: 03.11.2026

Abstract: Cloud platforms are now underpinning key Saudi public services, financial services and industrial ecosystems. The very nature of the cloud's architecture can however amplify the risk of identity compromise, configuration errors and dependency failures. Much of the existing literature focuses on threat detection, compliance and disaster recovery in isolation with little guidance on how to integrate these into a coherent set of accountable operating decisions. This critical integrative review of the literature from 2020 to 2025 develops a Saudi-specific framework for integrating anticipatory intelligence, preventative controls, incident command and continuity assurance. It uses a structured thematic synthesis to examine the peer-reviewed literature on cloud attack surfaces, zero trust architectures, organisational capabilities, response learning and recovery governance, and this is augmented by an analysis of the specific Saudi cybersecurity requirements. The review concludes that resilience is not so much about the sheer number of controls that are in place, as it is about the demonstrable ownership of security decisions across the cloud provider, the tenant, the supplier and the various business functions. Its primary contribution is an integrated, testable governance framework in which the level of service criticality determines the priority of controls, the level of incident escalation and the requirement for independently demonstrated restoration capability. The recommended indicators are detection and containment latency, privileged access exposure, backup integrity, recovery point achievement and critical service restoration. The framework is a conceptual synthesis and not an empirically validated intervention. It would be necessary to test it in practice across the various Saudi sectors before any claims of effectiveness could be made.

Keywords: Cloud Cybersecurity, Saudi Arabia, Cyber Resilience, Risk Governance, Threat Intelligence, Business Continuity, Recovery Assurance.

Copyright © 2026 The Author(s): This is an open-access article distributed under the terms of the Creative Commons Attribution **4.0 International License (CC BY-NC 4.0)** which permits unrestricted use, distribution, and reproduction in any medium for non-commercial use provided the original author and source are credited.

1 INTRODUCTION

Saudi Arabia's digital economy transformation programme is increasingly dependent on cloud-hosted government platforms, enterprise systems, data services and the associated infrastructure. The benefits of cloud adoption in terms of speed of deployment and elasticity are real, but this comes at the cost of creating new dependencies on identity providers, application programming interfaces, outsourced operations and

the global technology supply chain. Security is therefore not simply a matter of securing the server. You have to be able to maintain legitimate services and you have to be able to answer the question of who is responsible in the event of an incident, whose evidence is going to be valid and what can be restored in the event of a failure across the interconnected stack. Research into government cloud adoption has shown that uncertainty around the protection of information continues to be a major impediment. Service layer reviews have also shown that the lines of responsibility are different depending on whether you're looking at infrastructure, platform or software [1, 2].

Cyber resilience is different from just prevention. It's about anticipation, resistance, absorption, recovery and adaptation all the while maintaining an acceptable level of business function [3]. A very mature organisation may in fact prevent the vast majority of the intrusions that are attempted against it, yet still be quite fragile if a privileged account can disable the backups or if the recovery scripts are reliant on the same compromised management plane. On the other hand you can have redundant infrastructure which will maintain availability even if the data is corrupted, with the corrupted data simply being replicated across the regions. The importance of this distinction is particularly relevant when you have cloud workloads underpinning public health, payment processing, transport or energy related administration where a brief interruption can have consequences which go well beyond the direct cost of the technology.

Evidence from the Saudi business sector points to the institutional nature of the problem. Comparative governance research has identified the difficulty of translating policy into actual awareness at the firm level [4]. A separate study from Saudi Arabia found a link between cybersecurity systems and sustainable organisational performance but did not find the proposed mediation and moderation effects [5]. These findings serve as a warning against assuming that a positive policy, a good security culture or a resilient organisation will necessarily lead to a robust recovery. They also underline the need to look at the technical and managerial evidence in tandem rather than to simply state that investment in one will automatically result in improvement in another.

This article is about the resulting integration issue: how can the Saudi cloud dependent organisations integrate their threat management, their regulatory accountability and their business continuity into a defensible decision architecture? It makes no claim to have interviewed any new participants or to have tested a platform in operation. Its value is in the critical synthesis of the available evidence, in identifying the conflicting assumptions and in specifying an architecture which future empirical work can go on to test.

2 Review Aim, Questions and Analytical Boundaries

The objective is to develop an operationally meaningful model for enhancing cyber resilience across the Saudi cloud infrastructure with a specific focus on service continuity in the face of credible adversarial disruption. The review is structured around four objectives: to identify attack vectors which would circumvent individual security controls; to assess governance arrangements which can confer clear ownership; to compare recovery arrangements in terms of business criticality; and to establish quantifiable relationships between threat indicators, decisions and the resulting restoration.

Three research questions: Firstly, what characteristics of cloud threats render traditional perimeter defences inadequate? Secondly, how should Saudi institutions balance the responsibilities of the provider, the tenant and the state without either duplicating or dropping essential controls? Thirdly, what recovery measures actually demonstrate the continued delivery of critical business functions rather than simply the existence of duplicated resources? The unit of analysis is the cloud-dependent organisational service (including its associated identities, data, suppliers and recovery dependencies). The review covers public, private, hybrid and multi-cloud environments but does not presume that the risk profiles of these are the same.

3 REVIEW METHODOLOGY

An integrative literature review was the appropriate methodology given that the research question spans technical engineering, organisational governance and continuity management. This is in contrast to a more narrow quantitative meta-analysis which would be unable to effectively deal with the range of evidence. The integrative approach allows for the comparison of very heterogeneous evidence (security taxonomies, observational studies, conceptual models and qualitative organisational studies). The evidence window was restricted to publications from January 2020 to December 2025. The search concepts combined cloud computing or cloud-native infrastructure with cyber resilience, threat intelligence, ransomware, zero trust, risk governance, disaster recovery, business continuity, third party security and Saudi Arabia. The searches were focused on scholarly publishers, research indexes and the cited literature with the official Saudi regulatory publications examined separately in order to provide some context as to the interpretation of the regulatory environment.

Eligibility required a direct link to either cloud-specific attack exposure, cyber-resilience performance, institutional security practice or recovery accountability. General awareness studies were only included if they related to workforce or governance. Conceptual recommendations were clearly separated from empirical findings and studies from outside of Saudi Arabia were treated as transferable hypotheses rather than as evidence of effectiveness in the local context. The three papers provided the basis for structure and context: the two Saudi studies formed the basis of the governance and performance assessment [4, 5], and the later regional review was only used to inform the presentation and was not part of the evidence base.

The synthesis proceeded in five stages. Firstly, the candidate studies were screened for relevance, date of publication and availability of bibliographic details. Secondly, the claims were coded under exposure, preventive control, detection, incident coordination, restoration, regulatory ownership and organisational learning. Thirdly, each claim was categorised as descriptive evidence, reported empirical association, design proposal or normative requirement. Fourthly, comparisons were made to identify agreement, contradiction and any missing operational interfaces. Finally, the relationships which were supported by converging arguments were expressed as control objectives and testable indicators. There are no fabricated screening totals, no pooled effect sizes and no statistically representative Saudi estimates. This is an analytically structured review rather than a complete census of all relevant publications.

Methodological limitations were taken into account during the synthesis. Vendor managed systems are not typically studied under the same threat assumptions; detection accuracy on a lab dataset is not necessarily indicative of continuity benefits; and self reported compliance is no guarantee of restored functionality in the event of an incident. Technical papers will often describe the threat without stating what the contractual decision rights are, whilst management papers can completely omit the architecture that leads to the operational loss. Confidence in any proposed linkage was therefore a function of the clarity of the mechanism, the appropriateness of the context and the consistency with the other evidence rather than simply the number of publications.

4 Cloud Threat Pathways and Defensive Limits

Cloud exposure starts at the boundaries that are often hidden away in a standard network diagram. Identity federation, service accounts, automation pipelines and admin interfaces all allow for legitimate activity across a range of environments. The same would be true of a compromise of credentials though, with the attacker potentially gaining access to the same breadth of opportunity. Security awareness in the Kingdom is still very much an issue. A large survey was carried out which showed poor authentication practices and a lack of phishing awareness in the respondents (though it should be noted they are not a representative sample of cloud admins [6]. Organisational readiness research has also demonstrated that technical capability is a function of management, environment and structure [7]. Access compromise therefore is as much a matter of engineering controls as it is of decisions around training, exception approvals and ownership.

Zero trust challenges the very notion that being inside the perimeter equates to being trusted. The alternative to static perimeter defences is described as continuous authorisation, least privilege and context aware validation [9–11]. It's important to stress though that zero trust is not a product that eliminates the need for recovery. A compromised privileged identity could still authorise a destructive operation if the authorisation policy, monitoring and human escalation all share the same unmitigated failure mode. The implementation of zero trust also brings with it a number of costs in the form of identity inventory, application compatibility, service discovery and logging. The right test is whether the design reduces the adversary's effective permissions and allows for timely revocation without in any way disabling legitimate emergency access.

Misconfiguration and insecure interfaces constitute a second category of exposure. Excessively permissive object storage, secret leakage, exposed management interfaces and inconsistent encryption key ownership are all examples of risks which can completely undermine otherwise robust workloads. Cloud security surveys have documented these risks across the various service models and it is precisely for this reason that the tenant's responsibilities cannot be deduced from the provider's general certification [2-13]. Continuous configuration assessment should therefore be tied into an authoritative inventory of service owners, classifications and deployment pipelines. Automated remediation must also have in place controls to prevent unsafe bulk changes as a bad policy pushed out simultaneously to production and standby can of course escalate the scale of the outage.

A third route is lateral dependency propagation. Malware or malicious updates introduced via a trusted software component can impact multiple services even in the absence of any direct external access. Work on supply chain attacks has shown that established trust relationships can be

very effective as a distribution mechanism for malicious activity [23]. In a cloud environment dependencies can extend out to package repositories, managed identity services, observability providers and infrastructure templates. Signing of artefacts, isolation of build credentials and verification of software provenance all serve to improve assurance but no one technical control is going to eliminate the concentration risk that exists when a number of ostensibly separate services all rely upon a single provider or a single administrative function.

Ransomware is a perfect example of the difference between compromise detection and recovery design. The literature is full of accounts of evolving encryption techniques, operational extortion and the sheer difficulty of distinguishing malicious mass modification from legitimate activity [18–20]. Good controls will separate production authority from backup deletion authority, will maintain version history, will protect the credentials required to restore and will ensure that the data restored is predated by the adversary's persistence. Testing that a backup job has completed is a complete non-starter as it provides a false sense of security. A successful copy may well contain compromised files, unusable encryption keys or dependencies which can't be recreated. Recovery design must therefore be built on the assumption that a malicious insider or an administrator account has been compromised, not simply that a piece of hardware has failed.

Detection tools present another trade-off. Behavioural analytics and threat intel can help to reduce the time it takes to detect an unfamiliar attack pattern, but telemetry can contain PII, the number of false positives can be overwhelming for the responder and model drift can lead to reduced reliability. Intelligence sharing research has identified legal, interoperability and data quality issues [14]. Analysis of the overlapping attack surface of the cloud, the edge and the connected device also suggests that signals need to be contextualised by asset criticality [15]. A real security operations centre would be looking to correlate identity events with workload configuration changes, data access anomalies and business service health rather than just looking at the number of alerts.

Table 1 separates out the principal threat mechanisms from the decisions required to stop them. It doesn't rank the threats by some unverified Saudi incidence figures (that would be meaningless anyway). The purpose is to establish causal traceability. An indicator becomes actionable only when you know the corresponding asset, the owner of that asset, the decision threshold and the authorised containment procedure. This is especially pertinent in the case of automated isolation where an indiscriminate response could itself be the thing that takes down a critical service.

Table 1: Cloud threat pathways and decision-linked assurance requirements

Attack pathway	Operational exposure	Required assurance decision
Privileged identity compromise	Federated access, broad administrative roles	Revoke sessions; isolate management authority; audit access
Cloud configuration drift	Exposed data, inconsistent policies	Verify asset ownership; restrict and safely roll back changes
Supply-chain compromise	Trusted software or pipeline carries malicious changes	Quarantine releases; establish artifact provenance
Ransomware and data corruption	Encryption, destructive changes, contaminated backups	Validate immutable copies; restore from trusted recovery authority
Provider or dependency outage	Shared control-plane or integration failure	Invoke independently tested alternate service or degraded mode

5 Governance and Regulatory Accountability

Governance is the mechanism by which security expectations become binding decisions. Evidence from Saudi Arabia and from a number of other business contexts shows the gap that can exist between the written controls, the level of awareness among stakeholders and actual practice [4–8]. National strategy work has highlighted the need to link capability development with clear outcomes in terms of competencies [8]. You simply can't hand cyber resilience over to an infrastructure provider or a security team. The governance body has to decide which services are essential, who is going to accept the residual risk and who is going to be the one to step in if a failure occurs. Without those decisions, the various departments will tend to focus on optimising their own controls and will leave the cross-functional dependencies entirely unowned.

Saudi regulation provides the necessary context. The Essential Cybersecurity Controls of the National Cybersecurity Authority set out the baseline for those within scope, and the Cloud Cybersecurity Controls provide a clear distinction between cloud service providers and cloud service tenants. The 2024 version of the Cloud Cybersecurity Controls was issued as part of the Authority's updated control programme and specifically deals with the cloud

tenant's responsibilities. For each service you're looking at you'll need to determine applicability, data location and any sector specific obligations against the current set of authoritative requirements. These are the controls we're referring to as the regulatory context in this instance and they are not one of the thirty DOI bearing academic papers. A legal assessment is still required where personal data, critical systems or cross border activity would give rise to additional obligations.

The key governance issue is not whether or not a provider publishes a security statement, but rather who has the authority and the evidence to act during a disruption. Stamenkov's analysis of the various cloud service models demonstrates that the responsibilities for business continuity and recovery are very much service and recovery model dependent [16]. With infrastructure as a service the tenant may well have a significant amount of responsibility for their own workload recovery. With managed software the provider may run the recovery infrastructure and the tenant will own the business workarounds, the communications and the acceptance of the restored records. The contract should clearly define the responsibilities in respect of logging access, incident notification, forensic support, recovery responsibilities, any subcontractor dependencies and an enforceable exit clause.

Control ownership can be represented as an evidence chain. Each critical service should have a named accountable executive, a named technical operator, a provider contact, a data owner, a risk acceptance authority and a recovery approver. The service register should have the relevant policy clause linked to the configuration evidence and a dated test. That shifts the audit effort from 'does the policy exist' to 'has the action in the policy been taken'. Things like unreviewed privileged grants, overdue vulnerability remediation, failed recovery drills and missing provider attestations are far more diagnostic than some unqualified % referred to as compliance.

Governance also comes down to incentives and resources. Saudi evidence links organisational cybersecurity arrangements to outcomes without necessarily confirming all of the theorised causal pathways [5]. Research into SMEs also shows the existence of constraints in terms of people, confidence and prioritisation [27]. This suggests that a one size fits all technical prescription may be a burden on organisations of very different capacity. Proportionate implementation should ensure that the most critical services are protected by mandatory controls, while allowing smaller organisations to make use of shared monitoring, tested managed recovery and centralised expert support. Proportionality is about the means of assurance not about the permission to ignore known dependencies.

6 Business Continuity as Demonstrated Service Capability

Business continuity turns cyber threats into consequences for the users and the critical functions of the business. A business impact analysis should identify the service dependencies, the maximum tolerable outage, the permissible amount of data loss, the minimum viable operations and the point at which manual alternatives become unsafe. The recovery time objective is the target for restoration, the recovery point objective is the acceptable level of data loss. Neither of those in themselves is any kind of evidence that the target can be achieved. Reviews of cloud enabled resilience suggest that the use of managed capabilities can reduce the effort required to implement the solution, though the benefit will be workload and architect dependent [17].

Resilience architecture requires separation across failure domains. Replication within an environment will protect against some of the component faults, isolated backup identities and logically independent recovery environments will help in the case of deliberate compromise. Geographic redundancy may reduce the exposure to a regional outage but two regions sharing a central identity, a shared domain name or shared deployment automation are not truly independent. The right design is to map out the common mode failures first and then make your choice of active-active, warm standby, pilot-light or backup and restore. The cheapest option is not necessarily the most resilient as complexity in itself can slow down containment and recovery.

Recovery verification should go well beyond just proving that the infrastructure is available. A real test starts with an adversarial scenario, defines which control plane is deemed untrusted and then restores a representative business transaction with both integrity and the appropriate user access. The test has to cover the source of the encryption keys, the immutable recovery images, the network requirements, any third party integrations and the accuracy of the restored business data. A pass is only achieved when there is demonstrable evidence that the named service can deliver its minimum acceptable function within the agreed time and data loss parameters. The difference between infrastructure being up and a service being actually usable is of the utmost importance when decisions are being made on the basis of accurate data or the exchange of data with external parties.

A continuity plan must also deal with operational command. The security team will be

focused on containment and evidential preservation, the operations team on service restoration and those objectives can be in conflict. There needs to be clearly defined escalation criteria that will determine when an Executive Incident Commander authorises isolation, whether or not degraded operations can be tolerated and what data validation criteria will be required before the service can be re-opened to the public. Incident learning research has shown that many organisations do not actually learn from their incident investigations [21]. The review therefore treats after action correction, supplier renegotiation and regular scenario re-design as fundamental continuity processes rather than as optional activities which can be documented.

Forensic readiness underpins continuity decisions. Cloud investigations are constrained by the ephemeral nature of logs, shared infrastructure and evidence that is ultimately in the hands of the provider [22]. Practitioner evidence again points to the need to tailor forensic procedures to the operational requirements [29]. It is essential that organisations agree preservation procedures in advance of an incident (audit log retention, time sync, chain of custody, lawful access etc). The collection of evidence must be such that the need for rapid restoration does not result in the only records available being wiped out. The records in question will be the only records that can be used to identify compromised identities and to determine whether the data that has been recovered can be trusted.

7 Proposed Integrated Cloud Resilience Framework

The review sets out an Integrated Saudi Cloud Resilience Framework based on five interlocking functions: classify, constrain, sense, coordinate and restore. Classification is about determining the importance of a service, the level of acceptable disruption and the relevant legal constraints. Constraint management is about reducing the attacker's attack surface through identity segmentation, configuration baselines, secure supply chains and independent recovery authority. Sensing is the combination of threat indicators with business impact. Coordination is the process of taking the verified signal and turning it into a time-bound operational decision. Restoration is the process of proving that acceptable service can be restored with data intact and controls restored. The functions are deliberately linked via shared service identifiers rather than being run as separate committees or dashboards.

Figure 1 shows the operating architecture. At its heart is the critical service register which holds the dependencies, owners and thresholds. Threat telemetry feeds into a monitored decision layer but does not in any way trigger a destructive action. Governance provides the authorisation boundaries, the jurisdictional rules and the provider commitments. The recovery layer provides the actual options and the evidence as to whether a service can be recovered. Incident command is the layer that ties it all together, ensuring that containment decisions are made with regard to continuity impact and that restoration decisions are made with regard to the level of residual adversarial access. There's an audit and learning loop which returns the verified results back to the architecture and policy design.

A formal prioritisation rule is better than an opaque composite risk score. For each service the owning team should assess business criticality, plausible attack vectors, dependency concentration, control effectiveness and demonstrated recovery readiness. They can be placed on ordinal scales in the workshops but the numbers must not be presented as calibrated probabilities without the supporting data. A service with a moderate likelihood of compromise but with a catastrophic untested recovery is perfectly capable of being ranked above a service which is attacked fairly frequently but which is also very easy to recover from. The framework is all about having explicit decision rationales so that risk acceptance is something that can be reviewed by the appropriate level of management and regulators.

Table 2 sets out the framework in terms of deliverables. The indicators are deliberately split between leading indicators (e.g. excessive permissions, outdated response playbooks) and outcome indicators (e.g. verified restoration, actual incident duration). Organisations should establish a baseline of the indicators prior to setting targets as universal percentages would be meaningless across the sectors. Where an organisation is not prepared to provide the requisite evidence, the contract should require appropriate assurance or the organisation should record that as a material dependency risk. Reporting should be continuous and should expose the uncertainty, the test exclusions and the deferred corrective actions rather than simply hiding them away in an overall maturity score.

Implementation should happen in four controlled waves. Wave 1 is all about establishing a service and provider dependency inventory with a focus on the most critical workloads. Wave 2 is about enforcing privileged access, logging, configuration assurance and protected restore credentials. Wave 3 is the integration of incident exercises with business impact analysis, contractual notification and forensic readiness. Wave 4 is about institutionalising scenario based testing, periodic independent review and evidence led remediation. This is not some kind of universal maturity ladder. An organisation with very

real recovery weaknesses may well need to get to the point of testing and fixing backup integrity before it can start to deploy any form of advanced threat analytics. The order of investment should be in line with the level of exposure and consequence.

Figure 2 shows an assurance cycle in which intelligence, control checks and recovery activities all produce evidence for the subsequent governance decisions. The model assumes that recovery and threat containment are inextricably linked: if you simply keep on restoring without actually removing the adversary's persistence you're just recreating the original risk. The value of the recovery process therefore rests on a managed handover between the forensic team, the security response team, the cloud operators and the business owners. That handover must be exercised in both the malicious and accidental disruption scenarios as a control that is designed to work in one scenario may be completely fragile in another.

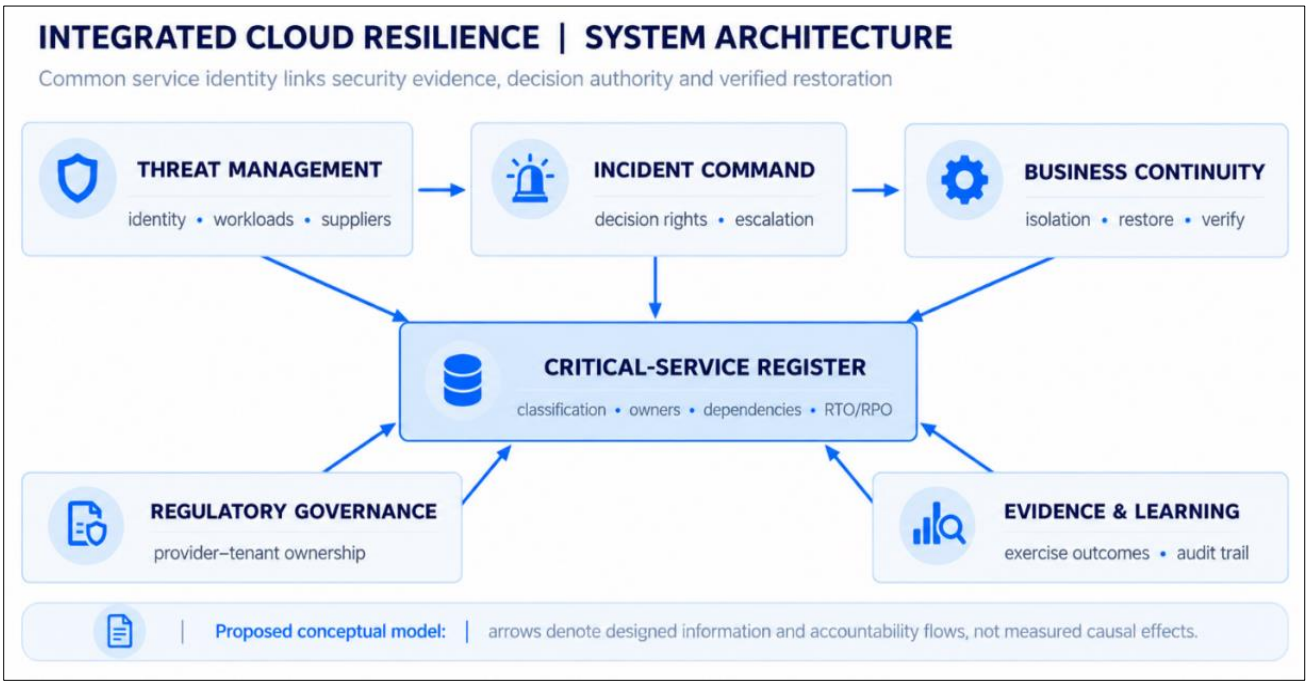


Figure 1: Integrated decision architecture for Saudi cloud-dependent services. Original conceptual representation



Figure 2: Closed-loop assurance sequence linking risk, response and tested business continuity. Original conceptual representation

Table 2: Accountability matrix and measurable cyber-resilience outcomes

Function	Accountable evidence	Candidate performance indicator
Classify	Signed impact analysis; service owner; dependency register	Share of priority services with current impact assessment
Constrain	Privileged-role review; configuration and backup policy	Open high-risk entitlements; failed baseline checks
Sense	Mapped detection rules; evidence provenance	Time from verified compromise to decision-ready alert
Coordinate	Incident commander; exercised provider escalation	Time to authorised containment; escalation completion
Restore	Independent recovery drill; integrity and function checks	RTO/RPO achieved; critical transactions successfully restored

8 Critical Discussion and Implications

The key synthesis is that you can't work out the security of a cloud service from the strength of the individual security controls that it has. The flexibility of the cloud creates correlated dependencies that persist across traditional organisational boundaries and organisational incentives can lead to the fragmentation of the authority required to manage those dependencies. Risk treatment research recommends explicit countermeasure selection and residual risk analysis [24] and critical infrastructure research has similarly stressed the importance of the relationship between technical assessment and governance [25]. The model proposed here adds another link in the chain: continuity outcomes become part of the normal risk acceptance process rather than being a separate exercise which takes place after the security controls have been selected.

There are still a number of tensions. Centralised monitoring gives you better visibility but can become an operational single point of failure. Automation reduces the response time but increases the impact of a bad policy. A multicloud architecture may spread your provider exposure but will also tend to multiply inconsistent identities and recovery dependencies. The research into zero trust implementations does support the idea of contextual verification, but doesn't prove that every architecture is going to give you better cost adjusted resilience [9–30]. The framework therefore does not endorse any vendor pattern. The decisions should instead be made on the basis of what failure modes each measure addresses and what new dependencies it introduces.

Measurement is another similar problem. The work on cyber-resilience metrics makes it clear that resilience must be about the system's behaviour in the face of a disturbance rather than just a list of pre-emptive checks [26]. In practice this would mean reporting the distribution of service restoration times across exercises, the proportion of scenarios in which the recovery objectives are achieved and the time taken to revoke the compromised permissions. The comparative interpretation of such metrics is still going to require consistent scenarios, consistent asset criticality and consistent reporting boundaries. High scores achieved in very easy simulations should not be compared to lower scores from more realistic exercises which might involve loss of identity infrastructure or the deliberate destruction of production data.

A methodological distinction to be made concerns the interpretation of association. The Saudi organisational study referenced did find associations between security arrangements and business performance but did not establish that each of the proposed resilience mechanisms actually mediated those outcomes [5]. The present synthesis therefore treats the proposed mechanisms as design propositions. This is a necessary restraint in order to avoid turning observed correlation, theoretical compatibility or control conformity into an unsupported assertion of causal efficacy. It is only through further field research that it will be possible to determine which of the elements actually do improve restoration performance and which simply result in more reporting.

An additional design requirement is around degraded operations. Some services are simply not going to be able to get back to full functionality within a reasonable period of a disruption. They may be able to achieve a socially useful outcome through a reduced transaction path though. The business owner will have to define in advance which functions can be turned off, which data can be queued and which decisions can be made by a human. This stops the technical team from treating a service that is only partially functional as being fully back. It also stops them from just re-attaching untrusted integrations to meet a headline uptime figure.

A second measurement issue is survivorship bias. The exercises which succeed are far more likely to be the ones that get presented to the execs than the ones which are inconclusive or which are simply abandoned. Assurance reports should record every exercise that was scheduled, its exclusions, its dependencies, its failures and its unresolved actions.

The reporting boundary must include any outsourced operations and any shared infrastructure which has a material impact on the protected service. Otherwise an apparently improving performance could just be the result of a narrowing of the test conditions. Governance needs to have a credible account of what has not been demonstrated because that residual uncertainty is itself a decision input.

These recommendations also place a requirement on tightly bounded information exchange. Threat indicators can be of great value in enabling a range of organisations to identify common adversary behaviour but indiscriminate sharing can result in the exposure of sensitive configurations or personal data. The information should be minimised, classified and disseminated under agreed authority with provenance and expiry information maintained. Indicators that cannot be validated against local telemetry should be used to inform investigation rather than automatically result in service isolation. This maintains the distinction between intelligence as evidence and intelligence as an instruction to act, a distinction which is absolutely critical for safely governed automation [14].

An issue with implementation is how do you verify that you've got a clean recovery state when the logs are incomplete? The restoration approval process should include verification of software integrity, identity key rotation, access policy review and reconciliation of the material business transactions. If full verification would breach the agreed outage threshold then management should authorise a precisely defined degraded mode rather than simply go ahead and waive the controls. The process must identify who is accepting the residual risk, how the uncertainty is communicated to the users and when the temporary exception ends. The restoration certificate should therefore be able to state that the infrastructure has been restored, that the information integrity has been verified and that the service functionality has been accepted. These are three separate assertions with different evidential requirements and different implications for public confidence.

There are significant implications for the Vision 2030 related transformation. Public bodies should treat assurance requirements as part of their procurement specification. Cloud providers should be able to produce evidence of segregation, breach notification and recovery support that is actually usable. Boards should fund joint exercises rather than separate security and continuity exercises. The workforce should be trained in incident decision making across the legal, operational and technical functions. The general reviews highlight the increasing complexity of the cyber threat landscape [28], but sustainable improvement will only be achieved through disciplined local validation rather than the assumption that the new technologies will automatically deliver the required level of risk reduction.

9 Limitations and Research Agenda

The review has three major limitations. The structured integrative search doesn't purport to be exhaustive in terms of database coverage and as such cannot be used to support any publication count or any assertion that all the studies in the review agree. A significant proportion of the relevant research is focused on organisations in countries other than Saudi Arabia so transferability is highly dependent on the sector, the contractual model, the language, the capability of the workforce and the degree of regulatory applicability. The framework itself has not been implemented or independently evaluated so the paper is offering a falsifiable design argument rather than any evidence of actual performance improvement.

Future work should assess the framework in government, health, finance and industrial service environments using matched disruption scenarios. Longitudinal work could look at detection to decision time, reduction in privileged access, business transaction restoration and recovery point achieved both pre and post governance integration. The research should note the failed tests and the evidence gaps that were imposed by the provider as much as the successful tests. Interviews with incident commanders and cloud tenants would be of value in explaining why a technically sound recovery capability can fail under the pressure of organisation.

Evaluation should also take into account service heterogeneity. The accuracy requirements, downstream systems and tolerable levels of disruption are all very different for a healthcare appointment system, a financial settlement system and a cloud based engineering repository. Shared metrics are only of any use for comparison in terms of governance if the definitions and test conditions are made explicit. In a future Saudi study the researchers could use the same incident scenarios and still have sector specific acceptance criteria, independent observers and the same level of documentary evidence. This would allow them to determine which intervention actually reduces recovery time and which simply shifts the workload or risk to another department. It would also allow them to determine whether contractual assurance is actually of any use when an outage affects a number of customers of the same cloud service provider.

10 CONCLUSION

Saudi cloud resilience requires a managed balance between adversarial risk, national governance and the practical ability to keep business services going. The literature indicates that cloud exposure is spread across identities, interfaces, suppliers and service dependencies and that formal security controls in themselves do not prove recovery. This review has developed a five function framework in which criticality, decision rights, evidence and tested restoration are all mutually accountable. The key proposition is that the outcome of the continuity exercises should lead to a change in the current security investment and policy posture rather than merely serve to validate a separate disaster recovery plan.

For practitioners the priorities are to identify the critical services, enforce the provider/tenant responsibility boundary, protect the independent recovery authority and test the usability of the restoration under credible hostile conditions. For researchers the next step is field validation within the Saudi organisational context with transparent metrics and documented failures. The framework is to be judged on the ability to recover and adapt the essential services in the face of disruption rather than on the number of controls that have been put in place or the number of certifications that have been obtained.

REFERENCES

1. Ukeje N, Gutierrez J, Petrova K (2024) Information security and privacy challenges of cloud computing for government adoption: a systematic review. International Journal of Information Security 23:1459–1475. https://doi.org/10.1007/s10207-023-00797-6
2. Parast FK, Sindhav C, Nikam S, Yekta HI, Kent KB, Hakak S (2022) Cloud computing security: a survey of service-based models. Computers & Security 114:102580. https://doi.org/10.1016/j.cose.2021.102580
3. Alhidaifi SM, Asghar MR, Ansari IS (2024) A survey on cyber resilience: key strategies, research challenges, and future directions. ACM Computing Surveys 56(8):196. https://doi.org/10.1145/3649218
4. Rawindaran N, Nawaf L, Alarifi S, Alghazzawi D, Carroll F, Katib I, Hewage C (2023) Enhancing cyber security governance and policy for SMEs in Industry 5.0: a comparative study between Saudi Arabia and the United Kingdom. Digital 3:200–231. https://doi.org/10.3390/digital3030014
5. Al-Somali SA, Saqr RR, Asiri AM, Al-Somali NA (2024) Organizational cybersecurity systems and sustainable business performance of small and medium enterprises (SMEs) in Saudi Arabia: the mediating and moderating role of cybersecurity resilience and organizational culture. Sustainability 16:1880. https://doi.org/10.3390/su16051880
6. Alzubaidi A (2021) Measuring the level of cyber-security awareness for cybercrime in Saudi Arabia. Heliyon 7:e06016. https://doi.org/10.1016/j.heliyon.2021.e06016
7. Hasan S, Ali M, Kurnia S, Thurasamy R (2021) Evaluating the cyber security readiness of organizations and its influence on performance. Journal of Information Security and Applications 58:102726. https://doi.org/10.1016/j.jisa.2020.102726
8. AlDaajeh S, Saleous H, Alrabaee S, Barka E, Breitinger F, Choo KKR (2022) The role of national cybersecurity strategies on the improvement of cybersecurity education. Computers & Security 119:102754. https://doi.org/10.1016/j.cose.2022.102754
9. Sarkar S, Choudhary G, Shandilya SK, Hussain A, Kim H (2022) Security of zero trust networks in cloud computing: a comparative review. Sustainability 14:11213. https://doi.org/10.3390/su141811213
10. Kang H, Liu G, Wang Q, Meng L, Liu J (2023) Theory and application of zero trust security: a brief survey. Entropy 25:1595. https://doi.org/10.3390/e25121595
11. He Y, Huang D, Chen L, Ni Y, Ma X (2022) A survey on zero trust architecture: challenges and future trends. Wireless Communications and Mobile Computing 2022:6476274. https://doi.org/10.1155/2022/6476274
12. Ahmadi S (2024) Systematic literature review on cloud computing security: threats and mitigation strategies. Journal of Information Security 15:148–167. https://doi.org/10.4236/jis.2024.152010
13. Salama S, Alamoudi Y, Alamoudi G, Albeshri F (2023) Cloud computing security issues and countermeasure: a comprehensive survey. International Journal of Computer Applications 185(14):54–61. https://doi.org/10.5120/ijca2023922832
14. Alaeifar P, Pal S, Jadidi Z, Hussain M, Foo E (2024) Current approaches and future directions for cyber threat intelligence sharing: a survey. Journal of Information Security and Applications 83:103786. https://doi.org/10.1016/j.jisa.2024.103786
15. Pawlicki M, Pawlicka A, Kozik R, Choraś M (2023) The survey and meta-analysis of the attacks, transgressions, countermeasures and security aspects common to the cloud, edge and IoT. Neurocomputing 551:126533. https://doi.org/10.1016/j.neucom.2023.126533
16. Stamenkov G (2025) Cloud service models, business continuity and disaster recovery plans,

and responsibilities. International Journal of Organizational Analysis 33(3):437–451. https://doi.org/10.1108/IJOA-12-2023-4127
17. Noronha LM, Coutinho C (2025) Leveraging cyber resilience using cloud services. Procedia Computer Science 263:548–556. https://doi.org/10.1016/j.procs.2025.07.066
18. Reshmi TR (2021) Information security breaches due to ransomware attacks: a systematic literature review. International Journal of Information Management Data Insights 1(2):100013. https://doi.org/10.1016/j.jjimei.2021.100013
19. Begovic K, Al-Ali A, Malluhi Q (2023) Cryptographic ransomware encryption detection: survey. Computers & Security 132:103349. https://doi.org/10.1016/j.cose.2023.103349
20. Gómez Hernández JA, García Teodoro P, Magán Carrión R, Rodríguez Gómez R (2023) Crypto-ransomware: a revision of the state of the art, advances and challenges. Electronics 12:4494. https://doi.org/10.3390/electronics12214494
21. Patterson CM, Nurse JRC, Franqueira VNL (2023) Learning from cyber security incidents: a systematic review and future research agenda. Computers & Security 132:103309. https://doi.org/10.1016/j.cose.2023.103309
22. Malik AW, Bhatti DS, Park TJ, Ishtiaq HU, Ryou JC, Kim KI (2024) Cloud digital forensics: beyond tools, techniques, and challenges. Sensors 24:433. https://doi.org/10.3390/s24020433
23. Andreoli A, Lounis A, Debbabi M, Hanna A (2023) On the prevalence of software supply chain attacks: empirical study and investigative framework. Forensic Science International: Digital Investigation 44:301508. https://doi.org/10.1016/j.fsidi.2023.301508
24. Sánchez-García ID, San Feliu Gilabert T, Calvo-Manzano JA (2023) Countermeasures and their taxonomies for risk treatment in cybersecurity: a systematic mapping review. Computers & Security 128:103170. https://doi.org/10.1016/j.cose.2023.103170
25. Ardebili AA, Lezzi M, Pourmadadkar M (2024) Risk assessment for cyber resilience of critical infrastructures: methods, governance, and standards. Applied Sciences 14:11807. https://doi.org/10.3390/app142411807
26. Ligo AK, Kott A, Linkov I (2021) How to measure cyber-resilience of a system with autonomous agents: approaches and challenges. IEEE Engineering Management Review 49(2):89–97. https://doi.org/10.1109/EMR.2021.3074288
27. Tam T, Rao A, Hall J (2021) The good, the bad and the missing: a narrative review of cyber-security implications for Australian small businesses. Computers & Security 109:102385. https://doi.org/10.1016/j.cose.2021.102385
28. Admass WS, Munaye YY, Diro AA (2024) Cyber security: state of the art, challenges and future directions. Cyber Security and Applications 2:100031. https://doi.org/10.1016/j.csa.2023.100031
29. Hargreaves C, Breitinger F, Dowthwaite L, Webb H, Scanlon M (2024) DFPulse: the 2024 digital forensic practitioner survey. Forensic Science International: Digital Investigation 51:301844. https://doi.org/10.1016/j.fsidi.2024.301844
30. Saleem M, Warsi MR, Islam S (2023) Secure information processing for multimedia forensics using zero-trust security model for large scale data analytics in SaaS cloud computing environment. Journal of Information Security and Applications 72:103389. https://doi.org/10.1016/j.jisa.2022.103389.