Enhancing Cyber Resilience in Saudi Arabia's Critical Digital Infrastructure: A Risk-Based Security Framework Supporting Vision 2030

Abstract

Saudi Arabia's drive to diversify its economy is contingent upon a suite of interconnected cloud platforms, industrial networks, digital government services and cyber-physical assets. Traditional cyber risk assessments tend to focus on the probability of compromise and are largely blind to the capacity of these critical services to tolerate disruption. This paper reviews how risk assessment, service dependency modelling and resilience engineering can be integrated into a practical framework for critical digital infrastructure in the context of Vision 2030. It is a structured integrative review of 30 DOI identifiable publications from 2020 to 2025, augmented by an assessment of the Saudi cybersecurity control documentation. The evidence is presented in the context of threat propagation, risk assessment technique, operational recoverability and governance accountability. Asset based scores fail to capture the dynamics of cascading disruption, recovery dependencies and industrial constraints. A six stage service centred framework is developed which links mission definition, dependency mapping, scenario analysis, proportionate controls, continuity exercises and adaptive assurance. A proposed measurement portfolio is presented which combines minimum service levels, duration of disruption, restoration integrity, supplier exposure and decision latency. The framework links national control requirements to outcomes that can be measured without in any way asserting that the outcomes are in fact achieved. The primary contribution of the paper is a transparent method by which protective and restorative investments can be prioritised in terms of their consequences for public services and industrial operations rather than in terms of isolated technical vulnerabilities.

UNDER PEER REVIEW IN IJAR

1. Introduction

Saudi Arabia's Vision 2030 is all about digital infrastructure and the importance of that infrastructure is far greater than its physical presence would suggest. The payment gateways, the industrial supervisory networks, the identity services, the health platforms, the telecoms and the cloud environments all underpin the activities that need to be running even if some of the components in the system do fail. The boundaries of these systems don't generally align with the boundaries of the organisations that own them. A utility will be reliant on the communications and identity providers. A hospital will be reliant on power and data exchange. A government platform will be reliant on shared software and on external vendors. Criticality is therefore a property of the services that are delivered and of the dependencies that exist rather than simply a function of the market value of the equipment. There's a considerable body of work on cyber-physical systems and infrastructure protection which shows that increased connectivity leads to a greater number of potential failure points but also opens up new avenues for coordinated control (Yaacoub et al., 2020; De Felice et al., 2022; Zhou et al., 2020).

A preventive security model is one which asks the question "can an adversary exploit a vulnerability?" A resilience model goes further and asks "can the essential service continue to operate safely during the incident?", "how long can degraded operation be tolerated?", "which restoration actions are reliable?" and "what changes will prevent the same failure from reoccurring?". These are different questions with different evidence needs. The difference becomes significant when the loss of system availability could have an impact on physical safety, public confidence or national economic activity. Comparative reviews have found that preparedness, absorption, recovery and adaptation are all part of the same lifecycle (Aghazadeh Ardebili et al., 2024), but do so with inconsistent measurement methodologies and varying degrees of integration with enterprise risk processes (Sepúlveda Estay et al., 2020; Alhidaifi et al., 2024; Araujo et al., 2024). Simply adding more technical controls does not in any way demonstrate that the resulting service is resilient.

The Kingdom has a very specific institutional approach to this issue. The National Cybersecurity Authority's Essential Cybersecurity Controls set out the broad expectations; its Critical Systems Cybersecurity Controls go on to set out additional requirements for nationally significant

UNDER PEER REVIEW IN IJAR

systems; and its Operational Technology and Cloud Cybersecurity Controls deal with the specific requirements of the OT and cloud domains. They provide the necessary steer, but if you just treat them as a set of separate compliance checklists you're going to miss the interdependencies and the competing recovery objectives. This is particularly the case when you have legacy industrial assets interacting with modern cloud-based analytics and remote maintenance. Work in hyperconnected oil and gas and industrial control detection has demonstrated that the operational dependencies and technical heterogeneity of the environment will have a bearing on the security decisions (Pettersen and Grøtan, 2024; Koay et al., 2023).

This paper is concerned with the development of a service-centric, risk-based cyber resilience framework for Saudi critical digital infrastructure. It is not simply an inventory of controls but a framework which sets out how the knowledge of dependencies should be used to inform the creation of disruption scenarios, investment decisions and demonstrable recovery capability. The review is guided by four objectives: to identify failure mechanisms which can propagate across digital and physical services; to compare the methods for assessing uncertain and cascading impacts; to reconcile the requirements of controls with the need for operational continuity; and to define auditable metrics which link resilience investment to service outcomes. The research question is therefore: how can Saudi operators embed risk assessment and continuity engineering into their security governance so as to protect the core functions on which the Vision 2030 programmes are predicated?

A key weakness is the distinction between contractual service availability and actual tested operational independence. A cloud service that is nominally redundant may still share identity infrastructure, software images, privileged admin accounts or communication pathways. Those shared dependencies can in fact prevent a failover when it is most needed. Supplier concentration adds another layer of complexity to the issue of accountability as one supplier's service can underpin a number of different critical organisations. Research into cyber resilience and continuity warns against assuming that the diversity of vendors is in any way equivalent to diversity of failure modes (Sepúlveda Estay et al., 2020; Osei-Kyei et al., 2021). The dependency

UNDER PEER REVIEW IN IJAR

analysis should therefore look at shared control planes and shared operational privileges rather than simply the number of systems that are installed.

2. Conceptual foundations and analytical boundaries

Risk assessment traditionally involves relating threatening events to their likelihood and consequences. That kind of analysis is useful in informing decisions about the protection of assets, but it breaks down when you have a low frequency event with poorly defined consequences or when the likelihood of the event is based purely on expert judgement. The nature of heavy tailed cyber loss distributions is a case in point in terms of why the average can be a very poor indicator of the true risk. Work on resilience performance curves also demonstrates that the time taken to recover and the degree of degradation that occurs provide information which is not captured in a simple risk ranking (Dacorogna et al., 2023; Poulin and Kane, 2021). For critical services the relevant loss is not simply the compromise of a device; it is the time dependent degradation of a function that society expects to be delivered.

A service centric approach defines four related quantities. Exposure is about the level of plausible attack access and vulnerabilities. Dependency is to do with the extent to which a service depends on other services. Consequence is the measure of safety, operational, economic or societal harm. Adaptive capacity is the ability to continue or to recover activities in the face of adverse conditions. A hospital's identity platform might be highly exposed and yet be its least recoverable dependency. On the other hand a stand alone process controller may be highly protected and yet be an unacceptable single point of failure. Research into complex infrastructure and Bayesian smart grid modelling has shown that the topological and conditional relationships have a significant effect on the system level risk (Mottahedi et al., 2021; Ibne Hossain et al., 2020).

UNDER PEER REVIEW IN IJAR

Cyber resilience also requires a clear distinction between availability and integrity. A service that has been restored which is then processing corrupted data or issuing unsafe physical commands is not in any sense successfully recovered. In an industrial context, segmentation, operator visibility, deterministic timing and validated fail-safe modes may well be more important than a rapid automated recovery. Industrial risk research has highlighted the danger of applying standard IT assumptions to industrial systems where the behaviour of ageing equipment, the nature of maintenance windows and the state of safety interlocks may all be very different (Khanna and Govindarasu, 2024; Kolosok and Gurina, 2022). The criteria for recovery therefore need to include trustworthy data, process stability and authorisation to operate as well as speed.

Governance provides the other boundary. You can't just hand cyber resilience over to a Security Operations Centre as that's not how the risk decisions are made. Procurement, engineering, safety, finance and service owners all have a bearing on those decisions. The Cybersecurity Framework 2.0 raises governance as a key element and the systems security engineering guidance treats anticipation, resistance, recovery and adaptation as design objectives (NIST, 2024; Ross et al., 2021). The value of that is in the methodology rather than the jurisdiction. The Saudi legal and regulatory requirements will always take precedence and the internationally developed guidance is purely as an optional analytical tool rather than a replacement for a rulebook.

Uncertainty has to be represented explicitly. We simply don't know with the same degree of confidence the attacker's capability, the operator's recovery skill or the state of the backup at the time of a future incident. A useful assessment therefore would separate evidence from actual drills, from engineering specs and from untested staff assumptions. That doesn't remove the uncertainty, it just changes the quality of the decisions that are made. Scenario ranges can be updated as a result of post incident reviews, component replacement or supplier testing. Sensitivity analysis should be used to determine whether the investment selected is still the preferred one when the key assumptions are varied. It's important that small shifts in the level of speculation don't cause the priority to be reversed without any real new evidence (Chandra et al., 2022; Dacorogna et al., 2023).

UNDER PEER REVIEW IN IJAR

3. Review design and method

The study is an integrative review using a structured approach to answer a design question that cuts across technical risk analysis, critical infrastructure engineering and organisational governance. An integrative approach is the right one given the range of evidence selected which includes analytical models, conceptual syntheses, sectoral case studies and standards based guidance that would not be appropriate to combine into a single effect size. The review corpus consists of 30 items with DOIs which were published between 2020 and 2025. The candidate papers were identified via a backwards search of the reference lists of the two provided papers, checking of bibliographic metadata and then through a process of purposive selection for relevance to the themes of infrastructure interdependencies, threat mechanisms, recoverability and security governance. The two papers provided in 2026 were only used to inform expectations around the presentation and the topic and were not included as part of the 2020-2025 evidence base.

Selection was driven by conceptual relevance rather than by sheer volume of citations. The sources had to provide an explicit method, an empirical observation or a defensible analytical position on one of the following: cyber resilience, cyber-physical risk, operational continuity, dependent infrastructure or cybersecurity governance. Purely consumer privacy issues, issues relating to unrelated cryptography or general technology adoption without any implication for the service were excluded. Three technical reports with persistent DOIs were selected as they deal directly with controls, engineering and governance and are therefore directly relevant to framework design. This is a deliberately curated set of sources rather than a full enumeration of the database results; no details of any unperformed screening stages, no reviewer agreement statistics and no fabricated search totals are reported.

UNDER PEER REVIEW IN IJAR

Each piece of work was assessed against five analytical questions: what is the asset in question; what disruption is being modelled; how are the consequences and uncertainties being expressed; what recovery or adaptation capability is being considered; and what organisational decision is the analysis leading to? Relevant statements were then placed into a thematic matrix covering service dependencies, scenario development, control selection, recovery validation and accountability. Contradictions were left in rather than being smoothed out. Evidence derived from a simulation was not described as national operational performance and a proposed architecture was not presented as an actual control. The synthesis then went on to compare the recurring insights across the utility, industrial, health, urban and cloud contexts.

Regulatory contextualisation was achieved using the publicly available Saudi cybersecurity control materials (the Essential Cybersecurity Controls 2-2024, the Critical Systems Cybersecurity Controls, the Operational Technology Cybersecurity Controls 1-2022 and the Cloud Cybersecurity Controls 2-2024). These core policy documents do not all have scholarly DOIs and were looked at separately from the 30-source DOI bibliography. Their function was to identify the relevant control domains and the institutional constraints, not to provide a quantified measure of effectiveness. The framework was developed abductively with the failure modes identified in the literature informing the architecture and the Saudi control requirements constraining the range of possible governance options. The findings are therefore to be understood as reasoned design propositions which will require field validation.

The review criteria were transparency of assumptions, transferability to mission critical services, treatment of dependencies, measure of recoverability and applicability to the Saudi operational context. An intrusion detection paper might be of value in informing the design of a monitoring system but in itself would not prove anything about organisational continuity. A healthcare risk model might be useful in demonstrating the impact of identity dependencies but would not in itself allow you to quantify the exposure of a Saudi hospital. Those limitations serve to prevent any overreaching of the methodology and define clearly the boundary between the reviewed evidence, the regulatory context and the original prescriptions of the article (Chandra et al., 2022; Smyrlis et al., 2024; Mehmood et al., 2023).

UNDER PEER REVIEW IN IJAR

Biblical validity and evidential weight were assessed separately. Each item selected was recorded with its year of publication, full title and persistent digital object identifier (DOI) and duplicate titles were collapsed. A check was made as to whether a DOI was present in the bibliography or in the official publication record at the time of preparation (this is a metadata check and does not in any way guarantee that all publisher landing pages will be accessible). The review deliberately does not attempt to pool together incomparable claims about detection accuracy, financial loss or recovery time. Its inferential test is whether the proposed measure is a plausible consequence of the identified failure modes and whether an operator could in fact obtain observations which would allow them to disprove the benefit which the measure is supposed to provide. Future verification is thus part of the framework rather than a post hoc justification.

4. Critical synthesis of the evidence

4.1. Interdependence is the primary unit of concern

Operators tend to structure their controls around administrative assets rather than the actual functional relationships through which disruptions propagate. Identity federation, clock sync, industrial remote access, cloud control planes and network routing are all examples of services which can fail across multiple independent services. A topology informed assessment therefore requires a directed dependency map which shows what upstream components each critical service depends on, how quickly an alternative provider can be brought in and whether those alternatives have the same underlying pre-requisites. Work on smart grid systems and on interconnected infrastructure has demonstrated the value of conditional modelling over simple rankings of individual devices (Ibne Hossain et al., 2020; Zhou et al., 2020; Osei-Kyei et al., 2021).

UNDER PEER REVIEW IN IJAP

Dependencies are not all the same. A weak dependency might just add a little bit of latency; an exclusive control channel might actually halt a physical process. The distinction requires consequence mapping at the level of the service output. The loss of a customer facing energy application is very different to the loss of operational dispatch even if they're on the same network. The same is true in healthcare and aviation investigations where the risk is very much a function of the workflow, the human decisions and the operational context rather than the software vulnerability per se (Smyrlis et al., 2024; Jovanović et al., 2020). In the Saudi context the interdependency registers should therefore show the safety critical boundaries, the national or contractual priority, the manual alternatives, the supplier control and the permissible outage times.

A dependency map must also show direction and timing. Some services will fail straight away if the supplier goes away; others will carry on until their credentials expire, their storage buffers are full or their manual staffing runs out. The length of that delay is what determines when you can actually contain the situation and whether a temporary operating mode is viable. In the course of an exercise you can of course annotate the edges with acceptable interruption windows, feasible alternatives and the quality of the supporting evidence. The map that emerges is not a permanent architectural drawing at all. It has to be versioned when a project brings in new sensors, moves workloads or changes its remote maintenance contracts. This is particularly important in the case of digitally expanding urban systems where a static diagram can quickly become misleading over the course of a procurement cycle (De Felice et al., 2022; Pourmadadkar et al., 2024).

4.2. Threat modelling must capture persistence and recovery compromise

Ransomware, compromised software updates, abuse of privileged accounts and manipulation of industrial telemetry can all lead to very specific degradation profiles. Attack graphs and trees will show you the multistep attack paths, but the assumptions they make about what assets are reachable and what the attacker can do at each stage need to be revisited (Lallie et al., 2020). An infrastructure aware assessment should blend the attack pathways with the operational states. The

UNDER PEER REVIEW IN IJAR

operational states would be normal, degraded, isolated, locally controlled, partially restored and verified. That last one is important in that the incident can continue even after the attack has finished if the credentials are still untrusted, the backups are compromised or the physical processes require a very careful recommissioning.

The literature points to scenario completeness being more important than apparent numerical precision. Expert scored probability-impact grids can still be of use in triage but are poor at capturing correlated suppliers, rapidly changing configuration or the effects of rare events with systemic consequences. There are methods in cyber loss modelling to deal with uncertainty and extreme exposure and power substation work has demonstrated the benefit of linking protective investment to resilience objectives (Dacorogna et al., 2023; Khanna and Govindarasu, 2024). The practical takeaway is to use ranges and scenarios, keep an assumptions register and never turn a subjective ordinal score into a false sense of monetary certainty. Scenario severity should be challenged through exercises rather than simply accepted because it has a decimal point (Ardebili et al., 2022).

4.3. Control compliance and operational assurance are complementary

Saudi cybersecurity controls define the governance and defensive responsibilities for the essential, critical, cloud and operational technology environments. The operator still needs to be able to demonstrate that the controls actually work together in the event of a real incident. A logged backup is not proof of recoverability. An approved supplier contract is not proof of offline access. A successful malware detection is not evidence that the emergency controls will in fact protect the physical safety of the site. Research into energy connected IoT and industrial systems has identified the problems of architectural complexity, inconsistent telemetry and recovery constraints which may not be apparent from a basic compliance inventory (Mehmood et al., 2023; Koay et al., 2023; Pettersen and Grøtan, 2024).

UNDER PEER REVIEW IN IJAR

This reinforces the assurance model which is based on demonstrable evidence. The controls need to be tied to a failure mode that you would expect to see, have a designated owner, a test procedure and retained evidence. Configuration checks are about prevention; adversarial exercises are about detection and containment; restoration drills are about data integrity and service level continuity; and independent governance reviews are about whether the residual risk is acceptable. The control catalogues and the various studies of product cybersecurity governance are useful in terms of categorisation, but they don't in any way obviate the need for local justification and assessment (Shaffique, 2024; Joint Task Force, 2020; National Institute of Standards and Technology, 2024). The governance question that remains is whether the system in question can deliver its minimum critical function, not whether its documentation is complete.

4.4. Emerging analytics require restrained deployment

ML systems have the potential to improve industrial anomaly detection by being able to spot the unfamiliar. Their performance is however entirely dependent on the quality of the training data, the level of observability and the changing nature of the operating conditions. The literature is full of references to the problems of attack diversity, false alarms and the differences between production equipment and the research datasets (Koay et al., 2023; Al-Hawawreh and Doss, 2023). As such critical infrastructure should treat any automated detection as a form of decision support with clear thresholds for human verification and safe containment. Where an incorrect automated action could lead to a loss of power, a loss of healthcare or a loss of transport then conservative fallbacks and operating authority approval are absolutely essential.

Governance research also cautions that security tools can themselves introduce new vulnerabilities, new model management burdens and new accountability issues. The literature on cybersecurity frameworks and AI deployment makes it very clear that you need to define auditability, operator responsibility and model lifecycle controls before you start to scale up automation (McIntosh et al., 2024). In reality the claimed accuracy of a detector is of considerably less importance than the proportion of high impact incidents that it finds early

UNDER PEER REVIEW IN IJAR

enough to be able to act on, the proportion of false positives that result in service loss and the time it takes to return an isolated system safely back to operation. Those are actual, measurable requirements rather than an argument for using the technology simply because it's new.

UNDER PEER REVIEW IN IJAR

Table 1. Cross-sector threat-to-service pathways and proposed resilience indicators.

Domain	Disruption pathway	Priority assurance control	Observable measure
Electricity / industrial OT	Loss or manipulation of dispatch data; unsafe remote action	Local process control, segmented remote access	Time below safe operating threshold; verified restart
Healthcare and emergency	Shared identity outage, ransomware or inaccurate restored records	Downtime procedures, offline authentication, record reconciliation	Clinical-service continuity; record-integrity pass rate
Cloud / digital government	Control-plane compromise or correlated provider failure	Independent recovery identity, clean configuration baselines	Verified recovery time; shared-dependency count
Transport / urban platforms	Telecom outage and cascade across field devices	Degraded-mode operations, tested communications fallback	Duration of essential mobility-service loss
Water and energy IoT	Telemetry spoofing or maintenance-vendor compromise	Isolated engineering access, validated sensing and fail-safe process	Containment time; unverified command exposure

UNDER PEER REVIEW

5. Proposed service-centred cyber resilience framework

5.1. Architecture and decision gates

Figure 1 is a six stage process which starts with the service rather than the asset register. Stage 1 is to identify the key outputs, the owners of those outputs, the safety constraints and the minimum acceptable capacity. Stage 2 is to map the technology, people, suppliers and physical infrastructure dependencies and to record the common cause vulnerabilities. Stage 3 is to develop a set of plausible scenarios (including the possibility of a malicious disruption during the recovery phase and the possibility of a simultaneous supplier failure). Stage 4 is to assess the options between preventative, absorptive and restorative measures. Stage 5 is to validate the measures through simulation, operational drills and evidence of safe re-entry. Stage 6 is to review the assumptions in the light of the exercises, incidents, design changes and procurement decisions. Each stage produces a decision record which allows the decision to be independently challenged rather than being the result of an untraceable technical prioritisation.

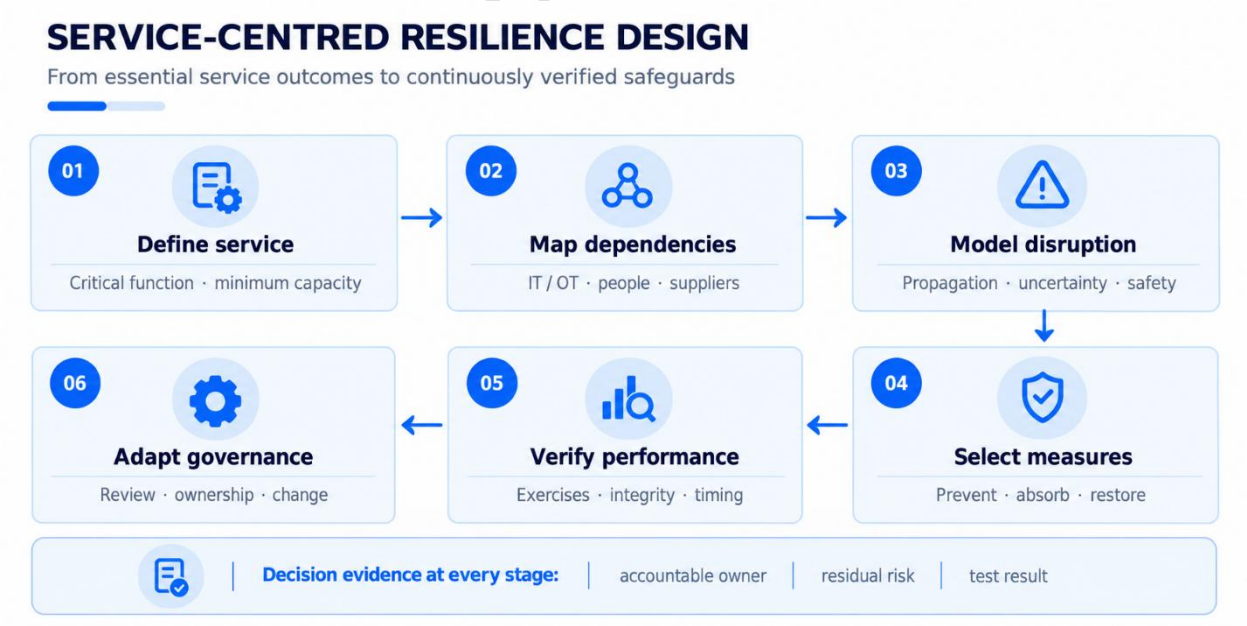


Figure 1. Proposed six-stage architecture. The illustration is an original conceptual synthesis, not a validated deployment.

UNDER PEER REVIEW IN IJAR

There's a governance gate between risk acceptance and operational execution. The service owner defines the acceptable outcomes. The engineering team confirms the physical feasibility. The cyber security team assess the threat environment. Executive management then go on to approve the residual risk. Supplier arrangements need to cover emergency credentials, restoration responsibilities, data access rights and the ability to test the subcontractor dependencies. Investment in cyber-resilience is made where the intervention changes the expected service loss profile or removes an unacceptable single point of failure. A protective tool which simply improves the vulnerability score without changing the mission impact may be of lower priority than a tested offline operating procedure (Ross et al., 2021; Sepúlveda Estay et al., 2020; Pourmadadkar et al., 2024).

Design varies by asset class. For data centres and digital government it's likely to be about immutable backups, recovery isolation, alternative identity and integrity verification. For the utilities and industrial processes it's all about segmented engineering access, local control and staged recommissioning. In the case of critical healthcare services you'd have very well rehearsed clinical downtime processes and a requirement to reconcile records on return to service. The differences are such that you wouldn't in any way be able to treat an industrial controller as a typical office workstation. They also provide a solid basis for making the case for prioritisation when spend is being competed against infrastructure modernisation or service expansion (Khanna and Govindarasu, 2024; Smyrlis et al., 2024; Pettersen and Grøtan, 2024).

5.2. Risk evaluation and measurement

The proposed risk register includes a service ID, disruption scenario, critical dependencies, confidence in likelihood, affected minimum service level, maximum tolerable interruption and control effectiveness evidence. The three recommended approaches are: ordinal screening (which identifies the most urgent exposures without any implication of numerical certainty), dependency based scenario analysis (which looks at the conditional and cross sectoral consequences), and time dependent performance analysis (which looks at the depth and duration of the degraded

UNDER PEER REVIEW IN IJAR

service). Bayesian structures can be used to update the assumptions as new dependency evidence becomes available and resilience curves can be used to express the operational impact of different restoration options (Ibne Hossain et al., 2020; Poulin and Kane, 2021; Chandra et al., 2022).

For service s, let q_s(t) be the verified fraction of the required service delivered at time t relative to the baseline, and let q_min be the minimum acceptable fraction during emergency operation. The service loss area over the assessment horizon T is L_s = integral from 0 to T of [1 - q_s(t)] dt (this captures both the depth and the duration of the disruption). Report separately the time below q_min, the time until verified restoration, the percentage of dependent suppliers with exercised contingencies and the share of recovered data which passes integrity checks. These must not be aggregated into a single, opaque resilience score. The units, thresholds and uncertainties must be visible to the risk owners.

Prioritisation can be made explicit without us claiming to have a universal formula. A decision panel should be looking at the various options in terms of the estimated area of avoided service loss, the recovery integrity, the safety implications, the level of implementation disruption and the confidence in the evidence. A very expensive isolation mechanism may be perfectly justified for an essential dispatch system and yet would be excessive for a non-critical information portal. Where two controls provide the same level of security improvement the one with the fewest new common dependencies and the most testable recovery path would probably be the preferred option. The rationale for any alternatives that are rejected should be recorded and the conditions under which they would be reconsidered should be established. This makes the approach repeatable across budget cycles and allows for the possibility that national security or life safety obligations may take precedence over a purely economic ranking (Ross et al., 2021; Janta et al., 2024).

Figure 2 shows why two responses with the same detection time can lead to very different continuity outcomes. The trajectories shown are indicative only and are not derived from the

UNDER PEER REVIEW IN IJAR

Saudi operator data. A well prepared service can absorb the attack via redundancy and then recover via a validated process, whereas a poorly prepared service may be below its critical threshold for a considerably longer period of time. The model also takes into account governance latency (the time that it takes the responsible people to approve safe isolation or safe re-entry). Time to detection, time to contained operation and time to verified recovery should all be recorded separately so that an early detection doesn't hide a prolonged service outage.

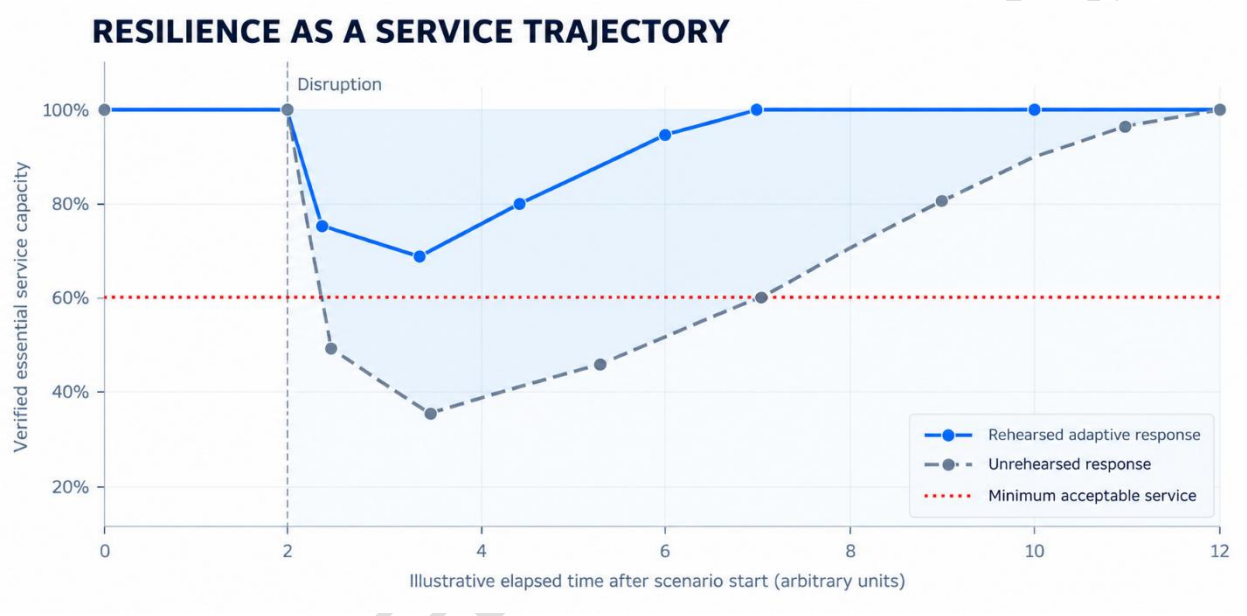


Figure 2. Conceptual service-capacity trajectories. Values and response shapes are illustrative, not observed data.

5.3. Worked application: an interconnected public-service platform

Imagine a hypothetical regional digital platform which serves up a hospital network, a transport operator and a utility's customer channels via a common cloud identity and telecoms platform. A privileged identity compromise leads to loss of admin access, malware triggers network isolation and an external support contractor is unable to immediately re-establish a trusted control plane. A traditional assessment might result in three high severity asset risks. The proposed approach instead looks at which of the clinical, mobility and utility functions are actually impacted, separates out the public from the essential and identifies which services can continue to operate using local credentials or pre-defined offline processes.

Scenario workshops define the three recovery priorities: protect the safety-critical local operations, restore the trusted identity of the emergency users and only re-introduce the customer systems once integrity has been verified. A tabletop exercise will determine whether the contractual recovery times are in line with the operational thresholds. A technical drill will assess the ability to rotate credentials and restore clean records. A governance drill will assess who is able to override a failing dependency. The scenario is entirely hypothetical and makes no claims as to any observed Saudi outages. The value of the scenario is that it exposes the coupled failures that are masked by the institution's own inventory of services and it produces a set of testable changes to the architecture, the agreements and the staff procedures.

Table 2. Governance, evidence and review arrangements for the proposed framework.

Accountable function	Required decision or check	Retained evidence	Review trigger
Mission owner	Specify minimum service and safety thresholds	Approved service-impact statement	Quarterly or material service change
Architecture / OT engineering	Validate dependencies and safe degraded modes	Dependency map; failover and recommissioning record	After every major architecture change
Security operations	Confirm detection-to-containment decisions	Scenario playbook; alert and response traces	Exercises and post-incident
Recovery / data owners	Verify backups and restoration integrity	Recovery run; reconciliation and integrity report	At planned resilience exercise
Procurement / suppliers	Test emergency access and supplier assumptions	Contractual recovery commitment; joint drill record	Renewal and significant supplier changes
Risk committee / executive	Decide residual risk, exceptions and remediation	Signed exception register; investment rationale	Annual review and threshold breach

UNDER PEER REVIEW IN IJAR

6. Discussion and implications for Vision 2030

The synthesis supports a strategic repositioning: critical digital infrastructure should be governed as a network of essential public and industrial services rather than a collection of independently secured technologies. This changes investment decisions. A higher-performing intrusion detector is worthwhile only if alerts trigger safe intervention; additional cloud redundancy helps only if identity and configuration dependencies are independent; and recovery funding delivers value only where exercise evidence demonstrates that the service can restart with integrity. Prior studies on complex-system resilience and energy recoverability reinforce the importance of modelling adaptations rather than assuming them (Janta et al., 2024; Mottahedi et al., 2021).

For Saudi operators, the framework implies a two-layer accountability structure. The first layer demonstrates conformity with applicable national controls; the second presents service-based resilience evidence to executives and operational authorities. The layers should share a common mapping so that controls, incidents, tests and improvement programmes are connected instead of duplicated. This arrangement is especially relevant to infrastructure programmes with multiple contractors, externally managed platforms and fast-changing technology. It also offers procurement teams a basis for requiring tested recovery commitments rather than accepting generic assurances of compliance or availability.

Several trade-offs remain unresolved. Replication may reduce outage risk but increase synchronised attack exposure if both environments share administrators. Automated containment may shorten attacker dwell time but stop a critical process if safety dependencies are poorly understood. Extensive monitoring may improve forensic visibility while raising privacy and operational overhead. Review evidence from cyber-physical systems, operational analytics and governance indicates that these trade-offs require context-specific assessment, not universal rules

UNDER PEER REVIEW IN IJAR

(Yaacoub et al., 2020; McIntosh et al., 2024; Al-Hawawreh and Doss, 2023). Risk owners should document the conditions under which a control's costs or secondary consequences exceed its benefits.

Cross-sector exercises offer a further implication. A hospital may test its own backups successfully while remaining unable to operate when a shared telecommunications or authentication provider fails. Joint exercises involving suppliers, engineers and service owners can reveal incompatible assumptions about responsibility, recovery order and access to emergency information. The outputs should be carefully scoped: common high-level lessons can inform sector coordination, while sensitive technical details remain protected. Governance should also define escalation when a vendor's preferred restoration sequence conflicts with a local operator's safety requirement. Without such agreements, an apparently mature control programme may still produce slow and contested operational decisions during a major incident (Jovanović et al., 2020; Alhidaifi et al., 2024).

7. Limitations and research agenda

This review intentionally samples conceptually relevant studies rather than undertaking exhaustive database screening. Its conclusions are therefore suitable for theory development and operational design, not for estimating the prevalence of Saudi vulnerabilities or effect sizes for individual controls. Sectoral studies differ in exposure assumptions, modelling resolution and outcome measures. Some findings derive from laboratory or simulated environments, and their transfer to live Saudi systems depends on engineering configuration and governance maturity. Further, policy documents evolve; organisations must verify the operative edition and applicability of national controls before implementation.

A robust next stage would test the framework in several Saudi sectors using a common yet non-sensitive reporting protocol. Researchers should establish baseline service dependencies, elicit

UNDER PEER REVIEW IN IJAR

operational thresholds, conduct controlled outage exercises and repeat measurements after prioritised changes. Comparative analysis could examine time below minimum service, restoration integrity, recovery-cost distribution and decision latency. Longitudinal follow-up would distinguish short-term training effects from structural improvement. Mixed technical and organisational evidence would be particularly valuable because a mathematically robust risk model cannot compensate for missing authority to make urgent operational decisions (Poulin and Kane, 2021; Alhidaifi et al., 2024).

8. Conclusion

Cyber resilience for Saudi Arabia's critical digital infrastructure requires a change in the object and time horizon of risk management. Service delivery, dependencies and trustworthy restoration must become explicit parts of security assessment, particularly where industrial, cloud and public-service environments converge. This structured review integrates recent evidence into a six-stage framework linking essential-service definition, dependency mapping, scenario analysis, control selection, operational verification and continuing adaptation. The framework adds decision accountability and interpretable continuity measures to familiar protective controls, while respecting the distinct constraints of national cybersecurity requirements.

The principal contribution is a testable design logic, not an unverified assertion of effectiveness. By measuring whether indispensable functions remain above agreed thresholds and return safely after disruption, operators can make more defensible investment and governance decisions. Future sector-based trials should examine both technical performance and organisational readiness, thereby establishing whether the proposed approach produces materially stronger continuity outcomes for Vision 2030's interconnected services.

UNDER PEER REVIEW IN IJAR

References

1. Aghazadeh Ardebili, A., Lezzi, M. and 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
2. Al-Hawawreh, M. and Doss, R. (2023). Enhancing security in industrial IoT: A taxonomy-driven approach to risk assessment. In 2023 IEEE 22nd International Conference on Trust, Security and Privacy in Computing and Communications, 434–443. https://doi.org/10.1109/TrustCom60117.2023.00074
3. Alhidaifi, S.M., Asghar, M.R. and Ansari, I.S. (2024). A survey on cyber resilience: Key strategies, research challenges, and future directions. ACM Computing Surveys, 56, 196. https://doi.org/10.1145/3649218
4. Araujo, M.S.D., Machado, B.A.S. and Passos, F.U. (2024). Resilience in the context of cyber security: A review of the fundamental concepts and relevance. Applied Sciences, 14, 2116. https://doi.org/10.3390/app14052116
5. Ardebili, A.A., Padoano, E., Longo, A. and Ficarella, A. (2022). The risky-opportunity analysis method (ROAM) to support risk-based decisions in a case-study of critical infrastructure digitization. Risks, 10, 48. https://doi.org/10.3390/risks10030048
6. Chandra, N.A., Ramli, K., Ratna, A.A.P. and Gunawan, T.S. (2022). Information security risk assessment using situational awareness frameworks and application tools. Risks, 10, 165. https://doi.org/10.3390/risks10080165
7. Dacorogna, M., Debbabi, N. and Kratz, M. (2023). Building up cyber resilience by better grasping cyber risk via a new algorithm for modelling heavy-tailed data. European Journal of Operational Research, 311, 708–729. https://doi.org/10.1016/j.ejor.2023.05.003
8. De Felice, F., Baffo, I. and Petrillo, A. (2022). Critical infrastructures overview: Past, present and future. Sustainability, 14, 2233. https://doi.org/10.3390/su14042233
9. Ibne Hossain, N.U., Nagahi, M., Jaradat, R., Shah, C., Buchanan, R. and Hamilton, M. (2020). Modeling and assessing cyber resilience of smart grid using Bayesian network-

UNDER PEER REVIEW IN IJAR

based approach: A system of systems problem. Journal of Computational Design and Engineering, 7, 352–366. https://doi.org/10.1093/jcde/qwaa029

10. Janta, P., Leeraphun, N., Thapmanee, K., Niyomna, P., Sintuya, H., Setthapun, W., Maneechot, P., Sriprapakhan, P., Chollacoop, N. and Silva, K. (2024). Energy resilience assessment: Incorporating consideration of recoverability and adaptability in risk assessment of energy infrastructure. Energy for Sustainable Development, 81, 101506. https://doi.org/10.1016/j.esd.2024.101506
11. Joint Task Force (2020). Security and privacy controls for information systems and organizations. NIST Special Publication 800-53, Revision 5. https://doi.org/10.6028/NIST.SP.800-53r5
12. Jovanović, A., Klimek, P., Renn, O., Schneider, R., Øien, K., Brown, J., DiGennaro, M., Liu, Y., Pfau, V., Jelić, M. et al. (2020). Assessing resilience of healthcare infrastructure exposed to COVID-19: Emerging risks, resilience indicators, interdependencies and international standards. Environment Systems and Decisions, 40, 252–286. https://doi.org/10.1007/s10669-020-09779-8
13. Khanna, K. and Govindarasu, M. (2024). Resiliency-driven cyber–physical risk assessment and investment planning for power substations. IEEE Transactions on Control Systems Technology, 32, 1743–1754. https://doi.org/10.1109/TCST.2024.3378990
14. Koay, A.M.Y., Ko, R.K.L., Hettema, H. and Radke, K. (2023). Machine learning in industrial control system (ICS) security: Current landscape, opportunities and challenges. Journal of Intelligent Information Systems, 60, 377–405. https://doi.org/10.1007/s10844-022-00753-1
15. Kolosok, I. and Gurina, L. (2022). Cyber resilience models of systems for monitoring and operational dispatch control of electric power systems. IFAC-PapersOnLine, 55, 485–490. https://doi.org/10.1016/j.ifacol.2022.07.084
16. Lallie, H.S., Debattista, K. and Bal, J. (2020). A review of attack graph and attack tree visual syntax in cyber security. Computer Science Review, 35, 100219. https://doi.org/10.1016/j.cosrev.2019.100219
17. McIntosh, T.R., Susnjak, T., Liu, T., Watters, P., Xu, D., Liu, D., Nowrozy, R. and Halgamuge, M.N. (2024). From COBIT to ISO 42001: Evaluating cybersecurity frameworks for opportunities, risks, and regulatory compliance in commercializing large

UNDER PEER REVIEW IN IJAR

language models. Computers & Security, 144, 103964. https://doi.org/10.1016/j.cose.2024.103964

18. Mehmood, A., Epiphaniou, G., Maple, C., Ersotelos, N. and Wiseman, R. (2023). A hybrid methodology to assess cyber resilience of IoT in energy management and connected sites. Sensors, 23, 8720. https://doi.org/10.3390/s23218720
19. Mottahedi, A., Sereshki, F., Ataei, M., Nouri Qarahasanlou, A. and Barabadi, A. (2021). The resilience of critical infrastructure systems: A systematic literature review. Energies, 14, 1571. https://doi.org/10.3390/en14061571
20. National Institute of Standards and Technology (2024). The NIST Cybersecurity Framework (CSF) 2.0. NIST Cybersecurity White Paper 29. https://doi.org/10.6028/NIST.CSWP.29
21. Osei-Kyei, R., Tam, V., Ma, M. and Mashiri, F. (2021). Critical review of the threats affecting the building of critical infrastructure resilience. International Journal of Disaster Risk Reduction, 60, 102316. https://doi.org/10.1016/j.ijdrr.2021.102316
22. Pettersen, S. and Grøtan, T.O. (2024). Exploring the grounds for cyber resilience in the hyper-connected oil and gas industry. Safety Science, 171, 106384. https://doi.org/10.1016/j.ssci.2023.106384
23. Poulin, C. and Kane, M.B. (2021). Infrastructure resilience curves: Performance measures and summary metrics. Reliability Engineering & System Safety, 216, 107926. https://doi.org/10.1016/j.ress.2021.107926
24. Pourmadadkar, M., Lezzi, M. and Corallo, A. (2024). Cyber security for cyber-physical systems in critical infrastructures: Bibliometrics analysis and future directions. IEEE Transactions on Engineering Management, 71, 15405–15421. https://doi.org/10.1109/TEM.2024.3489273
25. Ross, R., Pillitteri, V., Graubart, R., Bodeau, D. and McQuaid, R. (2021). Developing cyber-resilient systems: A systems security engineering approach. NIST Special Publication 800-160, Volume 2, Revision 1. https://doi.org/10.6028/NIST.SP.800-160v2r1
26. Sepúlveda Estay, D.A., Sahay, R., Barfod, M.B. and Jensen, C.D. (2020). A systematic review of cyber-resilience assessment frameworks. Computers & Security, 97, 101996. https://doi.org/10.1016/j.cose.2020.101996

UNDER PEER REVIEW IN IJAR

27. Shaffique, M.R. (2024). Cyber Resilience Act 2022: A silver bullet for cybersecurity of IoT devices or a shot in the dark? Computer Law & Security Review, 54, 106009. https://doi.org/10.1016/j.clsr.2024.106009
28. Smyrlis, M., Floros, E., Basdekis, I., Prelipcean, D.B., Sotiropoulos, A., Debar, H., Zarras, A. and Spanoudakis, G. (2024). RAMA: A risk assessment solution for healthcare organizations. International Journal of Information Security, 23, 1821–1838. https://doi.org/10.1007/s10207-024-00820-4
29. Yaacoub, J.P.A., Salman, O., Noura, H.N., Kaaniche, N., Chehab, A. and Malli, M. (2020). Cyber-physical systems security: Limitations, issues and future trends. Microprocessors and Microsystems, 77, 103201. https://doi.org/10.1016/j.micpro.2020.103201
30. Zhou, Y., Yu, F.R., Chen, J. and Kuo, Y. (2020). Cyber-physical-social systems: A state-of-the-art survey, challenges and opportunities. IEEE Communications Surveys & Tutorials, 22, 389–425. https://doi.org/10.1109/COMST.2019.2959013

UNDER PEER REVIEW IN IJAR