

Cyber Resilience Framework for Distributed Giga-Project IT Infrastructure in Saudi Arabia under Vision 2030: Integrating ISO 27001, ISO 22301, Hybrid-Cloud Security, Business Continuity and Incident Recovery .

Abstract:

Saudi giga-projects depend on distributed digital infrastructure that spans construction sites, corporate offices, command centres, operational technology, private data centres, public cloud services, edge platforms and international suppliers. This distribution enables rapid delivery, yet it also creates tightly coupled dependencies across identity, networks, data, applications and third parties. A local cyber incident can therefore become a programme-wide operational disruption if trust boundaries, recovery priorities and decision authority are unclear. This review develops an integrated cyber-resilience framework for giga-project information technology under Vision 2030. It synthesises evidence published from 2020 to 2025 across information-security management, business continuity, hybrid-cloud architecture, zero trust, critical-infrastructure guidance and Saudi cybersecurity requirements. The framework joins ISO 27001 governance with ISO 22301-aligned continuity practices, service-based risk analysis, segmented hybrid-cloud controls, clean recovery, coordinated incident command and evidence-led assurance. The review argues that resilience cannot be demonstrated by control compliance, backup completion or nominal system availability in isolation. It must be shown through the continued delivery of prioritised business and operational outcomes during compromise, followed by trusted restoration within agreed tolerances. Two original models are proposed: a distributed resilience architecture that separates governance, service, platform and recovery layers; and an incident-to-recovery assurance loop that links detection, containment, continuity decisions, clean restoration and learning. The resulting approach is designed for Saudi giga-project environments characterised by high change velocity, multiple delivery partners, geographically dispersed operations and nationally significant service obligations.

Keywords: *cyber resilience; Saudi Vision 2030; giga-projects; ISO 27001; ISO 22301; hybrid cloud; business continuity; incident recovery; critical infrastructure*

33 1. Introduction

34 Vision 2030 has accelerated the creation of large, technology-intensive programmes in tourism,
35 transport, energy, logistics, urban development, entertainment and public services. Their delivery
36 models differ from those of stable enterprises. A giga-project may establish temporary offices,
37 construction compounds, digital twins, smart assets, security operations, cloud landing zones, shared
38 enterprise platforms and operational command systems at the same time. New contractors, systems and
39 sites are onboarded continuously, while governance must keep pace with changing ownership, design
40 maturity and operational criticality. Digital capability is consequently not a supporting utility alone; it
41 is part of the mechanism through which the programme coordinates work, protects people, controls
42 assets and demonstrates progress [1-5].

43 This operating model creates a distinctive risk system. Distributed teams rely on identity federation,
44 remote administration, software-as-a-service, cloud-hosted collaboration, site connectivity, internet-of-
45 things platforms and vendor-managed technology. Dependencies cross organisational and geographic
46 boundaries, but accountability frequently remains fragmented across project companies, delivery
47 partners, cloud providers and specialist operators. A compromise of a shared identity platform,
48 network management plane, integration service or backup administrator can therefore affect several
49 projects simultaneously. Traditional perimeter security is insufficient because the effective perimeter
50 changes whenever a new workload, contractor, interface or site is introduced [2-9].

51 Cyber resilience provides a more appropriate objective than prevention alone. Contemporary
52 resilience guidance describes the ability to anticipate, withstand, recover from and adapt to adverse
53 conditions while preserving essential purpose [17,18,27]. This definition changes investment priorities.
54 A control is valuable not only because it reduces the probability of compromise, but also because it
55 restricts propagation, preserves decision options, protects recovery assets and supports reliable
56 restoration. For a giga-project, the relevant outcome is not that every digital service remains
57 continuously available. It is that safety, command, financial control, stakeholder communication and
58 critical operations continue at an acceptable level while affected services are isolated and rebuilt.

59 Information-security and business-continuity disciplines address different parts of this problem.
60 ISO 27001 provides a management system for governing security risks, responsibilities, objectives,
61 assurance and improvement [10]. ISO 27002 and ISO 27005 expand the control and risk-treatment
62 logic [11,12]. ISO 22301 is operationalised through the guidance in ISO 22313, which links business
63 impact, continuity strategies, procedures, exercises and continual improvement [13]. Crisis-
64 management guidance adds strategic command, communications and decision structures [14]. Used
65 independently, these disciplines can produce parallel registers, exercises and committees. Used
66 together, they can create a service-centred resilience system in which cyber controls, continuity
67 priorities and recovery actions share the same business impact assumptions.

68 Saudi requirements add national context and enforceability. The Essential Cybersecurity Controls
69 and associated implementation guidance establish expectations for governance, defence, vulnerability
70 management, incident response, continuity and third-party security. Cloud, critical-system,
71 operational-technology and data controls extend these expectations to distributed infrastructure and
72 sensitive information [2-8]. Cloud service regulations also shape provider responsibilities and service
73 governance [9]. The challenge is therefore not a shortage of frameworks. It is the absence of an
74 integrated operating model that converts them into coherent decisions across a fast-moving
75 programme.

76 **2. Aim and Objectives**

77 The aim of this review is to construct an evidence-centred cyber-resilience framework for
78 distributed giga-project information technology in Saudi Arabia. The study has four objectives: to
79 identify the concentration and propagation risks created by distributed delivery; to integrate
80 information-security, continuity and crisis-management requirements around critical services; to define
81 hybrid-cloud containment and clean-recovery controls; and to propose measurable assurance gates for
82 incident response, restoration and continual improvement. The review is technology-neutral and treats
83 architecture, governance and operational behaviour as inseparable parts of resilience.

84 **3. Review Methodology**

85 A structured integrative review was selected because the topic spans standards, national controls,
86 engineering guidance and peer-reviewed research. The review covered sources issued between January
87 2020 and December 2025. Search concepts combined cyber resilience, critical infrastructure, hybrid
88 cloud, zero trust, business continuity, incident recovery, crisis management, supply-chain security,
89 Saudi cybersecurity controls and giga-project digital infrastructure. Priority was given to authoritative
90 standards, government publications, regulator documents and research that proposed, evaluated or
91 synthesised resilience mechanisms. The attached study informed the organisation of the review,
92 particularly its movement from contextual conditions to governance implications and an integrated
93 framework, but its 2026 evidence was not included in the cited evidence set.

94 Sources were retained when they met three conditions. First, they addressed at least one element of
95 the research problem: governance, service criticality, hybrid-cloud security, containment, continuity,
96 crisis coordination, recovery or assurance. Second, they provided traceable requirements, methods,
97 empirical findings or implementation guidance rather than unsupported commentary. Third, their
98 recommendations could be translated into decisions relevant to distributed programme environments.
99 Pure product marketing, undated material, duplicate editions and sources outside the defined period
100 were excluded. Thirty sources formed the final evidence base.

101 The evidence was coded into six domains. Governance evidence covered leadership, risk
102 ownership, policy, compliance and assurance. Service evidence addressed business impact,
103 dependency mapping and minimum acceptable operations. Architecture evidence covered zero trust,
104 network segregation, cloud controls, identity and management-plane protection. Continuity evidence
105 included recovery objectives, alternative operating modes and exercises. Incident evidence covered
106 detection, command, containment, communication and restoration. Ecosystem evidence addressed
107 suppliers, managed services, software provenance and concentration risk [10-26]. Coding was
108 iterative: new categories were retained only where they changed the proposed framework or its
109 decision criteria.

110 Synthesis used pattern matching rather than statistical aggregation because the sources differed in
111 scope and measurement. A proposed practice was treated as strong when it appeared across at least
112 two evidence families and could be connected to an observable test. For example, least-privilege
113 access is a security principle, but it becomes a resilience mechanism only when it limits the reach of
114 compromised credentials and preserves independent recovery administration [15,17,20]. Similarly, a
115 continuity plan is not considered effective because it is documented; it must be exercised against
116 realistic loss of identity, cloud, network, data or supplier services [13,14,21].

117 **4. Distributed Giga-Project Risk System**

118 The central resilience problem is dependency density. Individual applications may appear non-
119 critical, yet their shared dependencies can be programme-wide single points of failure. A site-access
120 platform may rely on a cloud identity provider, a mobile network, an integration gateway and a central
121 data service. A payment process may depend on enterprise resource planning, bank connectivity,
122 digital signatures and contractor records. A command centre may rely on telemetry from several
123 operational domains. Service mapping should therefore begin with outcomes and trace backwards
124 through applications, data, identities, networks, facilities, people and suppliers. This exposes hidden
125 concentration before an incident reveals it [13,17-19].

126 Change velocity amplifies this risk. Temporary connectivity, emergency firewall rules, rapid cloud
127 deployments and short supplier engagements can leave behind excessive trust. Assets may move from
128 construction to commissioning and then to operations without a corresponding change in ownership or
129 control baseline. Resilience governance needs lifecycle gates that define when a service becomes
130 production, when its data classification changes, when third-party access expires and when recovery
131 evidence becomes mandatory. The objective is to prevent temporary delivery exceptions from
132 becoming permanent attack paths.

133 Geographic distribution also creates asymmetric conditions. Remote sites may have limited
134 connectivity, environmental exposure, small support teams and dependence on a single carrier. Central
135 cloud services may be robust but unreachable during local disruption. Conversely, a local operational

system may continue safely while central dashboards fail. Continuity strategies should distinguish central, regional, site and asset-level autonomy. Each level requires a defined minimum operating mode, local decision authority and a controlled method for reconciling data after connectivity returns.

The supplier ecosystem is the fourth source of propagation. Integrators often hold broad access because they support several platforms, while software suppliers and managed service providers may share tools across clients. Supply-chain guidance emphasises risk knowledge, contractual controls, access restriction, incident obligations and verified provenance [25,28]. For giga-projects, the additional requirement is portfolio visibility: leaders must know where the same vendor, identity domain, remote-support tool, software component or recovery provider creates common-mode exposure across otherwise separate projects.

5. Integrated Governance Architecture

An integrated governance model should place critical services, rather than standards or technologies, at its centre. ISO 27001 establishes the management disciplines required to define scope, leadership, risk treatment, objectives, competence, monitoring and improvement [10]. ISO 22313 contributes the logic for business impact analysis, continuity strategies, procedures and exercising [13]. ISO 22361 clarifies strategic crisis leadership and communication [14]. The Saudi controls add national requirements and sector-specific expectations [2-8]. These sources can be combined through a single resilience register in which each critical service has an owner, impact tolerance, security classification, principal threats, minimum operating mode, dependency map, recovery design and evidence status.

The register should distinguish risk acceptance from continuity acceptance. A service owner may accept a residual vulnerability because remediation would delay delivery, but that decision does not automatically accept prolonged outage or data loss. Security risk, operational impact and recovery capability should be recorded separately and brought together at release gates. This separation prevents a common governance error in which a low probability score hides intolerable consequences or a strong backup claim hides an untested dependency.

Accountability should operate at three levels. The governing body defines risk appetite, national obligations and critical outcomes. A programme resilience council aligns security, technology, operations, continuity, safety, legal, communications and procurement. Service resilience teams maintain technical and operational evidence. The council does not replace incident command; it establishes the conditions under which incident command can act. These include delegated authority to isolate a site, disable federation, suspend a supplier connection, invoke manual operations or restore from a clean environment without waiting for normal change approval.

Control mapping should be performed once and reused. A single control objective, such as privileged-access protection, can satisfy information-security, cloud, critical-system and continuity expectations when evidence is designed appropriately. Evidence may include identity architecture, approval records, phishing-resistant authentication coverage, session logs, emergency-access tests, account recertification and proof that recovery administrators are independent from production administrators [2-5,10-12,15-16]. This reduces duplicated audits and encourages controls that work in practice rather than controls designed for one checklist.

Assurance should be risk-based and time-sensitive. High-change environments require continuous evidence for configuration, identity, vulnerability, backup and logging controls, while governance, crisis leadership and supplier responsibilities can be reviewed at defined intervals. Dashboards should not collapse resilience into one score. They should show exposure, service impact, containment capability, recoverability and evidence freshness separately. A green average can otherwise conceal an expired recovery test for a nationally significant service.

The governance architecture must also recognise the boundary between project and operational responsibility. During construction, project teams may own systems that later support public or industrial operations. Handover should therefore be treated as a resilience transition, not a documentation event. Acceptance criteria should include asset ownership, support capability, secure configuration, monitoring coverage, dependency records, recovery tests, supplier exit provisions and residual-risk approval by the future operator.

6. Hybrid-Cloud Security and Containment

Hybrid cloud should be designed as a set of bounded trust domains rather than a seamless pool of resources. The environment may include private cloud, public cloud, software-as-a-service, edge clusters, operational networks, site data rooms and partner environments. Each domain has different administrators, telemetry, failure modes and legal constraints. Zero-trust architecture removes the assumption that network location establishes trust and requires explicit evaluation of identity, device, workload, transaction and policy context [15,20]. In a giga-project, this principle should apply equally to users, service accounts, application programming interfaces, automation pipelines and management tools.

Identity is the primary control plane. Federation simplifies collaboration, but a single compromised identity provider can become a programme-wide outage and compromise mechanism. Critical workloads should therefore have defined federation break-glass procedures, local recovery identities, hardware-backed credentials and separation between production and recovery administration. Privileged access should be time-limited, purpose-bound and recorded. Machine identities require the same discipline: certificate issuance, secret rotation, workload identity and key-management dependencies must be included in service maps and exercises [11,15-17].

204 Network segmentation provides the second containment layer. Management, security tooling, user
205 access, application tiers, data services, backup, operational technology and supplier access should not
206 share unrestricted reachability. Segmentation must be based on service intent and data sensitivity
207 rather than only on site or address range. Policy should be tested with positive flows, negative probes
208 and simulated credential compromise. Network infrastructure guidance also supports management-
209 plane isolation, secure routing, hardened administration and visibility of anomalous paths [19,23]. The
210 objective is not maximal fragmentation; it is a bounded blast radius that still supports required
211 operations and recovery.


212 Cloud-native controls provide the third layer. Separate accounts or subscriptions, policy guardrails,
213 controlled interconnection, encryption, workload protection, logging and immutable configuration
214 history create organisational and technical boundaries. The Cloud Cybersecurity Controls and Cloud
215 Controls Matrix emphasise shared responsibility, tenant governance, data protection, identity,
216 monitoring, portability and provider assurance [3,26]. A cloud service should not be approved only
217 because the provider holds certifications. The programme must understand which resilience activities
218 remain its responsibility, how evidence will be obtained and what alternatives exist if the service or
219 management plane is unavailable.

220 Data resilience requires more than replicated storage. Replication can rapidly reproduce corruption,
221 malicious deletion or unsafe configuration. Critical datasets need classification, integrity checks,
222 retention rules, protected recovery points and independent administration [7,8,16]. Recovery copies
223 should be logically or physically separated from production credentials and should support restoration
224 to an environment that can be trusted. The programme should identify the authoritative source for each
225 dataset during normal operation, degraded operation and recovery. This prevents conflicting writes and
226 uncontrolled reconciliation after a site or cloud partition.

227 Observability is the fourth containment layer. Logs, flow records, endpoint telemetry, cloud events,
228 identity signals and application health must be normalised sufficiently to support cross-domain
229 detection. Central visibility is important, but local sites need enough monitoring and response
230 capability to operate when central services are unreachable. Detection rules should consider
231 programme context, such as unusual vendor access across several projects, rapid creation of cloud
232 privileges, unexpected data transfer between construction and operational zones, or changes to backup
233 retention. Threat-landscape evidence shows that ransomware, credential abuse, supply-chain
234 compromise and exploitation of exposed services remain persistent concerns [24].

235 Finally, hybrid-cloud resilience requires portability without creating uncontrolled duplication.
236 Critical services should define which components can fail over, which can operate manually, which
237 require a second provider and which should remain deliberately single-homed because distributed
238 consistency would create greater risk. Architecture decisions should be linked to impact tolerance and

239 tested recovery capability, not to a general preference for multi-cloud. Complexity is itself a failure
240 source; resilience is achieved when alternative paths are usable, governed and rehearsed.



241
242 **Figure 1. Integrated cyber-resilience architecture for distributed giga-project services, platforms**
243 **and recovery environments.**

244 **Table 1. Integrated control-to-evidence matrix**

Domain	Primary framework logic	Saudi alignment	Operational evidence	Accountable owner
Governance	ISMS scope, risk treatment, objectives and review	ECC governance and implementation guidance	Approved service register, risk decisions, evidence freshness	Governing body / CISO
Service continuity	Impact analysis, continuity strategy and exercises	ECC continuity and critical-system expectations	Impact tolerance, minimum mode, exercise results	Service owner / BCM
Identity	Least privilege, authentication and access lifecycle	ECC and CCC access controls	MFA coverage, privileged-session logs, break-glass test	IAM owner

Domain	Primary framework logic	Saudi alignment	Operational evidence	Accountable owner
Hybrid cloud	Shared responsibility, segregation, logging and portability	CCC and cloud-service regulations	Landing-zone guardrails, provider evidence, exit test	Cloud platform owner
Data integrity	Classification, protection, retention and restoration	DCC and national data standards	Immutable recovery points, integrity validation, authority matrix	Data owner
Incident and crisis	Detection, response, escalation, communications	ECC incident management and national obligations	Playbooks, contact channels, decision log, notifications	Incident commander
Supply chain	Due diligence, access restriction and incident clauses	ECC third-party and critical-system controls	Access expiry, software provenance, joint exercise evidence	Procurement / supplier owner

7. Business Continuity and Recovery Engineering

Business continuity begins with an explicit service impact tolerance: the maximum disruption the organisation can absorb before safety, legal, financial, delivery or reputational consequences become unacceptable. Recovery time and recovery point objectives are useful technical parameters, but they should be derived from this tolerance rather than selected independently. The business impact analysis should identify peak periods, manual alternatives, minimum staffing, data dependencies, external obligations and the sequence in which capabilities must return [13]. For a giga-project, priorities may change by phase; payroll, access control, commissioning records and public-facing services can have different criticality during construction, launch and operation.

Continuity strategies should provide several operating modes. Normal mode uses the full digital service. Degraded mode preserves the critical outcome with reduced functionality, limited users or delayed synchronisation. Manual mode uses controlled procedures and later reconciliation. Recovery mode restores technology in a trusted sequence. These modes should be designed with operational teams, because an apparently simple manual workaround may require forms, local authority, trained people, secure storage and communications that do not exist during a real disruption.

261 Recovery architecture should separate survival from restoration. Survival controls keep essential
262 services running or support a safe stop. Restoration controls rebuild affected services from known-
263 good components. Backups, golden images, infrastructure code, application packages, certificates and
264 configuration baselines must be protected as a recovery system. NIST cyber-resilience guidance
265 emphasises diversity, redundancy, segmentation, deception, non-persistence and substantiated integrity
266 as ways to preserve options under attack [17]. These techniques should be selected according to the
267 service threat model rather than applied uniformly.

268 Clean recovery is particularly important after credential theft or ransomware. Restoring servers into
269 the same compromised identity, network or management environment can recreate the incident. A
270 clean recovery environment should have independent access, trusted tooling, validated images,
271 restricted connectivity and a controlled path for data restoration. Recovery teams need criteria for
272 determining the last known-good point, validating restored data, rotating secrets and proving that
273 persistence has been removed. The decision to reconnect should be based on evidence from security,
274 application, data and business owners, not on infrastructure availability alone [16,17,21].

275 Exercises should progress from discussion to technical proof. Tabletop exercises test authority,
276 communications and assumptions. Functional exercises test specific procedures such as disabling
277 federation or invoking local operations. Recovery tests rebuild applications and data. Integrated
278 exercises combine cyberattack, supplier failure, site disruption and public communication. Crisis
279 guidance recommends strategic coordination across decision-makers and stakeholders [14]. For
280 distributed projects, exercises should also test language, time-zone, contractual and cross-company
281 barriers that can slow action even when technical procedures are sound.

282 Recovery sequencing should follow service dependencies. Identity, name resolution, time, keys,
283 network control, logging and integration platforms may need restoration before business applications.
284 However, central infrastructure should not be restored automatically if a site can continue safely in
285 isolation. The recovery plan should specify authoritative control planes, restoration order, validation
286 checkpoints and abort criteria. A dependency-aware sequence reduces the risk of technically
287 successful but operationally unusable recovery.

288 Continuity evidence should be retained as an assurance record. It includes the approved impact
289 analysis, recovery design, exercise scope, actual restoration times, achieved data points, unresolved
290 defects and owner acceptance. Evidence must show not only that a test completed, but that the
291 intended business outcome was available and trustworthy. This turns continuity from an annual
292 documentation activity into a measurable engineering discipline.

8. Incident Command and Coordinated Recovery

Incident response and business continuity should share triggers but retain distinct responsibilities. Security teams determine the nature and extent of compromise, preserve evidence and direct technical containment. Technology teams maintain platforms and execute restoration. Business teams select degraded operations and priorities. Crisis leaders manage strategic decisions, communications, legal obligations and external coordination. The incident playbooks should identify when these streams converge, who has final authority and how decisions are recorded [14,21,29-30].

A distributed incident should be managed through a common operating picture. This does not require every organisation to use the same tool, but it requires agreed facts: affected services, sites, identities, data, suppliers, operational consequences, containment status, recovery readiness and decision deadlines. Separate technical tickets are insufficient because they hide cross-project impact. A resilience coordinator should maintain the service view while specialised teams maintain detailed evidence.

Containment decisions must consider continuity consequences. Disconnecting a site may stop lateral movement but interrupt safety monitoring or access control. Disabling federation may protect cloud services but block emergency users. The response team should therefore use pre-agreed containment envelopes. An envelope defines which communications can be cut immediately, which require service-owner approval, which local capabilities must remain and what evidence is required before reconnection. Preparation allows rapid action without improvising under pressure.

Communication is an operational control. Staff, contractors, regulators, customers and the public may require different messages, but all communications must be based on a consistent incident assessment. Crisis-management guidance emphasises timely, credible and coordinated communication [14]. Contact lists and channels should remain available when corporate email or collaboration platforms are affected. Supplier contracts should define notification times, evidence-sharing requirements and participation in joint response.

Recovery should be treated as a series of controlled authority transfers. A restored identity service becomes authoritative only after security validation. A recovered database accepts writes only after integrity and reconciliation checks. A site rejoins the network only after endpoint, route and policy evidence is approved. These gates prevent premature restoration driven by schedule pressure. The final gate is business validation: users confirm that the service performs its critical outcome and that data is complete enough for controlled operation.

After-action review should distinguish detection, containment, continuity and recovery performance. A rapid detection may still be followed by slow business decisions; a fast restore may still use an untrusted image. Lessons should be linked to owners, deadlines and evidence of closure.

327 Recurrent findings across projects should change enterprise architecture, supplier requirements or
328 programme governance rather than remain local corrective actions.

Incident-to-recovery assurance loop

Every transition is authorised by service impact and verified evidence



329

330 **Figure 2. Evidence-centred incident-to-recovery assurance loop for distributed operations.**

331 **Table 2. Resilience metrics and indicative decision bands**

Metric	Purpose	Green	Amber	Red / action
Critical-service dependency coverage	Expose hidden concentration	$\geq 95\%$ current	85-94% ageing or	$< 85\%$; halt major release
Privileged phishing-resistant MFA	Limit credential propagation	$\geq 98\%$	95-97%	$< 95\%$; remove or restrict access
Negative segmentation test pass rate	Verify denied paths	100%	98-99% with bounded defect	$< 98\%$; isolate affected domain
Protected recovery-point success	Confirm recoverable data	100% scheduled success	One failed copy with alternative	No verified clean point; escalate

Metric	Purpose	Green	Amber	Red / action
Clean-environment recovery test	Prove trusted restoration	Within approved tolerance	Within 120% of tolerance	Beyond 120% or incomplete
Mean time to contain high-impact event	Measure blast-radius control	Within scenario target	Up to 150% of target	Beyond 150%; redesign playbook
Business validation after restoration	Confirm critical outcome	All critical checks pass	Controlled workaround accepted	Critical check fails; do not reconnect
Evidence freshness	Prevent obsolete assurance	Within defined cycle	Up to 30 days overdue	More than 30 days overdue

9. Assurance Metrics and Implementation Roadmap

Resilience metrics should measure capability and consequence. Governance indicators include the proportion of critical services with current owners, impact tolerances and dependency maps. Exposure indicators include privileged-account age, supplier-access exceptions, internet-facing assets and configuration drift. Containment indicators include the number of trust zones reachable from a compromised identity, time to revoke access and negative segmentation test results. Recovery indicators include clean-environment readiness, achieved restoration time, achieved recovery point and business validation success. Evidence freshness should be visible because a recovery test from an earlier architecture may no longer be meaningful [18,22].

Implementation can proceed in four stages. The first stage establishes governance and identifies critical services. The second maps dependencies and creates minimum operating modes. The third builds hybrid-cloud boundaries, independent recovery administration, protected data and cross-domain monitoring. The fourth exercises integrated scenarios and converts results into release gates. This sequence avoids the common mistake of purchasing recovery technology before deciding which outcomes must be preserved.

The roadmap should prioritise shared services that create portfolio-wide concentration: identity, connectivity, cloud landing zones, security tooling, integration, backup and supplier remote access. Improvements in these areas produce resilience benefits across multiple projects. Local services can then be addressed according to impact tolerance. Each project should receive a minimum control

baseline, but higher-criticality services should have stronger segmentation, more independent recovery and more frequent exercises.

The first figure presents the layered architecture, while the second shows the incident-to-recovery assurance loop. Table 1 translates the principal frameworks into operational evidence. Table 2 provides indicative metrics and decision bands. Thresholds are starting points; they require calibration to project phase, service criticality and regulatory obligations.

10. Discussion and Research Agenda

The synthesis indicates that cyber resilience is an organisational capability expressed through architecture and operations, not a separate technology domain. ISO 27001 can govern risk and assurance, but it does not by itself determine the order in which a compromised service should recover. ISO 22301-aligned continuity can define priorities, but it does not by itself contain malicious activity. Hybrid-cloud controls can isolate workloads, but they cannot decide which business outcome has precedence. The proposed framework links these decisions around a common service model.

Several tensions remain. Centralisation improves standardisation and visibility but increases concentration risk. Multi-cloud and regional autonomy create alternatives but increase complexity and inconsistency. Automation reduces manual error but can propagate faulty policy quickly. Fine-grained segmentation limits lateral movement but can delay troubleshooting and restoration. These tensions should be resolved with tested evidence rather than universal rules. The correct architecture is the one that preserves the required outcome under credible failure conditions.

A mature programme should also maintain a resilience design authority that reviews significant changes before they create shared exposure. The authority should include architecture, security, continuity, operations, safety, data, procurement and legal expertise. Its purpose is not to centralise every decision, but to challenge assumptions about common dependencies, recovery independence, supplier concentration and evidence quality. Reviews should focus on services with low impact tolerance, broad administrative reach, sensitive data or dependencies used by several projects. Decisions should be documented as reusable design principles, approved patterns and explicit exceptions with owners and expiry dates. This creates institutional memory that survives contractor turnover, project reorganisation and the transition from construction to operations. It also enables leaders to compare planned resilience with observed incident and exercise performance. Where a control repeatedly fails, the design authority can determine whether the cause is technology, process, staffing, authority, supplier behaviour or unrealistic service expectations. Findings should be converted into architecture changes, revised contract clauses, new monitoring requirements or additional exercises rather than remaining as isolated actions. The authority should publish a concise quarterly view of portfolio concentration, overdue recovery evidence, recurring exceptions and unresolved cross-project risks. This learning function is essential because distributed programmes evolve faster

387 than static baselines. Resilience must therefore be governed as a living capability, refreshed whenever
388 architecture, suppliers, operational purpose, data sensitivity or threat conditions materially change. Its
389 annual work plan should prioritise shared identity, connectivity, cloud, integration, security monitoring
390 and recovery services because weaknesses in these platforms can propagate rapidly across the entire
391 digital portfolio.

392 Research should examine giga-project resilience longitudinally across design, construction,
393 commissioning and operation. Useful studies would quantify dependency growth, supplier
394 concentration, recovery performance and the effect of governance maturity on incident outcomes.
395 Further work should test cross-project scenarios involving identity compromise, cloud control-plane
396 outage, ransomware, site isolation and simultaneous physical disruption. Saudi-specific research can
397 also evaluate how national control requirements can be represented as machine-readable evidence
398 without reducing assurance to automated compliance.

399 **11. Conclusion**

400 Distributed giga-project information technology requires a resilience model designed for rapid
401 change, shared platforms, remote sites and complex delivery ecosystems. The proposed framework
402 integrates information-security governance, continuity management, crisis leadership, hybrid-cloud
403 containment and clean recovery around critical service outcomes. It replaces fragmented compliance
404 with one resilience register, one dependency model and evidence-based release and restoration gates.

405 The practical test of the framework is straightforward: when compromise occurs, the programme
406 must know what matters, what depends on it, what can be isolated, how essential operations will
407 continue, which recovery source can be trusted and who may authorise restoration. Saudi national
408 controls and international guidance provide the necessary foundations. Their value is realised when
409 they are engineered into observable behaviour, exercised across organisational boundaries and
410 improved after every material change or incident.

411 **References**

- 412 [1] Kingdom of Saudi Arabia. Vision 2030 Annual Report 2024. Riyadh; 2025.
- 413 [2] National Cybersecurity Authority. Essential Cybersecurity Controls (ECC 2:2024). Riyadh; 2024.
- 414 [3] National Cybersecurity Authority. Cloud Cybersecurity Controls (CCC 2:2024). Riyadh; 2024.
- 415 [4] National Cybersecurity Authority. Guide to Essential Cybersecurity Controls Implementation
416 (GECC-1:2023). Riyadh; 2023.
- 417 [5] National Cybersecurity Authority. Guide to Cybersecurity Controls for Critical Systems
418 Implementation (GCSCC-1:2023). Riyadh; 2023.
- 419 [6] National Cybersecurity Authority. Operational Technology Cybersecurity Controls (OTCC-
420 1:2022). Riyadh; 2022.

421 [7] National Cybersecurity Authority. Data Cybersecurity Controls (DCC-1:2022). Riyadh; 2022.

422 [8] Saudi Data and Artificial Intelligence Authority, National Data Management Office. Data
423 Management and Personal Data Protection Standards. Riyadh; 2021.

424 [9] Communications, Space and Technology Commission. Cloud Computing Services Provisioning
425 Regulations. Riyadh; 2023.

426 [10] International Organization for Standardization. ISO/IEC 27001:2022 Information security,
427 cybersecurity and privacy protection - Information security management systems - Requirements.
428 Geneva; 2022.

429 [11] International Organization for Standardization. ISO/IEC 27002:2022 Information security,
430 cybersecurity and privacy protection - Information security controls. Geneva; 2022.

431 [12] International Organization for Standardization. ISO/IEC 27005:2022 Information security,
432 cybersecurity and privacy protection - Guidance on managing information security risks. Geneva;
433 2022.

434 [13] International Organization for Standardization. ISO 22313:2020 Security and resilience - Business
435 continuity management systems - Guidance on the use of ISO 22301. Geneva; 2020.

436 [14] International Organization for Standardization. ISO 22361:2022 Security and resilience - Crisis
437 management - Guidelines. Geneva; 2022.

438 [15] Rose S, Borchert O, Mitchell S, Connelly S. Zero Trust Architecture. NIST Special Publication
439 800-207. Gaithersburg, MD; 2020. <https://doi.org/10.6028/NIST.SP.800-207>.

440 [16] Joint Task Force. Security and Privacy Controls for Information Systems and Organizations.
441 NIST Special Publication 800-53 Revision 5. Gaithersburg, MD; 2020.
442 <https://doi.org/10.6028/NIST.SP.800-53r5>.

443 [17] Ross R, Pillitteri V, Graubart R, Bodeau D, McQuaid R. Developing Cyber-Resilient Systems: A
444 Systems Security Engineering Approach. NIST Special Publication 800-160 Volume 2 Revision 1.
445 Gaithersburg, MD; 2021. <https://doi.org/10.6028/NIST.SP.800-160v2r1>.

446 [18] National Institute of Standards and Technology. The NIST Cybersecurity Framework (CSF) 2.0.
447 Gaithersburg, MD; 2024. <https://doi.org/10.6028/NIST.CSWP.29>.

448 [19] Chandramouli R, Butcher Z. Guide to a Secure Enterprise Network Landscape. NIST Special
449 Publication 800-215. Gaithersburg, MD; 2022. <https://doi.org/10.6028/NIST.SP.800-215>.

450 [20] Cybersecurity and Infrastructure Security Agency. Zero Trust Maturity Model, Version 2.0.
451 Washington, DC; 2023.

452 [21] Cybersecurity and Infrastructure Security Agency. Federal Government Cybersecurity Incident
453 and Vulnerability Response Playbooks. Washington, DC; 2021.

454 [22] Cybersecurity and Infrastructure Security Agency. Cross-Sector Cybersecurity Performance
 455 Goals, Version 1.0.1. Washington, DC; 2023.

456 [23] National Security Agency. Network Infrastructure Security Guidance. Fort Meade, MD; 2022.

457 [24] European Union Agency for Cybersecurity. ENISA Threat Landscape 2024. Athens; 2024.

458 [25] European Union Agency for Cybersecurity. Good Practices for Supply Chain Cybersecurity.
 459 Athens; 2023.

460 [26] Cloud Security Alliance. Cloud Controls Matrix, Version 4.0. Seattle, WA; 2021.

461 [27] Tzavara V, Vassiliadis S. Tracing the evolution of cyber resilience: a historical and conceptual
 462 review. *International Journal of Information Security*. 2024;23:1695-1719.
 463 <https://doi.org/10.1007/s10207-023-00811-x>.

464 [28] Roshanaei M. Resilience at the core: critical infrastructure protection challenges, priorities, and
 465 cybersecurity assessment strategies. *Journal of Computer and Communications*. 2021;9(8):80-102.
 466 <https://doi.org/10.4236/jcc.2021.98006>.

467 [29] Fysarakis K, Lekidis A, Mavroeidis V, et al. PHOENIX: a European cyber resilience framework
 468 with artificial-intelligence-assisted orchestration, automation and response capabilities for business
 469 continuity, incident response and information exchange. 2023.
 470 <https://doi.org/10.48550/arXiv.2307.06932>.

471 [30] Papastergiou S, Mouratidis H, Kalogeraki EM. Cyber security incident handling, warning and
 472 response system for the European critical information infrastructures (CyberSANE). *Proceedings of*
 473 *the 6th International Conference on Information Systems Security and Privacy*. 2020:837-843.