CYBER RESILIENCE

Cyber Resilience in Multi-Vendor Enterprise Networks: A Framework for Secure Integration, Threat Prevention, and Operational Continuity in Saudi Arabia

Ataur Rahman Atif Mohammed

Published November 2026

Abstract

Enterprises in SA are increasingly reliant on networks which are a mix of equipment from a number of vendors, cloud platforms, security products and managed services. The benefit of this is clearly in terms of procurement flexibility but at the cost of complexity in the areas of identity enforcement, configuration assurance, incident attribution and service restoration. This is a critical integrative review of how interoperability can be achieved without the assumption that the various products in the mix will have the same policy semantics or the same level of operational visibility. A literature review from 2020 to 2025 is combined with relevant engineering guidance in a transparent, theme-based synthesis. There is a clear distinction drawn between preventive cybersecurity and cyber resilience (the ability to anticipate, absorb, adapt to and recover from a disruptive event whilst maintaining the integrity of the essential services). The review identifies five interdependent capabilities: authoritative asset and dependency mapping, vendor neutral policy translation, identity based access control, correlated threat telemetry and independently testable continuity arrangements. On that basis it goes on to propose a layered Secure Integration and Continuity Framework which links enterprise governance to technical controls and to measurable service outcomes. Particular attention is paid to integration interfaces, admin privileges, outsourced maintenance, control plane concentration and recovery dependencies. The framework is in the context of Saudi regulatory expectations and the operational realities of the very rapidly expanding range of digital services. It does not claim to have demonstrated any performance improvements in practice, it simply sets out a number of falsifiable propositions and evaluation metrics for future field studies. The core proposition is that diversity can only become a resilience enabler to the extent that an organisation can demonstrate equivalent security outcomes across the vendor mix and can restore the essential services when the shared coordination systems have failed.

Keywords: cyber resilience; multi-vendor networks; zero trust; secure integration; operational continuity; Saudi Arabia

1 Introduction

Saudi enterprise networks almost never come out as a single, coherent single vendor solution. You end up with a mix as a result of acquisitions, phased modernisation, specialist operational requirements and cloud. You end up with a mix of routers, switches, firewalls, wireless controllers, application gateways, identity directories and management platforms. The advantage of that mix is that you're not reliant on one vendor. The disadvantage is that you end up with inconsistent interfaces, overlapping roles and responsibilities and configurations which are extremely difficult to compare. Your network may look perfectly segmented in an architecture diagram with no apparent paths between the various services, but there may very well be an undocumented management path

which allows lateral movement between those services. The real risk is therefore not the diversity of the technology per se, it's the gap between what you intend to protect and what you actually do protect across the independently managed components.

The stakes are far higher than just data confidentiality. A compromised privileged account could bring a payment workflow to a complete standstill. A bad routing update or a failed policy deploy could take a hospital service offline even if the application data is absolutely fine. Cyber resilience is about the ability to maintain and restore mission capability after controls have been bypassed or have failed. Ross et al. (2021) identify anticipation, resistance, recovery and adaptation as the engineering concerns. AlHidaifi et al. (2024) make it very clear why availability and recovery need to be explicitly included in resilience measurement. Prevention and continuity therefore need to be part of a single operational design rather than two separate programmes run by different teams.

Zero trust research is all about challenging the trust we've got from the concept of being 'inside' an internal network. Rose et al. (2020) make their access decisions on the basis of subject and resource. Syed et al. (2022) go on to document the practical problems of trying to implement that principle in real architectures. Multi-vendor estates add another layer of complexity in this respect: how do you enforce one intent when the various products have their own interpretation of identity, session state, routing zones and policy precedence? The controls may be perfectly sound in theory but they can end up in conflict when you translate them through the various vendor APIs. Ribeiro et al. (2022) point to security intent extraction from existing device configurations as a very promising approach to this very issue of management.

A crucial distinction to be made is between heterogeneity which diversifies failure and heterogeneity which conceals failure. The former is where the various technologies can carry the essential traffic through different failure modes. The latter is where the operator has no way of knowing whether the protections are actually equivalent or whether the recovery paths are in fact independent. The architectural decision therefore has to be based on common dependencies and assurance evidence rather than simply the number of manufacturers. This distinction also obviates the fallacy that compliance with a common procurement specification in any way guarantees continuity in the event of a coordinated incident.

Saudi Arabia is a very good example in that respect as you've got digital public services, industrial connectivity, financial platforms and all sorts of large scale development projects which all operate within extremely complex supplier ecosystems. The requirements from the NCA are very much around governance, defensive controls, resilience and the obligations you have to your third parties. That said, no international guidance in any way whatsoever replaces the need for a bespoke legal and regulatory assessment of the organisation as the scope of application will be very much different from sector to sector, from system classification to system classification, from contractual obligation to contractual obligation and from location to location.

The body of work that exists is of course very useful in that it provides some good answers. There's a lot of work on network automation (intent to action), a lot of work on intrusion detection (anomalous traffic) and a lot of work on continuity (the recovery of services). The degree to which these can be combined into a testable architecture for the operation of a heterogeneous enterprise is not well developed. This paper therefore sets out to answer the question of how an organisation can make consistent security decisions across a range of vendors, how it can detect a failure prior to it becoming systemic and how it can restore services when its normal management tools are them-

selves not available. Its contribution is a framework which integrates technical assurance with service based continuity rather than a claim of superior detection capability which is as yet unproven.

2 Review aim, objectives and methodology

The study is about producing an evidence-based architectural explanation of cyber resilience in multi-vendor enterprise networks in Saudi Arabia. The four objectives that will form the basis of the synthesis are: to identify the security failures that arise from heterogeneous interfaces and administrative domains; to assess the controls that provide cross-vendor assurance rather than just product specific assurance; to link the detection and response capabilities to the recovery of critical services; and to define an evaluation model which can be tested in a real enterprise environment. The scope of this is of course much more targeted than a generic survey of all the available network security technologies.

A critical integrative review was used as the evidence base as the range of studies is such that they are a mixture of technical experiments, conceptual models, standards and organisational work and their results cannot be aggregated into a single effect size. The searches were built around combinations of multi-vendor networking, network configuration, zero trust, software defined networking, third party cybersecurity, supply chain assurance, incident response and cyber resilience. The focus was on English language publications from 2020 onwards. Publisher records and DOI publications were checked for bibliographic identity, topical relevance and date of publication. Authoritative technical documents were included where a design principle required a formal definition or implementation guidance.

Eligibility was on the basis of direct explanatory value to at least one of four questions: interface trust and configuration consistency; exposure and threat propagation; coordinated response and recovery; or assurance and governance. Studies which simply described consumer privacy, described cryptographic algorithms in isolation from the context of enterprise operations or reported performance results with no useful security implication were not included in the core synthesis. The review contains both conceptual and methodological evidence and the strength of a survey is not equated with that of a controlled deployment. No claim is made that the reference set that we end up with is a complete census of the work in that period.

A structured extraction matrix was employed to aid the synthesis. For each of the sources selected the matrix captured the publication year, the proposed technical mechanism, the assumed trust boundary, the validation setting, the principal limitation and the relevance to a Saudi multi-vendor scenario. Where a paper contained a lab demonstration the interpretation was restricted to the topology and assumptions of that demonstration. The normative security documents were used to inform the control requirements but were not treated as experimental corroboration. This approach helped to avoid the risk of equating a recommendation, a conceptual taxonomy and an observed outcome with the same level of evidence.

The analysis was done in four phases. Firstly each publication was coded at the level of its primary unit of analysis (device, identity, flow, software component, service or organisation). Secondly its proposed mechanism was mapped to one of the five functions (prevention, detection, containment, restoration, learning). Thirdly the implementation dependencies were captured (need for shared identifiers, vendor APIs, skilled operators, independent recovery infrastructure etc). Fourthly the contradictions were examined, particularly in the case where automation can be extremely effective but creates new concentration of risk. Leivadeas and Falkner (2023) have a description of closed loop

intent assurance and Mubarkoot et al. (2023) have a discussion of the requirement for independently verified evidence in the case of compliance requirements and policies.

The synthesis is interpretive rather than statistical (no study weights, no pooling of incident rates, no invented number of screened records). Evidence quality is discussed in terms of whether a source has a formal control definition, comparative empirical data, an implementation demonstration or a conceptual argument. Findings are only converted into architectural propositions where the boundary conditions of the finding can be stated. This is entirely appropriate for a review which is intended to produce a falsifiable framework. The absence of a systematic database audit of course means that we can't reach any firm conclusions about the frequency of publication or the extent of research coverage.

3 Heterogeneity as a security and continuity problem

A multi-vendor environment is not just about different brands of physical hardware. The operational diversity is such that you can have different configuration languages, different auth protocols, different software versions, different telemetry schemas, different SP responsibilities and different support models. Each interface can be perfectly valid in its own right and yet the overall picture can result in completely unexpected access paths. Security policy in particular is prone to semantic drift in that a policy which is stated to be blocking admin access may well be using different address objects, different processing order, different logging defaults or different failure behaviour on different enforcement points. Ribeiro et al. (2022) make the point that it is essential to translate the existing configuration into a higher level of security intent prior to any form of consistent assurance being possible.

The management plane is a very common source of correlated exposure. Centralised automation systems will invariably require elevated privileges on a very large number of devices. A compromise of the management plane can turn a tool that was intended to reduce configuration errors into a very rapid mechanism for the distribution of bad changes. Melis et al. (2023) look at the offensive and defensive aspects of SDN and point out that the very nature of programmability creates both security opportunities and attack opportunities. The solution is to have a good architectural design which separates policy approval from policy execution, limits the identities of the automation and provides for an independent review of proposed changes.

Third party dependence expands the trust boundary. Your maintenance contractors, your firmware distributors, your cloud service providers and your integration partners can all have direct access into your production networks. Collier and Sarkis (2021) make it clear that you should be managing supply chain risk without assuming any level of trust in your partners. Boyens et al. (2022) take a similar view, placing supply chain risk in the context of enterprise wide governance rather than as some kind of procurement tick box. The practical outcomes of this are time bounded access, software provenance, partner specific escalation paths and clear ownership of the evidence during an incident.

Visibility collapses when every product reports an incident in its own language. The timestamps will be different, the address translation will hide the source and the event categories will be completely incompatible. Wang et al. (2023) show that encrypted traffic analysis can in fact recover useful behavioural signals, though the success of that analysis is very much a function of the representation, the data context and the evaluation conditions. Landauer et al. (2020) point to the same issues with the clustering of heterogeneous security logs. A central dashboard is not going to fix it. Trustworthy

correlation requires asset identity, synchronised time, provenance, field mapping and data quality checks.

Operational logs also point to a governance issue in that retention periods, classification and access controls can differ from supplier to supplier resulting in a very fragmented view of what actually occurred. Cascavilla et al. (2021) speak to the problem of the fragmentation of the cybercrime intelligence sources and the importance of preserving provenance and context as evidence moves across organisational boundaries. Enterprises should specify a minimum event contract which would cover the timestamp, the device ID, the actor ID, the configuration version and any collection failures. The contract should be practical enough to be applicable to older equipment and stringent enough to allow for a full reconstruction of the incident.

Continuity adds another layer of coupling. You can have access links from two different providers in the same physical site. You can have multiple backup repositories all on the same directory service. You can have different firewall vendors all on the same orchestration platform. Apparent redundancy is therefore of absolutely no value. You have to map the critical dependencies through power, facilities, network transit, name resolution, credentials, encryption keys and operator access. It's not about the number of suppliers. It's about whether a failure could take out multiple apparently independent recovery paths.

4 Secure integration and continuity framework

The proposed Secure Integration and Continuity Framework is based on five capabilities with governance and independent assurance running through all of them. The first of these is a living inventory of devices, software, interfaces, administrators and the business service dependencies. Discovery is achieved through a mixture of network observation, configuration repository data and owner attestations (again none of those on their own is going to be able to reliably distinguish a legitimate gateway from a long forgotten one). Each critical service is then linked to the network paths, the identities and the supplier contracts that it requires in order to operate. This enables controls to be prioritised on the basis of service consequence rather than simply on the number of vulnerabilities.

The second capability is to take the security intent and turn it into vendor specific enforcement, whilst still preserving the verifiable invariants. The policy owner defines the outcomes (e.g. administrators have restricted reachability, certain application flows are permitted, backup admin is separated). The adapter layer then takes those outcomes and translates them into the native config. The translation has to be tested on a representative set of devices, across different software releases as the rule precedence or protocol handling may have changed in those releases. Leivadeas and Falkner (2023) make the point about the need for assurance throughout the intent lifecycle and Mehmood et al. (2023) make a similar point that intent processing requires validation rather than just simple auto-deployment. A translation that compiles but allows an unintended path has not done its job.

The third capability is identity centric access control applied to users, workloads, service accounts and maintenance channels. Continuous verification doesn't mean you're going to be re-authenticating on every single transaction; it means that the authorisation is responsive to meaningful changes in identity, device posture, the resource being requested and the level of risk. Rose et al. (2020) have the foundational model and Chandramouli and Butcher (2023) take that and extend the access enforcement to cloud native applications across locations. Policies should be evaluated at both the service

and network boundary. There should be a separately governed, strongly authenticated emergency admin path which is viable in the event that the normal identity infrastructure is down.

The identity layer should also be able to tell the difference between a human authorising something and a machine trusting something. Workload credentials can be around for a very long time and can traverse the entire deployment pipeline. Emergency accounts can in fact go around all the normal controls. Alevizos et al. (2022) look at a range of approaches to improving endpoint trust but further crypto mechanisms won't be able to overcome poor privilege boundaries. Credential inventories should contain who owns the credential, what interfaces it's allowed to use, its rotation schedule and its emergency revocation procedure. A service identity that's used to read monitoring data for instance should not also be one that can change the firewall policy.

The 4th capability is the integration of structured telemetry with threat containment. Instead of assuming that all devices will generate the same level of security events the collectors record the source reliability, collection latency, identity resolution and parsing confidence. The correlation rules then look at changes in privilege, changes in configuration, east-west traffic and service availability. Katsis et al. (2021) look at the impact of policy semantics and policy testing on whether or not the actual access paths end up conforming to the policy. Of course the detection decision has to be human accountable in that you don't want an automated block to be the thing that stops an emergency or revenue critical service. The containment action can range from simply disabling a vendor support token all the way up to the point of isolating a network segment with pre-approved exceptions being documented prior to the incident.

The fifth capability is about treating recovery as an architectural property. You need immutable or otherwise protected configuration baselines, offline restoration procedures and tested alternate management access that are independent of the components which are most likely to fail in combination. Stamenkov (2022) makes a distinction between layered continuity and recovery plans with the recovery plan providing a clear link between the enterprise level crisis decisions and the technical recovery. The recovery sequence should put identity, routing, name services and the key application dependencies before the lower priority conveniences. Backup success is of no guarantee of recoverability until you've demonstrated integrity, restoration authority and service functionality in an exercise.

Governance closes the loop. Policy exceptions, supplier changes and open gaps are all logged against named risk owners. Assurance outcomes feed into the subsequent refinement of the inventory and deployment standards. The framework is deliberately vendor neutral (not vendor agnostic) as its efficacy is predicated on knowing where the native behaviour of a product cannot achieve the required invariant. Figure 1 shows the relationship between the five capabilities and their individual evidence streams.

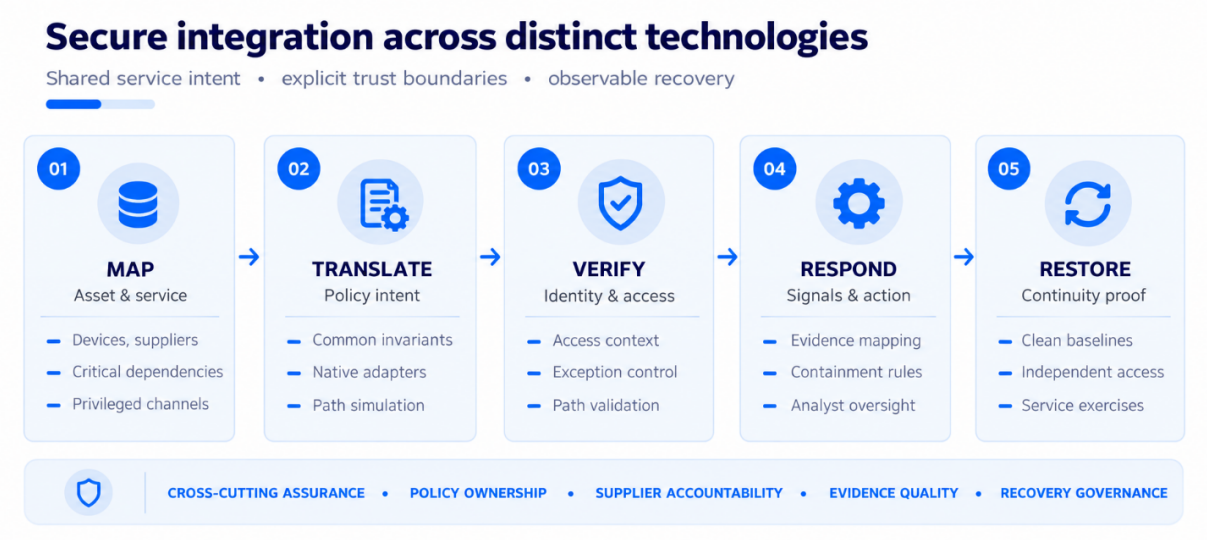


Figure 1. Proposed Secure Integration and Continuity Framework. The diagram is an original conceptual synthesis; stages are capabilities rather than a prescriptive product stack.

5 Threat prevention, containment and continuity engineering

Prevention starts with the eradication of unmanaged trust. Access to control interfaces must be via dedicated management paths, using authenticated service identities and the principle of minimal roles. Administrative actions must be traceable to individuals even when the technicians are using shared operational tooling. Buck et al. (2021) outline both the conceptual advantages and the practical difficulties in implementing zero trust. Teerakanok et al. (2021) go on to explain why the migration to zero trust must take into account the existence of legacy assets and the operational constraints that they impose. In a heterogeneous estate it is rarely practical to simply replace all the incompatible equipment immediately; compensating controls in such a situation should be documented with an expiry date and owner.

Software integrity is another equally important preventative boundary. An approved device configuration can still be compromised by a vulnerable controller plugin or an unaudited integration script. Scarfone et al. (2022) recommend secure development practices throughout the software lifecycle. Reichert and Obelheiro (2024) point to supply chain blind spots in software distribution and in the way dependencies are managed. Enterprise teams should have a full inventory of their firmware, application packages, connectors and supporting libraries and then link the vulnerability notifications to the actual versions that are in use. Third party provenance attestations can go some way to reducing uncertainty but they don't replace the need for behavioural monitoring of privileged supplier connections.

Threat prevention is much more believable when the control requirements can be tested continuously. Appropriate test cases would be a failed authentication from an unmanaged device, a blocked lateral movement between sensitive network zones, a rejected unauthorised policy change and proof that an emergency revocation has reached all the relevant gateways. He et al. (2022) look at the role of identity, access control and trust assessment; Cao et al. (2024) go on to discuss the orchestration and automation challenges. In the case of multi-vendor networks the acceptance criterion should be based on observable behaviour along a traffic path rather than on whether or not a vendor specific configuration command was actually issued.

Detection has to deal with encryption, inconsistent event formats and limited inspection capacity. Wang et al. (2023) point out that analysis of encrypted traffic is highly dependent on the analytical objective and the evaluation methodology. Umer et al. (2022) demonstrate how network intrusion signals can be separated from process level anomalies in industrial systems which is a very pertinent point for any organisation with operational technology connected to its corporate systems. Anomalies in traffic must be considered in the context of identity activity, configuration history and application health. A high volume of alerts is no indication of a good detection system if the analysts are unable to link the alerts to the relevant services or identify who has the authority to take action to contain the issue.

Containment decisions have to take into account blast radius. Disconnecting a compromised site may stop the spread but could also remove a key telemetry feed or cut off a vital remote operation. Site specific playbooks should define the permitted emergency actions, the business approval thresholds and the alternatives for services which we simply can't stop safely. Policy controllers should be able to support a staged rollout, maintain transaction logs and provide for an automatic rollback in the event that pre-defined guardrails are breached. Rollback isn't always safe though. Rolling back to a previous configuration may in fact re-introduce the very vulnerability that we were trying to contain. A human incident commander will be required to resolve such conflicts using the validated service dependencies and current threat intel.

Recovery begins prior to technical restoration. The response team need to establish a credible operating picture and have a clean control channel in place (particularly if there's any suspicion of compromise of the normal orchestration system). Ross et al. (2021) make a distinction between recovery and adaptation whereas Nelson et al. (2025) embed response into the ongoing risk management process rather than as a discrete emergency sequence. The requisite preps would be secure stored device configurations, credential re-issue procedures, supplier contact trees, validated software images and the documented prerequisites for re-connecting the various segments. The recovery team will need to validate the integrity of the route, the identity of the endpoint, the application transactions and the logging prior to declaring the service restored.

Organisations should be measuring time to detect material disruption, time to authorise containment, the proportion of critical paths with tested isolation and time to restore minimum acceptable service. AlHidaifi et al. (2024) give a sound argument for measuring resilience dynamically rather than just measuring the static security posture. Figure 2 shows an example of a service availability trajectory under disruptive conditions. The curves are entirely hypothetical (not actual measured performance) and are there to illustrate that the depth and duration of the service loss need to be considered in tandem. Table 1 takes the principal failure modes and turns them into evidence that an auditor or operator can actually go out and inspect.

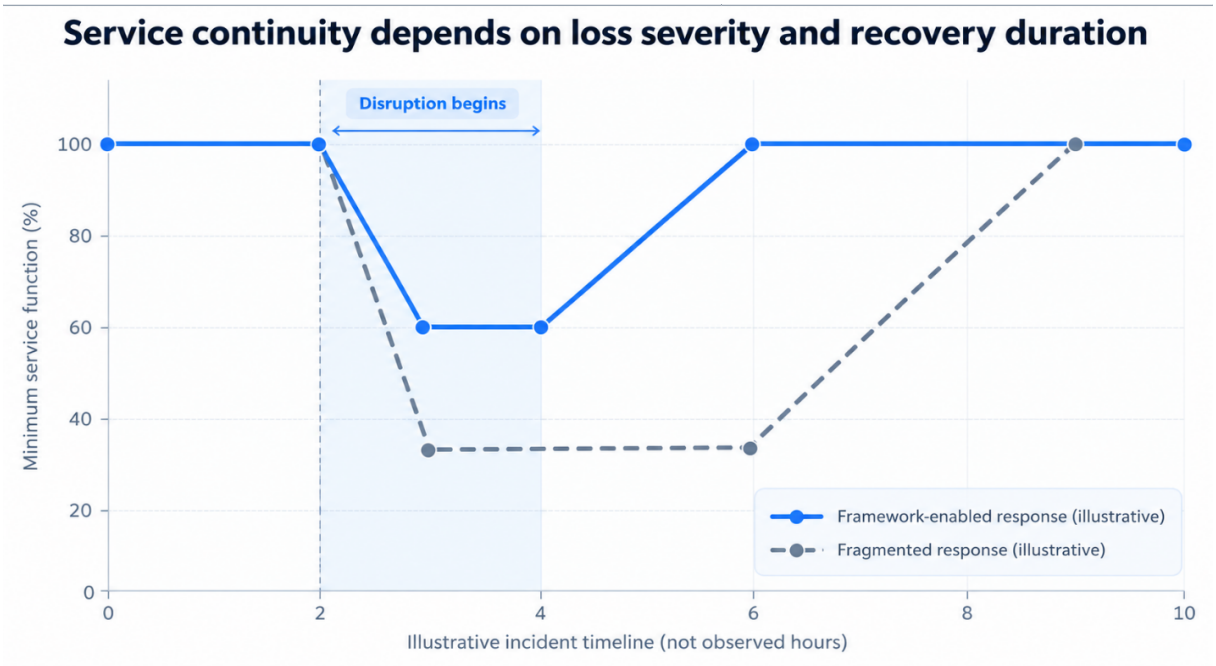


Figure 2. Illustrative service-function trajectories following disruption. Values and timelines are hypothetical visual aids, not observed results or estimated effect sizes.

Table 1. Representative multi-vendor failure modes, control responses and verifiable evidence.

Failure mechanism	Exposure	Control objective	Audit or operational evidence
Policy semantic drift	Vendor rules differ in precedence and object interpretation	Translate invariants; test intended and prohibited paths	Signed policy baseline; path-test trace; exception register
Privileged supplier access	Shared support accounts span network boundaries	Bound access by identity, approval, duration and resource	Token records; privileged-session log; revocation drill
Orchestration compromise	Broad controller privileges distribute malicious changes	Separate approval and execution; constrained roles; offline recovery	Change provenance; role audit; restoration exercise
Telemetry fragmentation	Incompatible logs delay incident reconstruction	Normalise identifiers, timestamps and provenance	Collection coverage; parsing errors; cross-source incident timeline
Correlated service failure	Distinct suppliers share identity or transit prerequisites	Map service dependencies; rehearse minimum-service restoration	Dependency map; failover evidence; achieved recovery threshold

6 Application to Saudi enterprise and critical-service settings

A Saudi deployment has to take the nation's cybersecurity obligations and turn them into engineering decisions without in any way equating compliance with resilience. The National Cybersecurity Authority publishes the Essential Cybersecurity Controls and specific requirements for critical systems and data. The Authority will have to work out which of those controls and requirements apply to its own activities, its suppliers and its information classes and then establish traceable mappings from each applicable requirement to policy, configuration, test results and accountable ownership. Pascoe et al. (2024) describe governance as a separate element of the cybersecurity risk

management process, one that enables an evidence based approach rather than a single compliance score.

Cross vendor estates are the norm when you're integrating old sites with cloud based apps, contracted networks and new industrial platforms. A sensible place to start would be a relatively small number of services of national or commercial significance where the degradation of those services would have a very clear operational impact. For a financial institution those might be authenticated transactions and settlement connectivity. For a health organisation it would be clinical communications and access to the patient systems. For an industrial operator it would be safe process visibility and authorised remote engineering. These are just illustrative service types, not in any way claims about a particular Saudi organisation.

Cloud dependencies are something that need to be looked at in detail as services that are marketed as standalone may in fact share an underlying identity, connectivity or administration layer. Sun (2020) points to a number of ongoing security and privacy issues in cloud computing with regard to control and responsibility boundaries. Saudi organisations should document who controls each recovery requirement and test the recovery when their primary cloud admin interface is not available. The assessment should cover locally hosted systems, regional service dependencies and external support arrangements and should not be based on any unsupported assumptions as to data residency or required architecture.

Data localisation, cross border support and managed service contracts all require a case by case legal analysis. A remote network vendor may well require diagnostic logs but the export of those logs is not necessarily permissible or required. The contract should set out the approved methods of access, the minimum level of event evidence required, the thresholds for notification and the responsibilities of the supplier in the event of a crisis (i.e. the responsibility for disabling the supplier's access). Boyens et al. (2022) treat supply chain governance as a lifecycle issue and Collier and Sarkis (2021) point to the risks of placing too much reliance on trust in extended supplier relationships. Contractual assurance should be subject to testing via exercises rather than simply being based on the promise of the document.

The most useful adaptation is institutional rather than superficial. Network engineering, security operations, business continuity, procurement and legal should all be involved in agreeing a single inventory of critical service dependencies and a single escalation path for conflicting operational decisions. Any supplier specific control which is unable to achieve the desired outcome should be assessed in terms of business impact, duration and the viability of a replacement path. That's the way to bring the national emphasis on secure digital development to bear on the reality of continuity, without in any way implying that the framework is an official certification or regulatory body endorsement.

7 Evaluation design and implementation roadmap

The framework needs to be subjected to empirical validation prior to any assertion of effectiveness. A credible study would involve a matched set of network scenarios in a reproducible emulation environment and then a monitored operational pilot. The scenarios would need to include policy drift between firewall platforms, compromise of a third party's management credential, outage of an automation controller, loss of a shared identity service and ransomware induced isolation of a network segment. The test configurations would need to record product versions, topology, back-

ground load, attack assumptions and recovery requirements. The emulation results would not be able to be extrapolated to national infrastructure without consideration of the operational constraints.

Five indicators are proposed. Policy conformance is the proportion of flows that are in line with the declared intent; detection latency is the time between a defined threat event and a validated alert; containment coverage is the proportion of critical paths that can be isolated without loss of essential functions; minimum service restoration time is the time to reach an agreed operational threshold; and evidence completeness is the extent to which key decisions and technical changes can be reconstructed. Shukla et al. (2022) provide the rationale for making explicit assurance claims and AlHidaifi et al. (2024) provide the rationale for service centric resilience indicators. The metrics should be presented with confidence intervals and at the scenario level rather than as a single composite figure which would mask serious problems.

Implementation should proceed in stages. Stage 1 is to map the services and the privileged relationships. Stage 2 is to standardise the policy intent and testing. Stage 3 is to integrate telemetry and constrained response. Stage 4 is to test disconnected recovery and supplier involvement. Each stage must have an operational exit criterion and an owner. The early pilots should focus on the critical paths with the lowest dependencies rather than trying to roll out across the whole device fleet at the same time. All changes must be rehearsed, versioned and where possible reversible.

To avoid misleading comparisons the tests need to have the same application traffic and functional requirements for the baseline and the framework. Repeated runs need to report the uncertainty that comes from network load, analyst availability and recovery sequence. The simulated attacker privileges must be made absolutely clear as the scenario of compromising a device console is very different to the scenario of having control of a central orchestrator. Negative results (failed policy translation, false isolation, excessive recovery time etc) need to be in the dataset. Only with this level of transparency can you be confident in determining whether the integration controls are actually providing an improvement in resilience or simply moving the operational risk elsewhere.

A very important falsifiable proposition is that automated, vendor-neutral policy assurance will reduce configuration inconsistency without increasing the time taken to restore service following a controller compromise. A second proposition is that having tested recovery paths which are independent of the production management plane will reduce the minimum time to restore service in the event of correlated failures. Neither of these propositions is established by this review. Both would require comparative measurement, transparent negative results and verification that any observed improvement is not achieved at the expense of increased operational fragility.

Table 2. Phased implementation with accountable owners and measurable exit criteria.

Phase	Accountable role	Required evidence	Minimum exit criterion
1 Establish dependencies	Service and asset owners	Critical-path register; approved supplier privileges	Every priority service has a signed dependency map
2 Assure policy equivalence	Network engineering and security architecture	Common policy invariants; vendor-native test results	Prohibited flows fail and necessary flows pass on tested devices
3 Integrate incident response	Security operations and service owners	Correlated telemetry; containment playbooks; audit trail	Scenario detection and authorised containment demonstrated

Phase	Accountable role	Required evidence	Minimum exit criterion
4 Prove independent recovery	Continuity leader, operations and suppliers	Clean baselines; alternate access; service restoration record	Minimum service restored under loss of primary control plane

8 Critical discussion and research limitations

The framework resolves one design tension but doesn't get rid of it. Centralised policy definition can lead to greater consistency whereas centralised control execution can lead to greater compromise. The architecture therefore needs to make the distinction between a common decision language and a single, essential runtime component. Independent approval, distributed enforcement and offline recovery are all useful in their own right. Cao et al. (2024) point to orchestration issues in zero trust systems and Melis et al. (2023) show why programmable networks are in need of scrutiny as both a form of protection and a potential attack surface.

Vendor neutral terminology is only useful if the semantics are precise. Two products can both support microsegmentation and yet have very different views on connection state, rule ordering or behaviour in the event of a management plane loss. The testing has to expose those differences. There's also the issue of evidence quality. A device can report that a rule has been installed without any actual proof that the intended traffic path is in fact blocked. That's why independent path testing, configuration snapshots and correlated service observations should all be used to provide assurance. That kind of testing can be very expensive in a regulated, safety critical or geographically dispersed environment.

Operational resilience also involves a range of trade-offs that go well beyond the engineering side. Aggressive containment can have a negative impact on short term service availability whereas a more conservative approach to containment may result in a larger eventual disruption. The appropriate decision will be driven by risk appetite and service criticality. A single score is not going to capture the nuances of these different scenarios. Future work should look at policy equivalence across firmware versions, recovery in the event of simultaneous credential and network failure, supplier led incident exercises and the cost of continuous verification on latency sensitive applications.

This is a very limited review. It's drawn from a selection of the literature and engineering guidance identified via DOI search rather than from a reproducible database sample. It contains no new measurements, no interviews, no field trials and the diagrams are all conceptual. So any conclusions regarding relative control performance are all propositions rather than actual findings. Saudi regulatory requirements are likely to change and the sectoral responsibilities will need to be verified by the appropriate professional at the time. The next step in the research would be a multi-case study reporting on configuration violations, detection performance and restoration outcomes in genuinely heterogeneous networks.

9 Conclusion

Cyber resilience in a multi-vendor environment is not about buying more security products or having more backup links. It's about the ability to express a common security intent, verify that intent has been implemented across the disparate technologies, be able to detect a material disruption and then be able to restore the essential services via a trusted alternative. This review brings those capabilities

together in a 5-capability framework covering dependency mapping, policy translation, identity based control, correlated telemetry and independently testable recovery.

The Saudi context makes ownership, supplier accountability and regulatory traceability of particular significance, though the architecture is of course applicable in any situation where network services cross organisational and technical boundaries. The proposed measures make a clear distinction between compliant configuration and demonstrated service continuity. Further experimental and operational validation is required before any assertions can be made as to effectiveness or cost benefits. The value of the work in the immediate term is that it provides a coherent, testable basis for making interoperability an enabler of resilience rather than an unmanaged source of correlated risk.

For the practitioner the first operational test is fairly trivial but very revealing. Pick one of the essential services, map out all the suppliers and credentials that it relies upon, deliberately disable a shared control and see whether an authorised team can recover the service with full evidence. That exercise will reveal the hidden fragility before the disruption actually happens.

References

[1] Alevizos L, Ta VT, Eiza MH (2022). Augmenting zero trust architecture to endpoints using blockchain: A state-of-the-art review. *Security and Privacy*, 5(1), e191. https://doi.org/10.1002/spy2.191

[2] AlHidaifi SM, Asghar MR, Ansari IS (2024). Towards a cyber resilience quantification framework (CRQF) for IT infrastructure. *Computer Networks*, 247, 110446. https://doi.org/10.1016/j.comnet.2024.110446

[3] Boyens J, Smith A, Bartol N, Winkler K, Holbrook A, Fallon M (2022). Cybersecurity Supply Chain Risk Management for Systems and Organizations. *NIST Special Publication* 800–161, Revision 1. https://doi.org/10.6028/NIST.SP.800-161r1

[4] Buck C, Olenberger C, Schweizer A, Völter F, Eymann T (2021). Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero trust. *Computers & Security*, 110, 102436. https://doi.org/10.1016/j.cose.2021.102436

[5] Cao Y, Pokhrel SR, Zhu Y, Doss R, Li G (2024). Automation and orchestration of zero trust architecture: Potential solutions and challenges. *Machine Intelligence Research*, 21, 294–317. https://doi.org/10.1007/s11633-023-1456-2

[6] Cascavilla G, Tamburri DA, van den Heuvel WJ (2021). Cybercrime threat intelligence: A systematic multi-vocal literature review. *Computers & Security*, 105, 102258. https://doi.org/10.1016/j.cose.2021.102258

[7] Chandramouli R, Butcher Z (2023). A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Location Environments. *NIST Special Publication* 800–207A. https://doi.org/10.6028/NIST.SP.800-207A

[8] Collier ZA, Sarkis J (2021). The zero trust supply chain: Managing supply chain risk in the absence of trust. *International Journal of Production Research*, 59(11), 3430–3445. https://doi.org/10.1080/00207543.2021.1884311

[9] He Y, Huang D, Chen L, Ni Y, Ma X (2022). A survey on zero trust architecture: Challenges and future trends. *Wireless Communications and Mobile Computing*, 2022, 6476274. https://doi.org/10.1155/2022/6476274

[10] Katsis C, Cicala F, Thomsen D, Ringo N, Bertino E (2021). Can I reach you? *Do I need to? New semantics in security policy specification and testing*. Proceedings of the ACM Symposium on Access Control Models and Technologies, 165–174. https://doi.org/10.1145/3450569.3463558

[11] Landauer M, Skopik F, Wurzenberger M, Rauber A (2020). System log clustering approaches for cyber security applications: A survey. *Computers & Security*, 92, 101739. https://doi.org/10.1016/j.cose.2020.101739

[12] Leivadeas A, Falkner M (2023). A survey on intent-based networking. *IEEE Communications Surveys & Tutorials*, 25(1), 625–655. https://doi.org/10.1109/COMST.2022.3215919

[13] Mehmood K, Kralevska K, Palma D (2023). Intent-driven autonomous network and service management in future cellular networks: A structured literature review. *Computer Networks*, 220, 109477. https://doi.org/10.1016/j.comnet.2022.109477

[14] Melis A, Sadi AA, Berardi D, Callegati F, Prandini M (2023). A systematic literature review of offensive and defensive security solutions with software defined network. *IEEE Access*, 11, 93431–93463. https://doi.org/10.1109/ACCESS.2023.3276238

[15] Mubarkoot M, Altmann J, Rasti-Barzoki M, Egger B, Lee H (2023). Software compliance requirements, factors, and policies: A systematic literature review. *Computers & Security*, 124, 102985. https://doi.org/10.1016/j.cose.2022.102985

[16] Nelson A, Rekhi S, Scarfone K, Souppaya M (2025). Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile. *NIST Special Publication* 800–61, Revision 3. https://doi.org/10.6028/NIST.SP.800-61r3

[17] Pascoe C, Quinn S, Scarfone K (2024). The NIST Cybersecurity Framework (CSF) 2.0. *NIST Cybersecurity White Paper* 29. https://doi.org/10.6028/NIST.CSWP.29

[18] Reichert BM, Obelheiro RR (2024). Software supply chain security: A systematic literature review. *International Journal of Computers and Applications*, 46(10), 853–867. https://doi.org/10.1080/1206212X.2024.2390978

[19] Ribeiro RH, Jacobs AS, Zembruzki L, Parizotto R, Scheid EJ, Schaeffer-Filho AE, Granville LZ, Stiller B (2022). A deterministic approach for extracting network security intents. *Computer Networks*, 214, 109109. https://doi.org/10.1016/j.comnet.2022.109109

[20] Rose S, Borchert O, Mitchell S, Connelly S (2020). Zero Trust Architecture. *NIST Special Publication* 800–207. https://doi.org/10.6028/NIST.SP.800-207

[21] Ross R, Pillitteri V, Graubart R, Bodeau D, 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

[22] Ross R, Pillitteri V (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

[23] Scarfone K, Souppaya M, Dodson D (2022). Secure Software Development Framework (SSDF) Version 1.1: Recommendations for Mitigating the Risk of Software Vulnerabilities. *NIST Special Publication* 800–218. https://doi.org/10.6028/NIST.SP.800-218

[24] Shukla A, Katt B, Nweke LO, Yeng PK, Weldehawaryat GK (2022). System security assurance: A systematic literature review. *Computer Science Review*, 45, 100496. https://doi.org/10.1016/j.cosrev.2022.100496

[25] Stamenkov G (2022). Layered business continuity and disaster recovery model. *Continuity & Resilience Review*, 4(3), 267–279. https://doi.org/10.1108/CRR-05-2022-0008

[26] Sun PJ (2020). Security and privacy protection in cloud computing: Discussions and challenges. *Journal of Network and Computer Applications*, 160, 102642. https://doi.org/10.1016/j.jnca.2020.102642

[27] Syed NF, Shah SW, Shaghaghi A, Anwar A, Baig Z, Doss R (2022). Zero trust architecture (ZTA): A comprehensive survey. *IEEE Access*, 10, 57143–57179. https://doi.org/10.1109/ACCESS.2022.3174679

[28] Teerakanok S, Uehara T, Inomata A (2021). Migrating to zero trust architecture: Reviews and challenges. *Security and Communication Networks*, 2021, 9947347. https://doi.org/10.1155/2021/9947347

[29] Umer MA, Junejo KN, Jilani MT, Mathur AP (2022). Machine learning for intrusion detection in industrial control systems: Applications, challenges, and recommendations. *International Journal of Critical Infrastructure Protection*, 38, 100516. https://doi.org/10.1016/j.ijcip.2022.100516

[30] Wang Z, Fok KW, Thing VLL (2023). Machine learning-powered encrypted network traffic analysis: A comprehensive survey. *IEEE Communications Surveys & Tutorials*, 25(1), 791–824. https://doi.org/10.1109/COMST.2022.3208196