© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722267

Cyber Resilience of IP/MPLS Core Networks in Saudi Arabia: Hardening, Threat Analysis, and Operational Continuity under Vision 2030

ZAKIUDDIN MOHAMMED

Abstract- ***Saudi Arabia's digital economy increasingly depends on IP/MPLS core networks that aggregate mobile, fixed, cloud, government and industrial traffic. Their importance creates a difficult resilience problem: the core must preserve deterministic forwarding and rapid convergence while resisting control-plane manipulation, data-plane abuse, management compromise, supply-chain weaknesses and operational error. This review synthesises research, standards and Saudi regulatory guidance published from 2020 to 2025 to develop an integrated cyber-resilience framework for national-scale carrier networks. The review examines layered attack surfaces across external routing, MPLS and segment-routing control, provider-edge services, router operating systems, automation platforms, telemetry pipelines and operational processes. It argues that resilience cannot be achieved by perimeter filtering alone. A defensible design combines authenticated administration, infrastructure access controls, route-origin validation, explicit peering roles, control-plane policing, label and interface validation, segmented management, cryptographic agility, continuous telemetry, tested rollback, diverse transport paths and recovery procedures governed by service criticality. The resulting framework links prevention, detection, containment, traffic restoration and organisational learning. It also aligns technical controls with Saudi cybersecurity obligations and Vision 2030 priorities for trusted digital infrastructure. The review concludes that the strongest operating model treats routing correctness, security assurance and continuity engineering as one measurable discipline rather than separate network, security and business-continuity functions [1–4].***

Keywords: IP/MPLS, Cyber Resilience, Core Networks, BGP Security, Segment Routing, RPKI, Network Telemetry, Operational Continuity, Saudi Vision 2030

I. INTRODUCTION

IP/MPLS cores are the circulation system of contemporary telecommunications. They carry internet transit, mobile backhaul, enterprise virtual private networks, cloud interconnects, emergency services and digital-government traffic through a relatively small number of high-capacity nodes and links. This concentration produces efficiency, but it also concentrates consequence. A malformed route, compromised router, control-plane flood, label error or automation defect can propagate across many services before customer-facing teams understand the common cause. In Saudi Arabia, this risk is amplified by rapid cloud adoption, nationwide digital services, smart-city programmes and dependence on resilient connectivity for critical sectors. Vision 2030 therefore makes core-network resilience an economic and public-service requirement rather than a narrow engineering preference [3,4].

Traditional carrier security often assumes that an MPLS domain is trustworthy because it is privately operated. That assumption is increasingly fragile. Provider networks now connect to cloud exchanges, managed-service platforms, internet peers, outsourced operations, virtualised network functions and application programming interfaces. Core routers also expose streaming telemetry, model-driven configuration and remote orchestration interfaces. Each connection improves agility while creating a path through which weak identity controls, stolen credentials, malicious updates or inaccurate data can influence forwarding. Zero-trust principles are relevant because they replace location-based confidence with explicit verification of users, devices, workloads and requested actions [2].

The threat is not limited to deliberate intrusion. Operational continuity is frequently challenged by incorrect policy, software defects, incomplete change testing, optical failure, power disruption and poorly coordinated maintenance. Security and reliability consequently overlap. A route leak may be accidental

yet operationally resemble an attack; a denial-of-service event may exploit legitimate protocol behaviour; and an emergency traffic-engineering action may create a new exposure if it bypasses established policy. Effective cyber resilience must therefore preserve confidentiality where needed, but it must give equal weight to route integrity, service availability, bounded failure domains and recoverability.

This paper reviews evidence from 2020–2025 and asks three questions. First, which attack and failure patterns most seriously threaten Saudi IP/MPLS cores? Second, which architectural and operational controls provide defence in depth without undermining convergence and service performance? Third, how should operators measure continuity, exercise recovery and demonstrate alignment with national controls? The aim is to produce a practical synthesis for carrier architects, security leaders, network operations centres and regulators. The paper contributes a layered threat model, a hardening blueprint and a continuity operating model that connect technical protection with service outcomes.

II. REVIEW METHODOLOGY

An integrative review method was selected because the subject spans protocol standards, cybersecurity frameworks, operational research and national regulatory instruments. Sources were limited to material published from January 2020 to December 2025. The evidence base included peer-reviewed studies on backbone anomaly detection and time-series analysis; standards defining network telemetry, segment-routing policy, routing-security mechanisms and MPLS operations; Saudi controls for communications providers, cloud, data, operational technology and critical systems; and digital-government guidance on continuity. The attached layered-security survey informed the organisation of threats by architectural plane rather than by isolated product category [1].

Search concepts combined IP/MPLS, carrier core, BGP security, route leak, RPKI, segment routing, control-plane protection, telemetry, anomaly detection, incident recovery and Saudi critical infrastructure. Sources were screened for direct relevance to one of five analytical domains: management and orchestration, routing control, MPLS or segment-routing forwarding, service edge and interconnection, and operations or recovery. Material was excluded when it pre-dated 2020, repeated older guidance without substantive update, focused only on endpoint security, or lacked enough technical detail to support a control recommendation.
The synthesis used a threat-control-outcome matrix. For each domain, the review identified assets, credible threat events, observable indicators, preventive safeguards, containment actions and continuity outcomes. Recommendations were then tested against three criteria: operational feasibility in a high-availability carrier environment, compatibility with standards-based routing, and traceability to Saudi governance expectations. This approach is not a quantitative meta-analysis because the evidence contains heterogeneous designs, standards and case-based evaluations. Instead, it develops a structured, auditable argument and identifies where empirical validation in Saudi production networks remains necessary.

III. IP/MPLS CORE ARCHITECTURE AND RESILIENCE CONTEXT

A national IP/MPLS core can be understood as interacting planes rather than a single trusted fabric. The data plane forwards labelled and unlabelled packets at line rate. The control plane builds reachability through interior gateway protocols, BGP, label distribution and traffic-engineering mechanisms. The management plane configures devices, software images, credentials and policy. A service plane creates Layer 2 and Layer 3 VPNs, internet access, mobile transport, cloud connectivity and shared functions. Above these planes, automation and assurance systems translate commercial intent into configuration and evaluate telemetry. Failure in one plane can migrate into another; a compromised orchestrator can alter control policy, while incorrect control state can blackhole the data plane.

Modernisation changes this relationship. Segment routing reduces intermediate per-path state and allows a headend to steer traffic through an ordered

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722267

policy [10]. SR-MPLS preserves the label-based data plane, while SRv6 binds instructions to IPv6 segment identifiers and supports programmable behaviours [6]. These capabilities improve traffic engineering, service chaining and recovery, but they increase the importance of protecting policy computation, segment lists and controller interfaces. An attacker who gains authority over a headend or path-computation process may not need to compromise every transit router; the attack can be expressed as an apparently valid policy.

The external routing boundary is equally important. BGP advertisements determine how national services are reached from other networks and where outbound traffic exits. Route leaks violate expected provider, customer or peer relationships, while origin hijacks falsely claim address space. BGP roles and the Only-to-Customer attribute provide a standards-based mechanism for preventing and detecting common leak patterns [9]. Route Origin Authorisations and validation reduce forged-origin risk, although poor maxLength choices enlarge the attack surface and validator availability remains a resilience concern [12,13,16].

Observability must therefore cover state, intent and forwarding. The network telemetry framework distinguishes data from devices, forwarding, control and management planes and supports uses ranging from service assurance to security analytics [8]. For MPLS and segment-routing environments, OAM must verify not merely that a packet returns, but that it traverses the intended path and terminates at the correct egress. SRv6 OAM and IPFIX export of segment information improve path visibility [11,22], while newer LSP ping extensions strengthen egress validation for MPLS and SR policies [28]. Figure 1 expresses these dependencies as a layered resilience architecture.

Service-edge and platform dependencies complete the architecture. Cloud interconnects, DNS, identity services, DDoS mitigation and customer-facing VPN gateways may sit outside the physical core while remaining essential to service delivery. Saudi cloud and data controls require lifecycle protection, access accountability, monitoring and secure hosting practices [4,17]. Where the core carries industrial or critical-facility traffic, operational-technology controls add requirements for isolation, controlled remote access and continuity-sensitive change [18]. A resilient design therefore maps logical services to every external platform and shared function on which restoration depends.

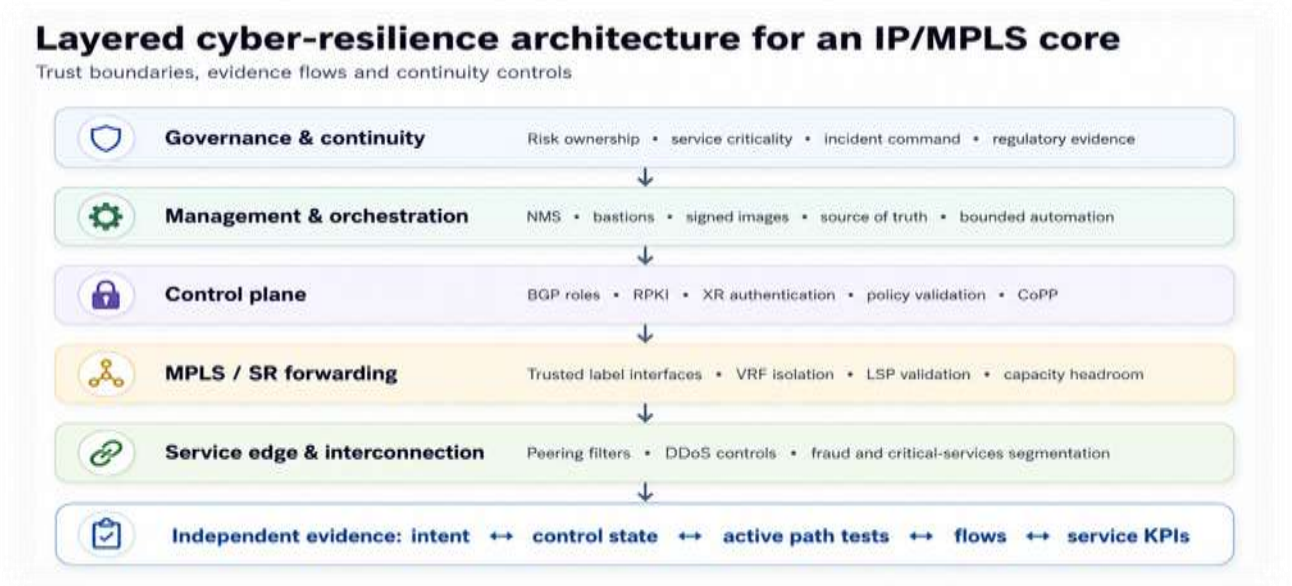

Figure 1. Layered cyber-resilience architecture for a national IP/MPLS core.

IV. THREAT ANALYSIS ACROSS THE CORE

Management and orchestration compromise is the highest-leverage threat because it can convert authorised tooling into a distribution mechanism for malicious change. Credential theft, exposed APIs, weak role separation, insecure jump hosts and poisoned software repositories may allow an adversary to alter routing policy, disable logging or implant persistence across many routers. The impact is magnified when a single automation account has unrestricted access to production, disaster-recovery and laboratory environments. Hardening must consequently begin with phishing-resistant authentication, privileged-access management, signed software, segregated administrative paths and transaction-level approval for high-risk changes.

Control-plane attacks target the mechanisms that create reachability. They include forged or replayed routing messages, session reset attempts, route leaks, prefix hijacks, excessive updates and manipulation of path attributes. TCP Authentication Option implementation testing supports stronger session protection than legacy mechanisms where platforms permit it [14]. Nevertheless, session authentication does not prove that a route is authorised or commercially valid. Operators need layered policy: strict import and export filters, maximum-prefix limits, BGP roles, route-origin validation, community controls and independent route monitoring. Data-plane validation is also necessary because a correct control-plane entry does not guarantee that packets follow the expected path [15].

Within the MPLS domain, abuse can involve unexpected labelled packets, label-stack manipulation, cross-VRF leakage, incorrect penultimate-hop behaviour, unauthorised service labels and deliberate exhaustion of forwarding resources. Provider-edge systems are particularly sensitive because they translate between customer, internet and core contexts. Infrastructure interfaces should not accept labelled traffic unless operationally required, and customer-facing interfaces should enforce explicit address, protocol and rate policies. VPN route targets, route distinguishers and import policies require automated linting because a small configuration error can violate tenant separation without compromising cryptography.

Denial of service can attack bandwidth, route processors, memory, adjacency tables or operational staff. High packet rates directed at router addresses, TTL-expiry processing, ICMP generation, BGP sessions, routing-protocol multicast or management services may shift traffic from hardware forwarding into slower control processing. Distributed attacks can also congest peering or international links before packets reach scrubbing systems. Control-plane policing must therefore be based on measured protocol needs, not copied templates. Limits that are too high offer little protection; limits that are too low transform legitimate reconvergence into self-inflicted failure.

Telemetry and assurance systems create another attack surface. Streaming data may reveal topology, interface utilisation, routing neighbours and customer relationships. False or delayed telemetry can also cause an automated controller to make harmful decisions. The problem resembles data poisoning described in layered IoT security research, but the operational consequence in a carrier core is incorrect network action rather than only inaccurate classification [1]. Telemetry collectors require mutual authentication, time integrity, schema validation, access control and independent comparison between configured intent, control state and sampled forwarding.

Supply-chain and lifecycle risks persist throughout the architecture. Router operating systems contain millions of lines of code, hardware components have long replacement cycles, and emergency upgrades may introduce new defects. Unsupported releases, inconsistent cryptographic capabilities and undocumented dependencies increase exposure. Resilience requires an accurate bill of materials, vulnerability intelligence, staged patching, golden images, tested downgrade paths and spare capacity. It also requires vendor diversity to be considered carefully: diversity may reduce common-mode failure, but too many platforms can increase configuration inconsistency and slow incident response.

Segment-routing extensions introduce protocol-specific verification needs. SRv6 OAM can test reachability of segment identifiers and sample traffic at segment endpoints, while IS-IS and OSPFv3 extensions advertise locators, behaviours and capabilities [11,20,21]. Incorrect or unauthorised advertisements can steer traffic towards unintended functions even when basic adjacency remains healthy. Operators should alarm on new locators, unexpected SID behaviour, candidate-path changes and maximum-stack-depth anomalies. RPKI policy also needs operational safeguards: restrictive authorisations reduce hijack exposure, but over-broad maxLength values and validator failure modes can weaken protection [12,13]. Exported SRv6 information can support forensic reconstruction when its integrity and retention are protected [22].

Human and process failures remain credible causes of national impact. A technically valid command entered in the wrong maintenance window can withdraw thousands of routes. A policy generated from an incorrect source-of-truth record may pass syntax validation while violating service intent. Fatigue during a prolonged incident can lead to uncontrolled changes and loss of forensic evidence. Table 1 therefore treats accidental, malicious and systemic events together. The unifying question is whether the network detects divergence early, contains the blast radius and restores priority services within agreed objectives.

Table 1. Threat-control-continuity matrix for IP/MPLS core networks.

Domain	Principal threats and failure modes	Primary hardening and detection controls	Containment and continuity response
Management and orchestration	Credential theft; API abuse; malicious or defective automation; software-image compromise	PAM; hardened bastions; signed software; scoped service identities; dual approval; immutable logs	Revoke identities; freeze changes; isolate controllers; restore golden image and validated intent
External routing	Prefix hijack; route leak; excessive updates; session reset; invalid export	RPKI; BGP roles; strict import/export policy; maximum prefixes; authentication; independent monitoring	Filter or withdraw routes; activate emergency communities; coordinate peers; validate global convergence
IGP and SR control	Unauthorised adjacency; false locator or SID; path-controller manipulation; control-plane flood	IGP authentication; passive interfaces; infrastructure ACLs; SID registry; policy simulation; CoPP	Isolate adjacency or controller; revert candidate path; protect stable control state; reroute priority traffic
MPLS and VPN forwarding	Label spoofing; labelled traffic on untrusted links; cross-VRF leakage; forwarding/control divergence	Trusted label interfaces; route-target linting; VRF invariants; LSP ping and egress validation	Block interface or label source; withdraw affected VPN policy; move service; verify forwarding before normalisation
Telemetry and assurance	Data poisoning; collector compromise; clock drift; loss of evidence; alert overload	Mutual authentication; time integrity; schema validation; multi-source correlation; retention and access controls	Fail over collectors; distrust automated actions; use active probes and customer KPIs; preserve evidence

Domain	Principal threats and failure modes	Primary hardening and detection controls	Containment and continuity response
Physical and capacity	Fibre cut; power loss; international congestion; common-mode platform failure	Route diversity verification; capacity headroom; redundant sites; spares; staged lifecycle management	Prioritise critical classes; use alternate gateways; shed non-essential load; restore with verified capacity

V. INTEGRATED HARDENING FRAMEWORK

The proposed framework begins with trustworthy governance of identity and change. Administrative access should use a dedicated out-of-band network, hardened bastions, device certificates and phishing-resistant multi-factor authentication. Personal accounts should replace shared credentials, and automation identities should be scoped by device group, command class and time. High-consequence actions—such as changing route-reflector policy, altering label allocation, disabling telemetry or modifying external export filters—should require dual approval and generate immutable audit records. Zero-trust guidance supports continuous evaluation rather than assuming that an authenticated session remains trustworthy throughout its lifetime [2].

At the routing boundary, operators should maintain an authoritative registry of prefixes, autonomous-system relationships, expected peers and approved export intent. BGP sessions need source filtering, generalised TTL security where appropriate, strong authentication, maximum-prefix thresholds and conservative timers. Import policy should reject bogons, martians, excessively specific routes and invalid origins according to documented risk rules. Export policy should be generated from the source of truth, independently compiled and tested against customer-provider relationships. BGP roles and route-leak signalling reduce reliance on unilateral configuration assumptions [9]. RPKI validation should use diverse, monitored validators with cached data and defined fail-safe behaviour, because validator or repository disruption can become a downgrade path [16].

Interior routing and label distribution require equivalent discipline. IS-IS or OSPF adjacencies should be confined to infrastructure links, authenticated with current algorithms and protected from unauthorised participation. Passive interfaces, adjacency limits and infrastructure access-control lists reduce exposure. Segment-routing locators, SIDs and policy candidates must be allocated from controlled registries. Standards for IS-IS and OSPFv3 SRv6 advertisement show how capabilities and locators enter the control plane [20,21]; therefore, incorrect advertisements must be treated as security-relevant events. Policy controllers should validate segment lists against topology, maintenance state, latency constraints and prohibited transit zones before installation.

The MPLS data plane should enforce explicit trust boundaries. Labelled packets should be accepted only from recognised core or inter-provider interfaces. Uniform and pipe TTL behaviour, ICMP tunnelling, entropy labels and explicit-null settings should be documented because they affect both troubleshooting and attack visibility. Provider-edge route import must be constrained to the intended VPN, and service-leaf or internet-edge roles should not be collapsed without compensating controls. LSP ping, traceroute and egress-validation procedures should be automated after change and during incident diagnosis to confirm synchronisation between control and forwarding state [28].

Router hardening should minimise reachable services. Unused protocols, web interfaces, discovery functions and legacy management services should be disabled. SSH, NETCONF, gNMI, SNMP and telemetry endpoints should bind only to management addresses and use modern cryptography. Configuration archives and software images require integrity checking. Secure boot and signed packages are valuable where supported, but operators must also plan for cryptographic agility because a national core cannot be upgraded instantly. Saudi cloud and communications-provider controls emphasise access

governance, secure configuration, vulnerability management and continuous compliance, which apply directly to controller and network-management environments [3,4].

Segmentation must extend beyond customer VPNs. The management plane, routing infrastructure, telemetry collectors, orchestration platforms, security tools and corporate IT should occupy separately controlled zones. A compromise in office productivity systems should not provide a route to router consoles. Similarly, a telemetry platform should have read-oriented privileges and should not silently inherit configuration authority. Where operational technology or critical-sector services traverse the core, their supporting management paths should follow stronger isolation and recovery requirements consistent with national OT and data controls [17,18].

Detection should combine deterministic rules with behavioural analytics. Deterministic controls identify invalid origins, unexpected peers, route-target changes, unauthorised commands and configuration drift. Behavioural methods analyse interface traffic, update rates, CPU, memory, latency, loss, label-stack depth and service KPIs. Graph-based analysis is useful because traffic anomalies at core and metro networks often reflect topology as well as time [5]. Deep time-series models can capture long-range dependencies but must be governed for false positives, data drift and explainability [7]. The most reliable design correlates routing, flow, device, optical and customer-experience signals instead of trusting one detector.

Automation must be bounded. Closed-loop remediation may safely quarantine a management account, suppress a known-bad prefix or shift traffic to a pre-authorised path. It should not invent wide-area routing policy during an uncertain incident. Every automated action needs eligibility conditions, maximum scope, rollback criteria and a human override. Digital signatures, versioned intent, pre-deployment simulation and canary rollout reduce the chance that automation amplifies error. Network changes should be evaluated not only for syntax and reachability but also for security invariants: no unauthorised route export, no cross-VRF import, no prohibited management exposure and no loss of telemetry.

Hardening also requires a governance model that makes risk ownership explicit. The NIST Cybersecurity Framework 2.0 places governance alongside identification, protection, detection, response and recovery [24]. Saudi essential controls similarly require approved policies, responsibilities, risk treatment, supplier oversight and continuous compliance [25]. Digital-government policies connect cybersecurity with continuity, data management and service accountability [26]. For core networks, these principles should be translated into named owners for peering policy, route-origin data, automation code, cryptographic keys, software exceptions, capacity risk and recovery evidence. Technical debt should be recorded with expiry dates and compensating controls rather than remaining an informal operational acceptance.

Finally, capacity is a security control. A network operating near saturation has little ability to absorb attacks, reroute around failures or perform maintenance. Operators should reserve headroom on core links, route reflectors, control processors, scrubbing paths and interconnection capacity. Diverse fibre routes must be physically verified rather than assumed from logical diagrams. Critical services should have pre-defined traffic classes and restoration priorities. The hardening framework is effective only when controls are tested under realistic load, failure and adversarial conditions rather than confirmed through configuration review alone.

VI. OPERATIONAL CONTINUITY AND INCIDENT RECOVERY

Operational continuity converts technical safeguards into measurable service survival. The first requirement is a service dependency map linking customer products and national functions to VPNs, route reflectors, provider-edge nodes, optical systems, DNS, authentication, cloud interconnects and management platforms. Without this map, recovery teams may restore visible infrastructure while leaving a hidden dependency unavailable.

© AUG 2026 | IRE Journals | Volume 10 Issue 2 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I2-1722267

Business impact analysis should define recovery-time and recovery-point objectives for management data, configuration repositories, telemetry, billing evidence and customer services. Saudi digital-government continuity guidance emphasises governance, risk assessment, exercises and continual improvement [19].

The response model should distinguish detection, verification, containment, traffic restoration, secure recovery and learning. Detection identifies abnormal state; verification tests whether the event is malicious, accidental or physical. Containment limits propagation through route filtering, session isolation, control-plane policies or management credential revocation. Traffic restoration uses pre-engineered paths, alternate peers, diverse sites and service prioritisation. Secure recovery restores trustworthy software, configuration and telemetry before normal optimisation resumes. Figure 2 shows this loop and the evidence required at each stage.

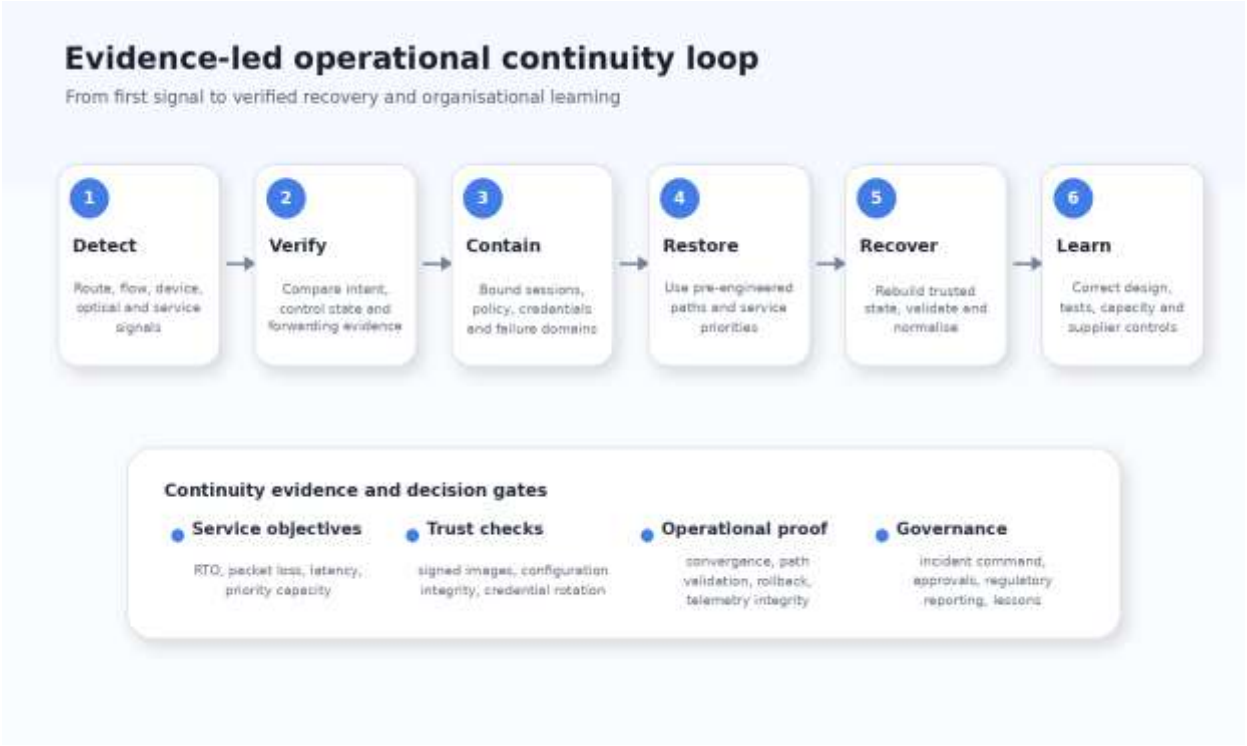

Figure 2. Evidence-led incident containment, restoration and learning loop.

Runbooks should be scenario-specific. A route leak runbook may freeze external policy changes, compare advertisements with the source of truth, apply emergency communities, contact peers and validate global withdrawal. A suspected router compromise may isolate management access, preserve volatile data, move services, reload a trusted image and rotate credentials. A control-plane flood may activate scrubbing, tighten policers, protect routing sessions and reduce non-essential diagnostics. A fibre-cut scenario focuses on capacity, latency and restoration priority rather than forensic containment.

Shared command structures prevent teams from applying contradictory remedies.

Testing must include more than device redundancy. Operators should exercise loss of a route reflector, corruption of the source of truth, unavailability of RPKI validators, failure of the out-of-band network, compromise of an automation account and simultaneous congestion on alternate paths. Tabletop exercises validate decisions and communications; laboratory tests validate configuration; controlled production drills validate real convergence and capacity. Evidence should include detection time,

decision time, route convergence, packet loss, customer-impact duration, rollback success and whether telemetry remained trustworthy.

Continuity also depends on people. Network operations, security operations, optical transport, cloud, customer care, legal, communications and executive leadership need a common severity model. During a national-scale event, one incident commander should own priorities, while technical leads retain authority within defined domains. Shift handover must capture hypotheses, commands, rejected options and residual risk. Training should include protocol knowledge and decision-making under uncertainty, because an incident may require restoring service before attribution is complete.

Post-incident review must connect root cause with control design. It is insufficient to record that an engineer entered a wrong command; the review should ask why the command was possible, why simulation failed, why blast-radius controls did not stop it and why detection took as long as it did. Corrective actions should update architecture, automation tests, monitoring thresholds, capacity plans and supplier obligations. Recovery metrics in Table 2 help convert lessons into accountable improvement and support evidence-based reporting to regulators and critical customers.

Table 2. Operational resilience metrics and decision thresholds.

Metric	Purpose	Example evidence	Decision use	Review cadence
Mean time to trustworthy detection	Measure how quickly material divergence is recognised and corroborated	Routing alert plus independent flow or active-path evidence	Escalate severity and select containment authority	Monthly and after every major incident
Blast-radius index	Quantify affected nodes, VPNs, regions and critical services	Topology and service dependency graph	Choose local isolation, regional reroute or national command	Per exercise and incident
Priority-service restoration time	Confirm that essential traffic is restored within agreed objectives	Convergence, loss, latency and application checks	Release emergency capacity and communicate recovery status	Quarterly drills
Rollback success rate	Test whether changes can be reversed without secondary impact	Versioned intent, configuration diff and post-checks	Approve automation scope and change-risk class	Every release cycle
Telemetry trust score	Assess completeness, integrity, time alignment and source diversity	Collector health, signature validation and cross-source agreement	Permit or suspend automated remediation	Continuous
Control effectiveness closure	Track completion and validation of corrective actions	Retest results, updated runbooks, audit evidence	Close incident actions and regulatory commitments	Quarterly governance review

VII. SAUDI GOVERNANCE AND VISION 2030 ALIGNMENT

Saudi operators work within a national framework that treats telecommunications and critical digital services as protected infrastructure. The communications-sector cybersecurity framework requires licensed providers to raise maturity and protect confidentiality, user interests and service quality [3]. The NCA's essential, cloud, data and

operational-technology controls add governance, access, vulnerability, monitoring, continuity and supplier expectations [4,17,18,25]. For an IP/MPLS core, compliance should be mapped to concrete evidence: approved routing policy, privileged-access records, configuration baselines, vulnerability decisions, recovery tests, telemetry retention and incident-response performance.

Alignment should avoid a checklist mentality. A control stating that access is restricted does not demonstrate that an automation token cannot change every router. A continuity policy does not prove that alternate paths have capacity. Operators need control objectives expressed as network invariants and service outcomes. Examples include: external routes are exported only to authorised relationships; labelled traffic enters only through trusted interfaces; management is unreachable from customer or corporate zones; critical services retain defined capacity after any single failure; and emergency changes remain attributable and reversible.

Vision 2030's expansion of digital infrastructure, cloud services and technology-enabled public services increases both the value and the exposure of national connectivity [29]. Resilience supports this agenda by reducing systemic outages, protecting confidence in digital transactions and enabling secure innovation. Investment should therefore prioritise observability, automation assurance, diverse connectivity, local operational capability and repeated exercises. These capabilities create a trusted platform on which cloud, smart-city, industrial and government services can scale without transferring unmanaged risk into the national core.

VIII. DISCUSSION, LIMITATIONS AND RESEARCH AGENDA

The review indicates that the central design problem is not choosing between security and performance, but engineering controls that preserve both. Carrier cores cannot rely on deep inspection of all transit traffic, and they cannot introduce manual approval into every convergence event. The appropriate controls operate at trust boundaries and decision points: peer relationships, administrative identity, route intent, label acceptance, policy installation and recovery authority. These controls are lightweight in the forwarding path but strict in the systems that define forwarding.

A second finding is that resilience requires independent evidence. Configuration state alone is inadequate because forwarding may diverge from intent. Telemetry alone is inadequate because collectors can fail or be deceived. Customer alarms alone are late and ambiguous. Strong assurance triangulates intended policy, distributed control state, active path tests, sampled flows and service outcomes. Network telemetry standards and recent ISP datasets create opportunities for more realistic anomaly research [8,27]. Saudi operators could contribute anonymised, labelled incident traces and testbeds that reflect desert geography, international gateway concentration, pilgrimage traffic, cloud regions and giga-project connectivity.

Anomaly classification deserves particular attention because a detector that merely labels a time series as abnormal gives responders little guidance. Research on telecommunications KPIs shows the value of separating detection from classification so that spikes, level shifts, trends and contextual anomalies can be routed to different workflows [23]. In a core network, classification should add topology, maintenance state and service criticality. The same traffic increase may indicate attack, legitimate event demand, rerouting after a cut or measurement error; response automation must not treat these causes as equivalent.

Several limitations remain. Public research provides limited evidence on real Saudi carrier topologies, attack frequency, convergence behaviour and recovery performance. Many standards specify mechanisms without quantifying operational trade-offs at national scale. Machine-learning studies often use enterprise or synthetic datasets that do not represent encrypted, high-capacity backbone traffic. The review therefore proposes controls based on convergence of standards and operational logic, but each operator must validate thresholds, capacity and automation in its own environment.

Future research should examine four areas. First, explainable cross-domain analytics should correlate routing updates, optical alarms, flow telemetry, device logs and customer KPIs while preserving data sovereignty. Second, digital twins should test policy and failure propagation before production deployment. Third, cryptographic agility for long-lived routers should be evaluated against performance and maintenance constraints. Fourth, autonomous remediation should be bounded by formal invariants and verified rollback. NIST's implementation guidance for zero-trust architectures demonstrates the value of testing integrated designs rather than treating products as isolated controls [30]. Similar multi-vendor demonstrations are needed for carrier routing and MPLS resilience.

The practical implication is organisational. Network engineering, cybersecurity and business continuity teams should share one resilience backlog, one service-criticality model and one exercise programme. Separate dashboards encourage local optimisation: security may report blocked attacks while customers experience outages; network teams may report availability while recovery evidence is weak. A unified model evaluates whether the organisation can prevent avoidable incidents, detect material divergence, contain propagation, restore priority services and learn faster than the threat and technology landscape changes. This integrated view gives boards, regulators and operators a basis for prioritising resilience investments and residual risks.

IX. CONCLUSION

Saudi Arabia's IP/MPLS cores are foundational to Vision 2030, yet their scale and programmability make small errors and targeted attacks capable of producing broad service impact. This review shows that cyber resilience depends on coordinated protection of management, control, forwarding, service and operational planes. The most defensible architecture combines strong administrative identity, routing-security policy, controlled label boundaries, segmented management, continuous multi-source telemetry, capacity headroom and rehearsed recovery. Standards such as BGP roles, RPKI validation, segment-routing policy and MPLS OAM provide important mechanisms, but mechanisms become resilience only when integrated with governance, service priorities and evidence-based exercises. Saudi regulatory controls provide a strong foundation; operators should translate them into measurable routing invariants and continuity outcomes. The enduring objective is not an impossible promise that the core will never fail. It is a demonstrated ability to keep failures bounded, decisions trustworthy, priority traffic moving and recovery controlled under both cyberattack and operational disruption.

REFERENCES

[1] Mrabet, H., Belguith, S., Alhomoud, A. and Jemai, A. (2020). A survey of IoT security based on a layered architecture of sensing and data analysis. Sensors, 20, 3625. https://doi.org/10.3390/s20133625.

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

[3] Communications, Space and Technology Commission (2020). Cybersecurity Regulatory Framework for Service Providers in the Information and Communications Technology Sector. Decision 424/1442.

National Cybersecurity Authority (2020). Cloud Cybersecurity Controls (CCC-1:2020). Riyadh, Saudi Arabia.

[4] Herrera, M., Proselkov, Y., Pérez-Hernández, M. and Parlikad, A.K. (2021). Mining graph-Fourier transform time series for anomaly detection of Internet traffic at core and metro networks. IEEE Access, 9. https://doi.org/10.1109/ACCESS.2021.3050014.

[5] Filsfils, C., Camarillo, P., Leddy, J., Voyer, D., Matsushima, S. and Li, Z. (2021). Segment Routing over IPv6 Network Programming. RFC 8986. https://doi.org/10.17487/RFC8986.

[6] Choi, K., Yi, J., Park, C. and Yoon, S. (2021). Deep learning for anomaly detection in time-series data: review, analysis, and guidelines.

IEEE Access, 9, 120043–120065. https://doi.org/10.1109/ACCESS.2021.3107975.

[7] Internet Engineering Task Force (2022). Network Telemetry Framework. RFC 9232. https://doi.org/10.17487/RFC9232.

[8] Internet Engineering Task Force (2022). Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages. RFC 9234. https://doi.org/10.17487/RFC9234.

[9] Internet Engineering Task Force (2022). Segment Routing Policy Architecture. RFC 9256. https://doi.org/10.17487/RFC9256.

[10] Internet Engineering Task Force (2022). Operations, Administration, and Maintenance in Segment Routing over IPv6. RFC 9259. https://doi.org/10.17487/RFC9259.

[11] Internet Engineering Task Force (2022). The Use of maxLength in the Resource Public Key Infrastructure. RFC 9319. https://doi.org/10.17487/RFC9319.

[12] Internet Engineering Task Force (2022). Policy Based on the Resource Public Key Infrastructure without Route Refresh. RFC 9324. https://doi.org/10.17487/RFC9324.

[13] Internet Engineering Task Force (2022). TCP Authentication Option Test Vectors. RFC 9235. https://doi.org/10.17487/RFC9235.

[14] Wirtgen, T. and Bonaventure, O. (2022). A first step towards checking BGP routes in the dataplane. Proceedings of the ACM SIGCOMM Workshop on Future of Internet Routing and Addressing, 50–57. https://doi.org/10.1145/3527974.3545723.

[15] Hlavacek, T., Jeitner, P., Mirdita, D., Shulman, H. and Waidner, M. (2022). Stalloris: RPKI downgrade attack. 31st USENIX Security Symposium, 4455–4471.

[16] National Cybersecurity Authority (2022). Data Cybersecurity Controls (DCC-1:2022). Riyadh, Saudi Arabia.

[17] National Cybersecurity Authority (2022). Operational Technology Cybersecurity Controls (OTCC-1:2022). Riyadh, Saudi Arabia.

[18] Digital Government Authority (2023). Guideline of Business Continuity Management in Digital Government, Version 2.0. Riyadh, Saudi Arabia.

[19] Internet Engineering Task Force (2023). IS-IS Extensions to Support Segment Routing over the IPv6 Data Plane. RFC 9352. https://doi.org/10.17487/RFC9352.

[20] Internet Engineering Task Force (2023). OSPFv3 Extensions for Segment Routing over IPv6. RFC 9513. https://doi.org/10.17487/RFC9513.

[21] Internet Engineering Task Force (2023). Export of Segment Routing over IPv6 Information in IP Flow Information Export. RFC 9487. https://doi.org/10.17487/RFC9487.

[22] Bordeau-Aubert, K., Whatley, J., Nadeau, S., Glatard, T. and Jaumard, B. (2023). Classification of anomalies in telecommunication network KPI time series. arXiv:2308.16279.

[23] National Institute of Standards and Technology (2024). The NIST Cybersecurity Framework 2.0. NIST CSWP 29. https://doi.org/10.6028/NIST.CSWP.29.

[24] National Cybersecurity Authority (2024). Essential Cybersecurity Controls (ECC-2:2024). Riyadh, Saudi Arabia.

[25] Digital Government Authority (2024). Digital Government Policies, Version 2.0. Riyadh, Saudi Arabia.

[26] Koumar, J., Hynek, K., Čejka, T. and Šiška, P. (2024). CESNET-TimeSeries24: time series dataset for network traffic anomaly detection and forecasting. arXiv:2409.18874.

[27] Internet Engineering Task Force (2024). Egress Validation in Label Switched Path Ping and Traceroute Mechanisms. RFC 9655. https://doi.org/10.17487/RFC9655.

[28] Saudi Vision 2030 (2025). Vision 2030 Annual Report 2024. Riyadh, Saudi Arabia.

[29] National Institute of Standards and Technology (2025). Implementing a Zero Trust Architecture. NIST Special Publication 1800-35. https://doi.org/10.6028/NIST.SP.1800-35.