

Advanced Network Security and Cloud Infrastructure Protection for Future Digital Economies

ABSTRACT:

Future digital economies depend on cloud platforms, software-defined networks, application programming interfaces, connected devices, and distributed workforces. This interdependence expands economic opportunities but additionally transforms identity compromise, configuration errors, vulnerable software components, and service disruptions into systemic risks. This review critically synthesizes research published from 2020 to 2025 on advanced network security and cloud infrastructure protection. It analyzes the integration of zero trust, continuous authorization, artificial-intelligence-assisted detection, cloud-native controls, DevSecOps, encryption, plus resilience engineering, rather than their isolated implementation. A structured review protocol, guided by PRISMA 2020, was applied to thirty peer-reviewed and credible sources selected for topical relevance, methodological clarity, and direct contributions to network or cloud security. The evidence was coded across threat, control, implementation, governance, and economic-resilience dimensions. Evidence suggests that zero trust provides a valuable architectural discipline, but its effectiveness depends on trustworthy telemetry, identity lifecycle governance, microsegmentation, policy automation, and human monitoring. Cloud protection is most effective when configuration assurance, workload security, software-supply-chain controls, and data protection are embedded throughout development and operations. Artificial intelligence enhances triage and anomaly discovery, yet model opacity, hostile manipulation, and false positives limit independent decision-making. The review proposes an adaptive security fabric that connects identity, network, workload, data, and recovery controls to measurable effects such as service continuity, regulatory assurance, innovation velocity, and digital trust. It concludes that future digital economies require interoperable security architectures, evidence-based metrics, skilled governance, and security-by-design investment across organizational and national ecosystems.

Keywords: zero trust architecture; cloud-native security; network security; DevSecOps; artificial intelligence; cyber resilience; digital economy; shared responsibility

1. Introduction

1.1 Context and research gap

Digital economies are progressively linked to cloud infrastructure. Payments, public services, industrial production, health platforms, education, logistics, and artificial-intelligence services now depend on distributed computing and connectivity. While this transition lowers entry barriers and accelerates innovation, it also concentrates operational reliances on identity providers, cloud control planes, application programming interfaces, container registries, telecommunications networks, and third-party software. Security thus assumes an economic role: it safeguards not only confidentiality, integrity, and availability, but also transactional confidence, market participation, continuity of vital services, and the capacity to innovate free of unacceptable systemic exposure.

Users connect from unmanaged locations, workloads are created and destroyed dynamically, data moves between regions, and machine identities frequently outnumber human accounts. Cloud services also divide responsibility between providers and customers, which can leave important configuration, identity, and data-protection duties misunderstood. Reviews of cloud security consistently identify access control, misconfiguration, multi-tenancy, data leakage, and governance as persistent concerns (Parast et al., 2022; Soveizi et al., 2023; Sun, 2020). Migration itself introduces organizational and technical risk when dependencies, controls, and accountability continue to be not assessed before deployment (Maniah et al., 2022; Al Hadwer et al., 2021).

The literature provides several promising approaches, such as zero trust architecture, microsegmentation, endpoint detection and response, cloud security posture management, confidential processing, policy as code, and DevSecOps. However, these solutions expose a divide between technical research and the integrated operating model necessary for future digital economies. The main challenge is both architectural and socio-technical: controls must operate collaboratively across identities, devices, networks, workloads, data, software pipelines, providers, and recovery processes.

60 **1.2 Aim and objectives**

61 This review fills this gap by synthesizing recent evidence into a unified protection model. The
62 aim is to determine how advanced network security and cloud infrastructure controls can be
63 integrated to support secure, resilient, and broad digital economic growth. Four objectives guide
64 the study.

- 65 • First, it classifies converging threat surfaces across identity, network, cloud-native workloads,
66 data, and software supply chains.
- 67 • Second, it evaluates the strengths and limitations of zero trust, artificial-intelligence-assisted
68 detection, cloud-native security, and DevSecOps.
- 69 • Third, it develops an adaptive security fabric linking technical controls to business continuity,
70 compliance, privacy, and digital trust.
- 71 • Fourth, it identifies governance, measurement, and research priorities for organizations and
72 decision-makers.

73 In contrast to sector-specific reviews focused on small and medium-sized enterprises, this paper
74 studies interconnected digital economies as a wider system. It applies a review logic that links
75 organized methodology, thematic analysis, implementation guidance, policy implications, and an
76 agenda for future research.

77

78 **2. Review methodology**

79 **2.1 Review design**

80 A structured review methodology was employed due to the topic's breadth across computer
81 networks, cloud computing, software engineering, cybersecurity governance, and digital-
82 economy resilience. Reporting logic was informed by PRISMA 2020, which emphasizes
83 transparent identification, screening, eligibility, and inclusion decisions (Page et al., 2021). The
84 review intended at providing a critical narrative synthesis rather than a statistical meta-analysis.
85 The included studies examined heterogeneous architectures, datasets, threat models, maturity

levels, and performance measures; therefore, pooling effect sizes would have produced misleading precision.

2.2 Search strategy and eligibility

Searches covered publications dated January 2020 through December 2025. The core discovery sources were Scopus, Web of Science, IEEE Xplore, ScienceDirect, SpringerLink, and publisher databases, with Google Scholar used only to identify additional cited works and credible standards. Search combinations included: “zero trust” AND cloud; “network security” AND digital economy; cloud-native security; container security; DevSecOps; continuous authorization; cloud misconfiguration; machine learning AND cloud security; 5G OR 6G AND zero trust; and software supply chain AND cloud.

Studies were included when they: were published in English between 2020 and 2025; presented empirical findings, a systematic review, an architectural contribution, or an authoritative technical standard; directly addressed advanced network security, cloud infrastructure protection, or their integration; and described enough method or technical detail to support appraisal. Items were excluded when they were primarily promotional, duplicated earlier versions, lacked a clear security contribution, focused only on generic privacy attitudes, or discussed pre-cloud perimeter controls without relevance to distributed environments. Thirty sources formed the final analytical corpus.

2.3 Quality appraisal and evidence synthesis

Quality appraisal considered publication venue, transparency of research design, clarity of assumptions, reproducibility of technical claims, acknowledgement of limitations, and pertinence beyond a single product. Empirical studies were assessed for dataset realism, baseline selection, and whether reported accuracy was accompanied by operational considerations such as latency, false positives, or deployment cost.

A structured extraction matrix was used to record threat domain, protected asset, architecture, control mechanism, evidence type, implementation dependency, limitation, and stated outcome. Thematic coding grouped findings into five categories: identity-centric and zero-trust control;

intelligent detection and response; cloud-native workload and configuration protection; secure software delivery; and governance for resilience. The synthesis followed a configuration logic, examining which combinations of controls are necessary under different trust, scale, and operating conditions, rather than identifying a single optimal technology. The protocol combines PRISMA-informed screening with thematic synthesis, avoids unsupported numerical claims, and defines a bounded corpus that can be audited against explicit eligibility criteria.



Figure 1. Structured review and synthesis protocol used in this study. Source: Authors' synthesis.

3. Evolving threat environment

3.1 Interconnected identity and cloud control planes

Cloud applications communicate through APIs, users authenticate through federated identity services, devices connect over public and private networks, and organizations depend on shared platforms and open-source components. This architecture increases scale, but it also allows a local weakness to become a cross-domain incident. They can exploit one weak credential,

exposed secret, vulnerable dependency, permissive role, or misconfigured storage service and then use legitimate administrative pathways to expand control.

Password theft, token replay, session hijacking, consent phishing, and abuse of dormant privileges can bypass perimeter devices because the resulting activity appears authenticated. Machine identities create an additional problem: service accounts, certificates, workload tokens, and API keys are frequently long-lived, over-privileged, or poorly inventoried. Zero-trust guidance responds by rejecting network location as sufficient evidence of trust and demanding explicit, context-sensitive authorization for each resource interaction (Rose et al., 2020; Syed et al., 2022). However, continuous verification is only as reliable as the identity data, device posture, and telemetry that direct policy decisions.

Cloud control planes are attractive targets because they can modify networks, workloads, identities, keys, and backups at scale. Public storage, permissive security groups, weak cross-account trust, unmanaged access keys, and disabled logging can expose assets without exploiting a software vulnerability. The changing nature of cloud based systems compounds the issue: containers and serverless functions may exist briefly, infrastructure is declared through code, and production states can drift from approved templates. Systematic studies of cloud based services emphasize that visibility, isolation, dependency management, and uniform policy enforcement remain difficult across rapidly changing components (Theodoropoulos et al., 2023; Rahaman et al., 2023).

3.2 Software-supply-chain and recovery risk

Dependencies, build systems, package repositories, container images, deployment tools, and managed services can introduce malicious or vulnerable code. A signed artifact may still be unsafe if the build environment was compromised, while a clean source repository does not guarantee that the deployed binary is authentic. Collateral damage can affect many customers simultaneously, making provenance, build isolation, vulnerability management, and rapid revocation critical parts of cloud protection.

Ransomware and destructive attacks increasingly target operational recovery as well as primary systems. Adversaries attempt to disable monitoring, encrypt or delete backups, steal data for extortion, and exploit trusted remote-management channels. Endpoint detection and response

systems can improve visibility, but empirical assessment shows that advanced attack techniques may still evade or degrade detection when telemetry, rules, or response coverage are incomplete (Karantzas and Patsakis, 2021). Resilience therefore requires protected backups, isolated recovery credentials, tested restoration procedures, and the ability to operate critical functions in a degraded mode.

3.3 Intelligent and telecommunications threats

Defenders use machine learning to identify anomalies in logs, network flows, authentication events, and cloud workloads. Reviews report useful potential for classification and prioritization, especially where event volume exceeds manual capacity (Nassif et al., 2021; Mohammad and Pradhan, 2021). Attackers, however, can automate reconnaissance, generate convincing social-engineering content, mutate malicious behavior, or manipulate models and their inputs. False positives exhaust analysts, false negatives create unjustified confidence, and opaque scores complicate accountability.

Fifth- and sixth-generation networks distribute functions across software-defined, virtualized, and edge environments. Network slicing, open interfaces, and multi-access edge computing support new services but also form dynamic trust relationships among operators, vendors, tenants, devices, and applications. Research on intelligent zero trust for 5G and 6G suggests that continuous assurance, automated policy, and machine learning can strengthen these environments, while also introducing challenges concerning latency, interoperability, and adversarial robustness (Ramezanpour and Jagannath, 2022; Nahar et al., 2024).

186

187

188

189 **Table 1. Converging threat domains, integrated controls, and digital-economy**
 190 **consequences.**

Threat domain	Typical failure mode	Integrated control set	Economic consequence
Identity and access	Stolen tokens, excessive privilege, weak machine identities	Phishing-resistant authentication; short-lived credentials; continuous authorization; privileged-access controls (Rose et al., 2020)	Fraud, unauthorized transactions, and loss of platform confidence
Cloud control plane	Public exposure, policy drift, over-permissive roles	Posture management; policy as code; protected audit logs; separation of duties (Parast et al., 2022)	Cross-service compromise and regulatory exposure
Cloud-native workloads	Vulnerable images, orchestrator abuse, weak isolation	Signed artifacts; admission policies; workload identity; runtime monitoring (Flauzac et al., 2020)	Production disruption and compromised shared services
Software supply chain	Malicious dependencies, compromised builds, unverifiable artifacts	Provenance; software bills of materials; isolated builds; rapid revocation	Correlated compromise across customers and sectors
Data and cryptography	Leakage, key misuse, unauthorized processing	Classification; encryption; key separation; fine-grained authorization	Privacy harm, intellectual-property loss, and cross-border restrictions

Availability and recovery	Ransomware, destructive administration, backup compromise	Segmentation; endpoint detection; immutable backups; rehearsed restoration (Karantzas and Patsakis, 2021)	Service outage, financial loss, and reduced public trust
---------------------------	---	---	--

191

192 4. Advanced network security architectures

193 4.1 Zero-trust decision and enforcement

194 NIST defines it as a shift from static, network-based perimeters toward protection of users,
 195 assets, and resources, with authentication and authorization performed before access is
 196 established (Rose et al., 2020). It is a decision discipline: every request is evaluated using
 197 identity, device, resource, environmental, and behavioral context; access is limited to the
 198 minimum required; and trust is continuously reconsidered. Surveys agree that this approach can
 199 reduce lateral movement and implicit trust, but they also identify deployment complexity, legacy
 200 integration, performance, and policy consistency as major obstacles (He et al., 2022; Sarkar et
 201 al., 2022).

202 Decision components aggregate identity assurance, device posture, asset sensitivity, session risk,
 203 and threat intelligence. Enforcement components apply the resulting decision at gateways,
 204 software-defined network controls, service meshes, endpoints, and application layers. Cloud-
 205 native access requires further granularity because workloads communicate through ephemeral
 206 identities and east-west traffic. NIST guidance for cloud-native applications emphasizes identity-
 207 tier and network-tier policies across multi-cloud environments, including application and service
 208 identities rather than only human individuals (Chandramouli and Butcher, 2023).

209 4.2 Microsegmentation and continuous authorization

210 Microsegmentation limits the blast radius by dividing resources into policy-controlled zones or
 211 service-level trust boundaries. Conventional segmentation based on subnets is often too coarse
 212 for containers, serverless functions, and dynamic scaling. Software-defined enforcement can use
 213 labels, workload identities, application context, and data sensitivity. Policies should describe
 214 permitted business flows, deny unforeseen paths, and be tested as code. During an incident,

dynamic segmentation can isolate affected devices or workloads while safeguarding essential services.

Context changes: a device can become noncompliant, a token can be used from an unusual location, an account can display impossible travel, or a workload can begin communicating with an unapproved destination. Cloud-based zero-trust access studies show the value of fine-grained controls for remote work, but also underline dependency on reliable identity and policy infrastructure (Mandal et al., 2021).

4.3 Automated and intelligent response

Orchestration can correlate identity, endpoint, network, and cloud signals and then trigger proportionate actions such as step-up authentication, token revocation, workload quarantine, or restricted access. Cao et al. (2024) argue that automation can make zero trust more adaptive, while stressing interoperability, data quality, and orchestration risk. Consequently, high-impact actions should use confidence thresholds, staged containment, rollback mechanisms, and human approval where uncertainty remains material.

Models can establish behavioral baselines, prioritize alerts, detect unusual privilege use, and identify relationships that static rules miss. Intelligent zero-trust research combines machine learning with policy engines to adapt access decisions to changing risk (Ajish, 2024; Hussain et al., 2024). Secure deployment requires explainable features, monitored drift, adversarial testing, privacy-preserving data treatment, and clear ownership of model outcomes. Federated approaches may allow multiple domains to improve models without centralizing raw security data, but they introduce poisoning, aggregation, and trust challenges.

4.4 Telemetry and managed trust

Endpoint and extended detection complete the architecture by observing activity after access is granted. Identity controls cannot prevent all malicious use of legitimate tools, vulnerable applications, or insider actions. Endpoint telemetry, network detection, cloud audit records, and application traces provide overlapping perspectives. Correlation across these sources can reveal a sequence that appears harmless in isolation: a new device registration, privilege escalation, unusual data access, and outbound transfer. The value lies less in collecting every event than in preserving high-quality, time-synchronized evidence linked to assets and identities.

Campbell (2020) argues that trust itself can be a vulnerability when it becomes an unexamined assumption. The best architecture treats every trust relationship as a managed risk with evidence, contractual responsibility, technical enforcement, and a recovery plan (Dhar and Bose, 2021).

5. Cloud infrastructure protection

5.1 Shared responsibility, configuration, identity, and data

Cloud infrastructure protection begins with an accurate shared-responsibility model. Providers secure physical facilities, foundational infrastructure, and specified managed-service layers, while customers remain responsible for identities, data, configurations, workloads, and application behavior to varying degrees. Organizations should therefore maintain a control responsibility matrix for every service, identifying the owner, implementation mechanism, evidence source, and escalation path.

Cloud security posture management and broader cloud-native application protection platforms can inventory resources, identify risky permissions, detect public exposure, compare settings with policy, and prioritize remediation. A technically noncompliant setting may be accepted for a documented business reason, while a compliant configuration may still be dangerous when combined with an unexpected trust relationship. Effective posture management therefore links findings to asset criticality, data classification, exploitability, exposure, and compensating controls. Infrastructure-as-code scanning should find issues before deployment, and runtime monitoring should identify drift afterward.

Privileged roles should be time-bound, purpose-limited, strongly authenticated, and separated from routine user accounts. Workloads should use short-lived, automatically rotated identities rather than embedded secrets. Cross-account and cross-tenant trust should be minimized and continuously reviewed. Cryptographic keys should be protected via dedicated management services, separation of duties, and auditable rotation. Data should be encrypted in transit and at rest, while highly sensitive use cases may require protection during processing through trusted execution environments or other confidential-computing techniques.

5.2 Cloud-native workload and orchestration protection

Containers share host resources and inherit risk from base images, registries, kernels, orchestration settings, and runtime privileges. Research on native container security identifies isolation, image integrity, vulnerability management, and runtime monitoring as central concerns (Flauzac et al., 2020). Later analysis similarly emphasizes layered precautions across build, deployment, and execution rather than dependence on a single scanner (Devi Priya et al., 2023). Minimum viable controls include trusted base images, software bills of materials, signed artifacts, admission policies, non-root execution, resource restrictions, secret isolation, network policies, and runtime behavioral detection.

A compromised control plane, permissive role, exposed dashboard, or malicious admission component can affect many workloads. Security should be anchored in hardened administrative access, namespace and tenant isolation, least-privilege roles, encrypted secrets, protected audit records, and policy enforcement at admission and runtime. Service meshes can provide mutual authentication and traffic policy between services, but they add operational complexity and should not obscure application-level authorization.

5.3 DevSecOps and software assurance

A systematic review by Rajapakse et al. (2022) identifies social resistance, skills gaps, tool integration, and unclear responsibility as recurring adoption challenges. A later multivocal review organizes DevSecOps around multiple dimensions, reinforcing that it is an operating model rather than a collection of pipeline tools (Zhao et al., 2024). Mature implementation combines threat modeling, code review, dependency analysis, secret scanning, infrastructure-as-code checks, artifact signing, deployment policy, runtime feedback, and rapid remediation. Gates should be risk-based: blocking every low-severity defect can slow delivery without materially reducing exposure, whereas exceptions for serious vulnerabilities should require accountable approval and expiry.

Rahaman et al. (2023) find that static-analysis research for cloud-native systems addresses important security challenges, yet tool fragmentation and coverage gaps remain. Controls should therefore be layered with dynamic testing, API testing, configuration validation, runtime protection, and manual review for high-risk logic.

5.4 Multi-cloud visibility and resilience

Organizations may use multiple providers for capability, regulation, resilience, or commercial leverage, but security controls and logging formats differ. A common control taxonomy can define required outcomes while provider-specific implementations satisfy them. Central visibility should not become a single point of failure; critical logs need durable, access-controlled storage outside the protected workload's administrative boundary. Recovery designs should account for region failure, identity-provider disruption, corrupted infrastructure code, and provider-specific dependencies.

6. Integrated protection model for digital economies

6.1 Fabric layers and service mapping

The evidence supports an adaptive security fabric rather than a sequence of independent defensive products. The fabric has four interacting layers. The observation layer collects identity events, endpoint state, network flows, cloud control-plane actions, application traces, data-access events, vulnerability information, and threat intelligence. The decision layer converts evidence into risk-informed policy through rules, analytics, and accountable human decision. The enforcement layer applies least privilege, microsegmentation, workload isolation, secure configuration, encryption, pipeline controls, and response actions. Resilience layer preserves critical services through redundancy, protected backups, tested restoration, crisis communication, and alternative operating procedures.

Integration should be guided by business services rather than technology inventories. A digital payment service, for example, depends on customer identity, APIs, databases, cryptographic keys, network connectivity, fraud monitoring, cloud regions, and external partners. Charting these dependencies reveals where preventive, detective, and recovery controls must intersect. It also supports meaningful metrics. Counts of blocked attacks or scanned assets can be useful operationally, but they do not prove resilience. Better indicators include coverage of strong authentication for privileged actions, percentage of high-risk service flows explicitly authorized, time to revoke compromised credentials, exposure time of critical misconfigurations, restoration time for priority services, and proportion of production artifacts with verified provenance.




Figure 2. Adaptive security fabric for future digital economies. Source: Authors' synthesis.

6.2 Governance and ecosystem accountability

Governance translates this architecture into sustained practice. Executive leaders set risk appetite and ensure that security investment protects economic priorities. Security teams define common controls and assurance methods. Platform and network teams build secure default services. Product owners classify business impact and accept residual risk. Developers and operators maintain controls through the lifecycle. Internal audit or independent assurance tests whether evidence supports claims. This distribution avoids treating cybersecurity as a specialist function detached from service design.

At ecosystem level, cloud providers, telecommunications operators, software suppliers, regulators, and customers share dependencies. Contracts should identify security responsibilities, notification duties, evidence rights, data location, subcontractor controls, recovery expectations, and exit arrangements. Sector-wide incident exercises can test coordination before a crisis. Common reporting formats can accelerate threat sharing without exposing unnecessary personal

or commercial data. Standards and certification are valuable when they improve comparable assurance, but checklist compliance should not replace threat-led testing.

6.3 Proportional implementation

The model also supports proportionality. Small organizations and public institutions may not operate large security operations centers, yet they still require strong identity, secure configuration, protected backups, and incident support. Managed services and secure cloud defaults can democratize advanced protection, provided customers retain visibility and contractual accountability. Risk-based implementation ought to prioritize services whose disruption would create the greatest economic, safety, or trust impact. Table 2 presents a phased roadmap originating from foundational control to adaptive, ecosystem-level resilience.

369 **Table 2. Phased execution roadmap for the adaptive security fabric.**

Horizon	Priority actions	Evidence and metrics	Target maturity
0–90 days	Map critical services, owners, identities, data, approved flows, recovery dependencies, and provider responsibilities. Remove exposed secrets and enforce strong privileged authentication.	Critical-service coverage; privileged MFA coverage; exposed-secret age; verified recovery owners.	Visible foundations
3–9 months	Introduce policy-as-code, cloud posture management, centralized durable logging, workload identity, secure deployment templates, and risk-based vulnerability gates.	Misconfiguration exposure time; log coverage; signed-artifact rate; high-risk remediation time.	Consistent control
9–18 months	Expand microsegmentation, continuous authorization, cross-domain detection, automated containment with rollback, and tested regional or provider recovery.	Unauthorized-flow reduction; credential revocation time; mean containment time; restoration performance.	Integrated resilience
Continuous	Exercise cross-sector incidents, review model drift and policy conflicts, test suppliers, update cryptographic inventory, and report service-level resilience to governance bodies.	Exercise findings closed; assurance exceptions expired; supplier evidence coverage; service resilience trend.	Adaptive ecosystem

370

371

7. Discussion

7.1 Consolidated findings

The synthesis indicates that future digital-economy security is determined by control interaction, not product accumulation. Zero trust reduces implicit access, but it cannot compensate for incomplete asset inventories or compromised policy infrastructure. Artificial intelligence accelerates analysis, but it depends on representative data, interpretable outputs, and disciplined response. Cloud posture tools reveal misconfiguration, but remediation may fail when ownership is unclear. DevSecOps finds defects earlier, but cultural and incentive problems can encourage superficial compliance. Resilience plans appear credible until restoration is tested under realistic identity, dependency, and communication restrictions.

A first major finding is that identity, network, and workload security are converging. Identity-aware access without network containment allows an authorized compromise to move broadly. Network segmentation without workload identity becomes brittle inside dynamic cloud environments. Workload controls without user and administrative governance leave the control plane exposed. An effective architecture therefore evaluates both the requesting principal and the requested resource, while enforcing policy close to the transaction. The emerging cloud-native zero-trust literature supports this twofold focus, particularly for multi-cloud applications and service-to-service communication (Chandramouli and Butcher, 2023; Liu et al., 2024).

A second finding is that automation must be bounded by assurance. Digital economies operate at a speed and scale that make manual control impossible, yet indiscriminate automation has the potential to amplify errors. An incorrectly propagated policy may block a national service, while a poisoned detection model might systematically misclassify activity. Safe automation requires tested decision logic, confidence-aware actions, observability, rollback, and clear human authority. Wang et al. (2025) demonstrate the research direction toward dynamic zero-trust access control, but operational adoption must also address policy conflict, explainability, and failure behavior.

A third finding is that cloud security should be treated as a lifecycle capability. It continues through secure architecture, infrastructure as code, software delivery, runtime monitoring, and recovery. Cloud adoption research shows that organizational readiness and governance materially

guide outcomes, not simply technical features (Al Hadwer et al., 2021). Likewise, migration risk reviews show the importance of systematic analysis before moving critical workloads (Maniah et al., 2022). This questions the common assumption that purchasing a managed cloud service transfers security responsibility.

A fourth finding concerns economic robustness. Security controls create value when they preserve reliable participation in digital markets. Strong identity reduces fraud and enables remote transactions. Secure APIs support platform ecosystems. Verified software supply chains improve trust between suppliers and customers. Tested recovery limits service interruption. Transparent governance supports regulatory confidence and cross-border cooperation. These outcomes are broader than breach prevention; they influence innovation velocity, investment, inclusion, and public willingness to use digital services.

7.2 Trade-offs and evidence limitations

Uninterrupted monitoring can conflict with privacy and workforce expectations. Strong segmentation may increase administrative burden and hinder troubleshooting. Multi-cloud strategies may reduce concentration risk while increasing configuration complexity. Strict pipeline gates may improve assurance but delay urgent fixes or product delivery. Encryption and confidential processing can add cost and latency. The right response is not to avoid controls, but to make trade-offs explicit, evidence-based, and revisable. Risk acceptance should state the affected service, expected benefit, exposure, compensating control, accountable owner, and review date.

The literature also has limitations. Many studies evaluate proposed architectures in simulations, testing environments, or narrow datasets. Security-vendor terminology changes rapidly, making comparisons difficult. Reported machine-learning accuracy is often not accompanied by deployment latency, analyst workload, model drift, or adversarial evaluation. Zero-trust studies vary in scope: some address remote access, others microsegmentation, device authentication, or telecommunications. These differences prevent simple ranking. They do, however, enable a more useful conclusion: successful protection requires a context-specific configuration of complementary mechanisms.

430

431 **7.3 Contribution and implications**

432 This review pushes the field forward by moving systematically from converging threats to
433 architectural controls, implementation limitations, governance choices, and research priorities. It
434 separates evidence appraisal from thematic synthesis, avoids unsupported screening statistics,
435 and connects technical mechanisms to measurable economic outcomes. The resulting systems
436 perspective encompasses cloud providers, telecommunications operators, software ecosystems,
437 public institutions, commercial organizations, and communities which depend on reliable digital
438 services across regional, national, and interdependent cross-border economic relationships.

439 For practitioners, the immediate implication is to begin with service dependency and trust
440 mapping. Every critical service should have an owner, an identity model, approved
441 communication paths, protected data flows, monitored control-plane actions, verified software
442 inputs, and a tested recovery method. For decision makers, the implication is to promote baseline
443 controls, incident coordination, workforce development, and transparent provider accountability
444 free of freezing innovation through overly prescriptive technology mandates. For researchers, the
445 priority is stronger real-world evidence that compares integrated architectures within operational
446 constraints.

447

448 **8. Research agenda**

449 Research from 2020 to 2025 establishes a strong architectural foundation, but several questions
450 remain open. First, continuous authorization needs common evaluation methods. Studies should
451 report policy accuracy, decision latency, service disruption, privacy impact, and resistance to
452 identity or telemetry manipulation. Testbeds should include multi-cloud, edge, and intermittent-
453 connectivity conditions rather than idealized enterprise networks.

454 Second, intelligent security requires trustworthy learning. Subsequent research should examine
455 federated and privacy-preserving analytics, explainable risk scores, adversarially robust models,
456 and methods for detecting poisoned telemetry. Evaluation should include analyst decision quality

and workload, not only classifier accuracy. Longitudinal studies are needed to understand model drift as organizations, users, and attackers change.

Third, software-supply-chain assurance needs usable provenance across source, build, artifact, deployment, and runtime. Research should test how signing, attestation, software bills of materials, and isolated builds perform in complicated dependency ecosystems. Economic studies should measure the cost and benefit of rapid vulnerability disclosure, coordinated patching, and supplier assurance for smaller organizations.

Fourth, post-quantum migration must be integrated with cloud and network modernization. The research problem is not simply selecting algorithms; it includes cryptographic inventory, protocol agility, key lifecycle, interoperability, performance, and long-lived data exposure. Fifth, resilience metrics should connect technical recovery to social and economic outcomes. Cross-sector exercises and anonymized incident datasets could reveal cascading dependencies that single-organization studies miss.

Finally, future-digital-economy research should represent varied regions and institutional capacities. Security architectures designed for highly resourced enterprises may not transfer directly to public services, small firms, or infrastructure with limited connectivity. Comparative studies should evaluate secure defaults, managed protection, workforce models, and regulatory approaches that make advanced security accessible without creating opaque dependency on a few providers.

9. Conclusion

Advanced network security and cloud infrastructure protection are basic capabilities for future digital economies. The review shows that no isolated control is sufficient. Zero trust contributes a disciplined approach to explicit, least-privilege, continuously evaluated access. Artificial intelligence helps interpret high-volume telemetry, but requires explainability, adversarial testing, and human accountability. Cloud-native protection depends on secure configuration, workload identity, container and orchestration controls, protected data, and continuous visibility. DevSecOps extends assurance across the software lifecycle, while resilience engineering assures that key services can recover when prevention fails.

The proposed adaptive security fabric integrates observation, decision, enforcement, and resilience around critical business services. Its value is measured through reduced exposure, faster containment and recovery, verified software provenance, reliable authorization, regulatory assurance, and sustained digital trust. Implementation should be phased, risk-based, and proportional to service impact. Organizations must make trust relationships explicit, assign control ownership, test automation, and protect recovery mechanisms. Providers and policymakers have to support interoperable standards, clear shared responsibility, incident collaboration, and access to skilled security capabilities.

Subsequent research should move beyond isolated prototypes toward operational comparisons of integrated architectures. It should examine dynamic authorization, trustworthy artificial intelligence, software-supply-chain provenance, post-quantum transition, privacy-preserving monitoring, and cross-sector resilience. A secure digital economy is not created by eliminating risk; it is created by making dependencies visible, constraining failure, learning quickly, and preserving confidence when systems are stressed. Security therefore belongs within economic strategy, service design, and institutional governance, not at the edge of technology operations.

References

- Ajish, D. (2024). The significance of artificial intelligence in zero trust technologies: A comprehensive review. *Journal of Electrical Systems and Information Technology*, 11, Article 30. <https://doi.org/10.1186/s43067-024-00155-z>
- Al Hadwer, A., Tavana, M., Gillis, D., & Rezania, D. (2021). A systematic review of organizational factors impacting cloud-based technology adoption using the technology–organization–environment framework. *Internet of Things*, 15, 100407. <https://doi.org/10.1016/j.iot.2021.100407>
- Campbell, M. (2020). Beyond zero trust: Trust is a vulnerability. *Computer*, 53(10), 110–113. <https://doi.org/10.1109/MC.2020.3011081>
- Cao, Y., Pokhrel, S. R., Zhu, Y., Doss, R., & Li, G. (2024). Automation and orchestration of zero trust architecture: Potential solutions and challenges. *Machine Intelligence Research*, 21(2), 294–317. <https://doi.org/10.1007/s11633-023-1456-2>

- Chandramouli, R., & Butcher, Z. (2023). A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Location EnvironmentsNIST Special Publication 800-207A). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-207A>
- Devi Priya, V. S., Sethuraman, S. C., & Khan, M. K. (2023). Container security: Precaution levels, mitigation strategies, and research perspectives. *Computers & Security*, 135, 103490. <https://doi.org/10.1016/j.cose.2023.103490>
- Dhar, S., & Bose, I. (2021). Securing IoT devices using zero trust and blockchain. *Journal of Organizational Computing and Electronic Commerce*, 31(1), 18–34. <https://doi.org/10.1080/10919392.2020.1831870>
- Flauzac, O., Mauhourat, F., & Nolot, F. (2020). A review of native container security for running applications. *Procedia Computer Science*, 175, 157–164. <https://doi.org/10.1016/j.procs.2020.07.025>
- He, Y. H., Huang, D. C., Chen, L., Ni, Y., & Ma, X. J. (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>
- Hussain, M., Pal, S., Jadidi, Z., Foo, E., & Kanhere, S. (2024). Federated zero trust architecture using artificial intelligence. *IEEE Wireless Communications*, 31(2), 30–35. <https://doi.org/10.1109/MWC.001.2300405>
- Karantzas, G., & Patsakis, C. (2021). An empirical assessment of endpoint detection and response systems against advanced persistent threats attack vectors. *Journal of Cybersecurity and Privacy*, 1(3), 387–421. <https://doi.org/10.3390/jcp1030021>
- Liu, C., Tan, R., Wu, Y., Feng, Y., Jin, Z., Zhang, F., Liu, Y., & Liu, Q. (2024). Dissecting zero trust: Research landscape and its implementation in IoT. *Cybersecurity*, 7, 20. <https://doi.org/10.1186/s42400-024-00212-0>
- Mandal, S., Khan, D. A., & Jain, S. (2021). Cloud-based zero trust access control policy: An approach to support work-from-home driven by COVID-19 pandemic. *New Generation Computing*, 39, 599–622. <https://doi.org/10.1007/s00354-021-00130-6>

Maniah, Soewito, B., LumbanGaol, F., & Abdurachman, E. (2022). A systematic literature review: Risk analysis in cloud migration. *Journal of King Saud University – Computer and Information Sciences*, 34(6), 3111–3120. <https://doi.org/10.1016/j.jksuci.2021.01.008>

Mohammad, A. S., & Pradhan, M. R. (2021). Machine learning with big data analytics for cloud security. *Computers & Electrical Engineering*, 96, 107527. <https://doi.org/10.1016/j.compeleceng.2021.107527>

Nahar, N., Andersson, K., Schelén, O., & Saguna, S. (2024). A survey on zero trust architecture: Applications and challenges of 6G networks. *IEEE Access*, 12, 94753–94764. <https://doi.org/10.1109/ACCESS.2024.3425350>

Nassif, A. B., Abu Talib, M., Nasir, Q., Albadani, H., & Dakalbab, F. M. (2021). Machine learning for cloud security: A systematic review. *IEEE Access*, 9, 20717–20735. <https://doi.org/10.1109/ACCESS.2021.3054129>

Page, M. J., McKenzie, J. E., Bossuyt, P. M., Boutron, I., Hoffmann, T. C., Mulrow, C. D., et al. (2021). The PRISMA 2020 statement: An updated guideline for reporting systematic reviews. *BMJ*, 372, n71. <https://doi.org/10.1136/bmj.n71>

Parast, F. K., Sindhav, C., Nikam, S., Yekta, H. I., Kent, K. B., & Hakak, S. (2022). Cloud computing security: A survey of service-based models. *Computers & Security*, 114, 102580. <https://doi.org/10.1016/j.cose.2021.102580>

Rahaman, M. S., Islam, A., Cerny, T., & Hutton, S. (2023). Static-analysis-based solutions to security challenges in cloud-native systems: Systematic mapping study. *Sensors*, 23(4), 1755. <https://doi.org/10.3390/s23041755>

Rajapakse, R. N., Zahedi, M., Babar, M. A., & Shen, H. (2022). Challenges and solutions when adopting DevSecOps: A systematic review. *Information and Software Technology*, 141, 106700. <https://doi.org/10.1016/j.infsof.2021.106700>

Ramezanpour, K., & Jagannath, J. (2022). Intelligent zero trust architecture for 5G/6G networks: Principles, challenges, and the role of machine learning in the context of O-RAN. *Computer Networks*, 217, 109358. <https://doi.org/10.1016/j.comnet.2022.109358>

Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-207>

573 Sarkar, S., Choudhary, G., Shandilya, S. K., Hussain, A., & Kim, H. (2022). Security of zero
574 trust networks in cloud computing: A comparative review. *Sustainability*, 14(18), 11213.
575 <https://doi.org/10.3390/su141811213>

576 Soveizi, N., Turkmen, F., & Karastoyanova, D. (2023). Security and privacy concerns in cloud-
577 based scientific and business workflows: A systematic review. *Future Generation Computer*
578 *Systems*, 148, 184–200. <https://doi.org/10.1016/j.future.2023.05.015>

579 Sun, P. J. (2020). Security and privacy protection in cloud computing: Discussions and
580 challenges. *Journal of Network and Computer Applications*, 160, 102642.
581 <https://doi.org/10.1016/j.jnca.2020.102642>

582 Syed, N. F., Shah, S. W., Shaghaghi, A., Anwar, A., Baig, Z., & Doss, R. (2022). Zero trust
583 architecture (ZTA): A comprehensive survey. *IEEE Access*, 10, 57143–57179.
584 <https://doi.org/10.1109/ACCESS.2022.3174679>

585 Theodoropoulos, T., Rosa, L., Benzaid, C., Gray, P., Marin, E., Makris, A., et al. (2023).
586 Security in cloud-native services: A survey. *Journal of Cybersecurity and Privacy*, 3(4), 758–
587 793. <https://doi.org/10.3390/jcp3040034>

588 Wang, R., Li, C., Zhang, K., & Tu, B. (2025). Zero-trust based dynamic access control for cloud
589 computing. *Cybersecurity*, 8, 12. <https://doi.org/10.1186/s42400-024-00320-x>

590 Zhao, X., Clear, T., & Lal, R. (2024). Identifying the primary dimensions of DevSecOps: A
591 multi-vocal literature review. *Journal of Systems and Software*, 214, 112063.
592 <https://doi.org/10.1016/j.jss.2024.112063>