Secure Cloud Platform Governance for Saudi Enterprises under Vision 2030: Integrating SAP BTP Administration, Identity Authentication, Access Control, and Cloud Cybersecurity Compliance.

Abstract

Saudi enterprises increasingly rely on cloud platforms to connect applications, orchestrate services and support national digital transformation. Yet platform adoption creates a governance problem that cannot be solved by authentication technology or contractual assurances alone. This review examines how SAP Business Technology Platform (BTP) administration, federated identity, authorisation and cybersecurity compliance can operate as one accountable control system under Saudi Vision 2030. A structured integrative synthesis of research and technical standards published between 2020 and 2025 was combined with contextual examination of Saudi regulatory instruments and SAP administration guidance. The analysis distinguishes vendor capabilities from enterprise responsibilities, maps identity and entitlement decisions to accountable control owners, and identifies evidence required for continuous assurance. It develops a layered operating framework in which account hierarchy, identity lifecycle, least privilege, integration boundaries, workload protection and regulatory traceability are linked through policy decisions and auditable records. The central finding is conceptual: secure platform operation depends less on accumulating isolated controls than on maintaining consistency between business authority, technical configuration and compliance evidence. Particular risks arise from persistent privileged accounts, mismatched group mappings, unmonitored service identities and ambiguous cloud responsibilities. The proposed framework provides decision criteria for Saudi enterprises, while its effectiveness remains to be evaluated through organisational case studies and longitudinal control testing.

Keywords: *SAP Business Technology Platform (SAP BTP); cloud platform governance; identity authentication; access control; zero-trust security; Saudi Arabia; cybersecurity compliance; Vision 2030*

UNDER PEER REVIEW IN IJAR

1. Introduction

Cloud platforms are becoming integral to enterprise transformation because they allow organisations to extend core systems without recreating every service in a local data centre. For Saudi enterprises, this capability is relevant to Vision 2030 objectives involving digital

services, operational productivity and resilient national infrastructure. However, deployment velocity also enlarges the population of administrators, connected identities, privileged interfaces and externally hosted data. Cloud governance therefore concerns the authority to make changes, the legitimacy of access and the ability to demonstrate that protective obligations have actually been fulfilled. It should not be mistaken for a collection of settings within a vendor console [1,2].

SAP Business Technology Platform supplies integration, application development, automation, data and extension services. A typical enterprise may connect its SAP estate with business applications, identity providers and on-premises workloads while maintaining distinct production and development subaccounts. Security failures in this setting can arise without exploitation of a software vulnerability. For example, an inherited administrative assignment may permit configuration changes long after an employee transfers role; alternatively, a service credential may allow a connector to access more back-end resources than a business process requires. These scenarios expose governance weaknesses at the junction of people, technical privileges and change procedures [3,4].

The scholarly discussion offers substantial knowledge about zero-trust architectures, cloud attack surfaces and identity management, but less agreement on how these ideas should be translated into the administrative units of a particular enterprise platform [5,6]. A generic instruction to enforce least privilege does not settle who may create a BTP subaccount, approve a trust configuration, modify role collections or retain emergency access. Nor does a cybersecurity control catalogue explain which evidence a tenant can obtain directly and which assurance must be requested from a cloud service provider. A review centred on these decision rights addresses a practical integration gap.

The aim of this article is to develop an accountable governance model for secure SAP BTP use by Saudi enterprises, joining platform administration with identity authentication, access control and cybersecurity compliance. Its objectives are to establish the relevant control boundaries; critically compare identity-centred security approaches; translate Saudi cloud and personal-data obligations into tenant activities; and specify testable assurance evidence. The study asks how control ownership should be allocated, where authorisation may drift from business intent, and which indicators distinguish effective governance from nominal policy adoption. It makes no claim to measure security performance in an actual SAP deployment.

UNDER PEER REVIEW IN IJAR

2. Conceptual and regulatory foundations

Zero trust provides a useful conceptual starting point because trust is not automatically established by a private network address, corporate device or previous session. Access decisions should consider the identity requesting a resource, the resource's sensitivity and relevant context. This does not imply continuous reauthentication of every user action; rather, policy should support appropriate verification, constrained sessions and reassessment when conditions materially change. The distinction matters in enterprise platforms where administrators, application users and machine identities use different routes to authority [1,8].

A second distinction concerns the legal status of instruments. International technical guidance may describe prudent security practice without being binding Saudi law; national control documents may be mandatory only for specified entities; and supplier documentation describes available functions rather than customer compliance outcomes. Treating these sources as interchangeable exaggerates the certainty of a control mapping. The register should therefore distinguish an applicable obligation, a voluntarily adopted benchmark and a vendor implementation recommendation, and preserve the rationale for each classification.

A platform governance model must also separate control design from control effectiveness. A rule collection can exist while remaining unreviewed; multifactor authentication can be configured while privileged recovery accounts remain outside its scope; and an audit dashboard can display events without a defined escalation owner. Control catalogues and assessment guidance therefore emphasise objective criteria, evidence, independent review and adaptation to organisational risk appetite. The relevant question is not simply whether a security feature exists, but whether it protects the intended asset under realistic operating conditions [2,7].

Saudi cloud governance adds jurisdiction-specific duties. The National Cybersecurity Authority's Cloud Cybersecurity Controls distinguish requirements for cloud service providers and cloud service tenants; the updated CCC-2:2024 edition also addresses changes involving data localisation. The Essential Cybersecurity Controls establish a wider cybersecurity baseline for entities within scope. These documents should not be read as imposing one identical obligation on every private enterprise: applicability depends on organisational classification, sector, service model and relevant contractual or regulatory requirements. Accordingly, a compliance register must record why each requirement applies, who owns it and what satisfies it.

Personal information creates a parallel governance stream. Saudi Arabia's Personal Data Protection Law became effective in September 2023, followed by the end of a general one-year

UNDER PEER REVIEW IN IJAP

adjustment period in September 2024. Controllers must address lawful processing, defined purposes, access restrictions, retention and relevant cross-border transfer conditions. Security controls contribute to compliance but do not establish the lawful basis for processing by themselves. Where BTP applications exchange customer or employee data, technical owners and privacy specialists must agree on data classifications, permissible recipients and the evidentiary consequences of integration design [4,17].

These layers reveal a fundamental governance distinction. Provider assurance covers the controls operated by the cloud supplier; tenant assurance covers choices made by the enterprise, including identity configuration, entitlements, deployed applications, interfaces and data handling. Shared responsibility does not mean that every control is jointly operated. It means boundaries must be explicitly documented, including dependencies when evidence crosses those boundaries. Ambiguous responsibilities otherwise create unowned gaps precisely where auditors expect an identifiable decision maker [12,19].

3. Review design and analytical method

A structured integrative review was used because the question combines platform-specific administrative practices, conceptual cybersecurity research and nationally applicable governance instruments. The review was neither a statistical meta-analysis nor an experimental study. Sources were selected from peer-reviewed publications and publicly available technical standards dated 2020–2025, with emphasis on access governance, zero trust, cloud-native protection, software security and assurance. Official NCA and Saudi data-protection instruments and SAP product documentation were consulted as primary implementation context rather than represented as peer-reviewed studies. The supporting bibliography contains thirty DOI-bearing works.

Selection prioritised sources with an identifiable publisher, publication date, persistent identifier and direct relevance to at least one governance function. Conceptually redundant papers were not treated as independent proof of effectiveness. Publications focused exclusively on cryptographic construction were used sparingly because their laboratory assumptions do not establish enterprise governance outcomes. The synthesis employed four interpretive categories: decision authority, technical enforcement, operational monitoring and regulatory evidence. A concept was retained when it explained a concrete BTP risk or clarified how an organisation might test a control. This approach enables transparent reasoning without claiming an exhaustive database census.

UNDER PEER REVIEW IN IJAR

Each candidate control was examined against a common set of questions: what asset is protected; who may approve the relevant action; where is policy technically enforced; which event demonstrates compliance; and how would failure be detected or reversed? Cross-source comparison distinguished normative recommendations from empirical findings and from this article's own proposed design. The resulting operating model is a reasoned synthesis, not a validated product reference architecture. No interviews, production logs, surveys or measured breach rates were collected, and the review does not invent observations to fill that evidentiary limitation.

The attached comparator informed the organisation of background, related work, methodology, discussion and conclusion, but its survey design was not reproduced because this manuscript analyses secondary evidence only. Numeric citations are used consistently, with direct support traceable to the corresponding DOI entries. Regulatory and product statements were separately checked against their issuing bodies to avoid treating broad academic recommendations as authoritative interpretations of Saudi law or SAP functionality.

4. Governance of the SAP BTP administrative surface

Platform administration begins with the hierarchy of global account, directories where used, subaccounts, environments and service instances. This arrangement is not merely a billing convenience. It defines boundaries through which administrators create resources, assign entitlements, establish trust and activate services. Enterprises should design these boundaries around ownership, data sensitivity and operational separation before distributing authorisations. A production subaccount hosting regulated workflows requires a different approval and change regime from a development area used for short-lived experimentation. Otherwise, organisational policy is undermined by technically convenient but indiscriminate shared administration [2,3].

The first governance requirement is an account ownership register. Every subaccount should have a named business sponsor, technical custodian, intended purpose, service inventory, data classification and approved connectivity pattern. The register must be reconciled to actual resources rather than maintained solely through manual declarations. Orphaned accounts and residual subscriptions are important because inactive workloads can retain identity trusts, service bindings and network permissions. A quarterly review is useful for stable accounts, but highly privileged changes require event-triggered checking instead of waiting for a scheduled review [4,17].

UNDER PEER REVIEW IN IJAR

Role design should distinguish platform users who manage BTP resources from business users who access deployed applications. These populations have separate trust configurations and different privilege implications. Platform administrators can influence foundational controls, including role collections and service configuration, while business users should receive application-scoped rights. Combining both populations under a broad group label makes effective access difficult to review. A robust approval route therefore links each privileged role collection to an administrative purpose, authorised approver, expiry or review rule and documented separation-of-duties assessment [7,18].

Administrative change is also a software supply-chain and configuration-management problem. A defensible baseline requires version-controlled infrastructure definitions where supported, peer-reviewed changes, protected deployment credentials and safeguards against direct production edits. However, automation does not remove human accountability. If a pipeline automatically restores a misconfigured privilege, the defect becomes repeatable at speed. Security testing must cover the policy definitions themselves, repository permissions, deployment identities and emergency rollback paths, not only the application packages transferred through the pipeline [10,11,12].

An important failure mode is inconsistent configuration between seemingly similar subaccounts. Two environments may use identical application code but diverge in identity trust, role mapping, permitted destinations or security logging. This creates an assurance problem when a successful development test is assumed to represent production. A reusable configuration baseline should define mandatory settings and approved variations, while periodic drift detection compares actual state with that baseline. Exceptions need an owner and expiry date, especially when introduced to maintain operational continuity during major system upgrades.

Service instances and integration destinations deserve independent scrutiny. A destination may hold authentication configuration or define access to a remote system; a connectivity service can bridge cloud applications to on-premises resources. The security property comes from limiting what the connection can reach and which identity can invoke it, not simply from encrypting the transport. Enterprises should record the target system, permitted operations, credential owner, rotation interval and expected transaction pattern. Local security teams must be able to revoke the route without waiting for a full application redeployment when a credential is compromised [14,16].

UNDER PEER REVIEW IN IJAR

Governance also requires a controlled exception mechanism. Teams may need additional access during migrations, outages or incident containment. Time-bounded elevation with independent approval is preferable to permanent administrative accumulation. Emergency access should use separately protected credentials, explicit activation records and retrospective review; it should not depend exclusively on the identity provider whose failure is being addressed. The organisation should test recovery of administrative access before an actual federation outage, because an untested recovery procedure creates apparent resilience rather than usable resilience [6,22].

The architectural result is a layered model. Strategic owners approve scope and risk; identity services establish who or what is requesting access; BTP administration translates authority into resources and role collections; workload controls protect services and connections; and assurance functions verify evidence against applicable obligations. Figure 1 represents these relationships without suggesting that any one SAP component independently provides legal compliance.



Figure 1. Layered governance architecture linking decision authority, identity control, SAP BTP administration, workload protection and independent assurance. Original conceptual illustration.

5. Identity authentication and authorisation

Authentication establishes confidence that a claimant controls the asserted identity; authorisation determines the operations that identity may perform. Conflating the two leads to a common control failure: successful corporate single sign-on is treated as evidence that

downstream application privileges are appropriate. Federated authentication can centralise user policy, but application entitlements still require mapping, approval and review. Contemporary identity guidance distinguishes proofing, authenticator strength and federation because each addresses a separate failure mode. Their assurance levels should inform risk choices rather than be copied mechanically into every business workflow [21,23,24].

For SAP BTP, SAP Cloud Identity Services—Identity Authentication can act as a central integration point between the platform and corporate identity providers. The enterprise must govern trust relationships, assertion attributes, token lifetimes and administrative fallback arrangements. Trusting an identity source is not equivalent to trusting every claim it emits. Account matching and group attributes require explicit conventions to avoid duplicate identities, name collisions or inappropriate role propagation. When identity-provider groups are mapped to BTP role collections, the organisation should test both positive grants and negative outcomes after group removal [18,22].

Provisioning addresses the temporal dimension that authentication cannot resolve. Recruitment, role transfer, contractor expiry and termination each alter legitimate access. A joiner–mover–leaver workflow should connect authoritative employment or contract records to group membership, BTP assignments and relevant application accounts. Automating account creation without reliable removal may worsen privilege accumulation. Effective testing therefore measures not only provisioned identities but also the maximum time between a termination event and the revocation of usable access across all connected services. Delayed session expiry and cached authorisation information must be considered explicitly [7,21].

Role-based access control offers an understandable administrative foundation, yet roles proliferate when organisational tasks vary by project, country or data sensitivity. Attribute-based rules may better capture contextual conditions, but they introduce dependency on accurate attributes, reliable policy evaluation and accountable rule ownership. The appropriate solution is often hybrid: a narrow role defines the business function while additional attributes constrain access to particular resources or transactions. Neither approach is intrinsically superior without tests for privilege creep, conflict resolution and the cost of explaining a decision to auditors [8,20,28,29].

Privileged access must be evaluated separately from ordinary employee access. Platform operators may create application trusts, rotate secrets, alter logging settings and grant privileges to others. Consequently, strong authenticators, dedicated administrative identities and protected privileged workstations deserve higher priority than indiscriminate additional login prompts for

UNDER PEER REVIEW IN IJAR

low-risk users. Risk-based authentication can escalate assurance when device or session conditions change, but it must not substitute for narrowing the authorised action. Continuous evaluation of sessions is useful only when the platform can actually enforce the resulting decision [1,6,18].

Machine identities require equally careful design. APIs, automation jobs, pipelines and service-to-service connections often authenticate without a person present. Shared technical users and long-lived secrets weaken attribution and make incident containment difficult. Wherever supported, narrowly scoped service identities, managed credentials and short-lived tokens should be preferred. The enterprise must inventory each workload identity, define its owner and purpose, inspect its effective privileges and schedule credential rotation. A dormant service identity that retains production access may be a more persistent exposure than an employee account regularly reviewed by human resources [14,15].

Authorisation testing should examine policy interactions rather than isolated roles. Consider a purchasing application in which a user can submit a supplier record, approve a payment change and administer the interface that transmits bank details. Each permission might appear acceptable individually while their combination creates a material fraud path. Separation-of-duties checks therefore require a cross-system view of effective access, including inherited collections and API permissions. Security teams need an intelligible record of why a decision was granted, whose authority approved it and which dependency could change that outcome [5,9].

Figure 1 also demonstrates why technical identity controls cannot be assigned solely to a central IAM team. Human resources owns identity status, business leaders own legitimate task authority, platform administrators implement mappings, and security assurance verifies behaviour. Identity governance becomes credible only when these roles exchange timely evidence and have a defined procedure for resolving contradictory records.

6. Cloud cybersecurity compliance and operational assurance

Cloud cybersecurity compliance is most persuasive when regulatory language becomes verifiable operating behaviour. A control statement such as 'access shall be restricted' is not a complete test. A tenant should specify protected assets, permitted principal types, privileged approval rules, monitoring events, exception limits and the interval at which evidence is reviewed. This translation reduces the distance between policy interpretation and system configuration. It also exposes cases where a vendor capability is available but not enabled, or where the customer requires assurance that only the provider can produce [2,17].

UNDER PEER REVIEW IN IJAR

For Saudi enterprises, an applicability assessment should start with the organisation's legal and regulatory scope and the current NCA control editions, rather than assume that every CCC clause imposes identical obligations on each tenant. The control register can then map relevant requirements to SAP BTP resources and neighbouring systems. An enterprise might classify identity federation and account configuration as tenant activities, while physical data-centre protection and underlying infrastructure operations depend on provider assurance. Both links should be visible in an evidence package, with the relevant service boundary and contractual responsibilities stated [3,12].

Data localisation presents an architectural decision rather than a label that can be attached after deployment. Storage location, processing pathways, support access, disaster recovery, logs and interfaces may have different geographic characteristics. A data-flow inventory should therefore identify each class of regulated information, the systems receiving it, transfer mechanisms and applicable approval or restriction conditions. The PDPL's transfer framework and sector-specific rules require specialist interpretation; an assurance model should preserve that uncertainty instead of claiming that all cross-border transfers are categorically forbidden or automatically permitted. Data classification must precede region and service selection [4,27].

Logging is a core evidentiary control, but collecting events is insufficient. Administrators should identify which actions are material—trust changes, privileged role assignments, secret rotations, connector alterations, policy overrides and authentication anomalies—and ensure relevant records can be correlated. Logs need controlled access, reliable timestamps, retention rules and tested forwarding to security monitoring. An audit trail loses credibility if the same administrator can both create a high-risk entitlement and erase the only record of that change. Separation between operational configuration and independent evidence preservation is therefore important [7,19].

Assurance evidence must also be evaluated for quality. A screenshot showing that multifactor authentication is selected demonstrates a configuration point, not necessarily that every privileged identity passed a sufficiently strong authentication challenge. Evidence should identify the collection period, population covered, exceptions, review method and any unsupported assets. Sampling may be justified for lower-risk processes, but privileged identity changes warrant complete event capture. When compliance is scored, gaps in coverage and evidence freshness should be disclosed rather than hidden inside a single favourable percentage.

UNDER PEER REVIEW IN IJAR

Incident readiness should connect platform telemetry to response authority. If a service token is exposed, the organisation needs an approved sequence for revoking credentials, evaluating affected integrations, preserving evidence, restoring service and communicating with stakeholders. Runbooks should identify which tasks are available to a tenant and which require provider support. Recovery must be tested against realistic dependencies, including a loss of identity federation or an unavailable on-premises connector. Regular exercises can reveal hidden assumptions that documentation reviews are unlikely to detect [16,30].

Application security extends the assurance boundary beyond access governance. Cloud-native applications may inherit weaknesses from dependencies, containers, API gateways and software delivery pipelines. Secure development controls should cover code review, dependency provenance, secret scanning, vulnerability remediation and approved deployment routes. Runtime monitoring is valuable but cannot reliably compensate for structural exposure caused by unrestricted network paths or unsafe application permissions. The most defensible approach combines preventive build-time gates with behavioural monitoring and rapid response, assigning different owners to each lifecycle phase [10,11,13,15,16].

Third-party assurance remains difficult because documentation can report general control maturity without proving protection of a particular tenant configuration. Enterprises should request service-specific attestations, identify exclusions, understand subcontractor dependencies and assess the age of evidence. Provider certifications can support confidence but do not validate the customer's own group mappings, data classifications or integration credentials. Conversely, a tenant's technically correct controls cannot establish the provider's underlying physical or infrastructure safeguards. An evidence matrix should make both limitations explicit [12,19,27].

Measurement should distinguish activity from outcomes. Counting the number of policies written, users enrolled or security alerts generated may reward administrative volume without reducing exposure. More informative measures include the proportion of privileged accounts with current approval, time to revoke terminated identities, percentage of high-risk changes with independent review, proportion of critical connections with named owners and elapsed time to close audit findings. These measures are proposed governance indicators; they are not statistically validated predictors of breach frequency. They should be interpreted alongside incident severity, service criticality and business disruption [4,17,26].

UNDER PEER REVIEW IN IJAR

7. Integrated governance framework and implementation pathway

The proposed framework treats every high-risk cloud decision as a traceable chain from requirement to evidence. At the policy layer, the organisation identifies business purpose, regulatory scope and risk tolerance. At the identity layer, authoritative sources establish the subject and permitted function. At the platform layer, BTP account structures, trust settings and role collections implement those decisions. At the workload layer, application permissions, network destinations and service identities further constrain behaviour. At the assurance layer, independent review tests whether observed configurations remain aligned with the approved baseline. A failure at any link breaks the chain of demonstrable control [1,7,17].

Implementation should begin with discovery rather than immediate procurement of another security product. The enterprise first inventories global accounts, subaccounts, privileged principals, identity providers, service instances and data flows. It then establishes ownership and assesses the applicability of NCA, privacy and sector requirements. The next phase removes unjustified access, standardises trust and role mapping, introduces protected changes and defines monitoring events. Only after these fundamentals are stable should advanced risk scoring or automated policy enforcement be expanded. This ordering reduces the danger of automating an undocumented privilege model [5,25].

Table 1 allocates accountability across the main control surfaces; Table 2 converts representative risks into practical tests and evidence. The proposed allocation is illustrative, not a substitute for an enterprise's formal responsibility matrix. In a regulated financial entity, for instance, sector oversight may add approval and reporting requirements beyond the generic cloud baseline. In a smaller organisation, one employee may perform multiple technical tasks, but independent authorisation and review of high-risk actions remain essential. The model should scale control independence rather than simply reproduce the organisational chart [3,4].

The framework's value is clearest when compared with two alternatives. A platform-only programme may standardise subaccounts yet fail to revoke users whose employment has ended; an identity-only programme may govern human accounts while leaving connector secrets and service permissions unmanaged. A compliance-only programme may produce policies but lack logs demonstrating enforcement. The integrated model closes these disconnects by requiring each important business authorisation to have an location, owner, revocation route and independent evidence trail.

Figure 2 depicts a repeating assurance cycle. The process begins with a scoped requirement, translates it into enforceable policy, observes live operation, compares evidence with expected

UNDER PEER REVIEW IN IJAR

behaviour and triggers corrective action. Feedback should update both configuration and the risk register. A rejected access request can demonstrate technical enforcement, whereas a permissible but unauthorised entitlement may require comparison with human approval records. Effective control assessment therefore combines machine-collected signals with business context and protects the integrity of both evidence sources [7,25].

Table 1. Proposed allocation of decision rights and assurance responsibilities across SAP BTP governance surfaces.

Control surface	**Primary accountable function**	**Evidence / critical check**
Global account and subaccounts	Platform owner + business sponsor	Owned inventory; approved boundaries; entitlement changes
Identity federation and provisioning	Enterprise IAM + HR	Trust configuration; group lifecycle; revocation testing
Role collections and privilege	Business approver + BTP administrator	Least-privilege mapping; privileged exception log
Connections, APIs and service keys	Integration or application custodian	Destination approval; secret lifecycle; target restrictions
Data protection and compliance	Privacy, legal and security leads	Classification; transfer assessment; control applicability
Control assurance and response	Independent security monitoring	Immutable evidence path; review and recovery exercises

Table 2. Illustrative risk-to-control-to-evidence matrix for ongoing review; tests must be tailored to actual system scope.

Risk scenario	**Required control behaviour**	**Testable assurance evidence**
Former user retains a session	Revoke identity, assignments and usable sessions promptly	Termination-to-revocation test; token/session checks
Unapproved privileged collection	Independent approval and time-bounded grant	Entitlement diff; approver record; exception expiry
Connector accesses	Restrict back-end target	Destination inventory; permissions;

excessive scope	and permitted operations	connection test
Service credential is exposed	Constrain scope; rotate and revoke safely	Secret owner; rotation event; response exercise
Sensitive data crosses a boundary	Classify flows and assess applicable transfer duties	Data-flow map; lawful basis and transfer assessment
Cloud configuration drifts	Detect divergence and enforce remediation ownership	Baseline comparison; findings register; retest



Figure 2. Closed-loop assurance cycle from requirement scoping and authorisation to enforcement, observation and corrective verification. Original conceptual illustration.

8. Critical discussion, limitations and research agenda

The principal contribution is a reconciliation of three perspectives often discussed separately: platform administration, identity-centric security and statutory accountability. Existing studies explain zero-trust principles and threats to cloud-native workloads, while standards enumerate controls and assessment procedures [5,8,15]. These bodies of knowledge do not by themselves determine how a Saudi enterprise should govern a BTP role collection or assign responsibility for a particular connector. The review's framework addresses that intermediate organisational design problem without presenting vendor-specific configuration as a universal security solution.

Several tensions remain. Centralising federation improves consistency but increases dependency on a shared identity service; granular entitlement design limits exposure but can generate administrative complexity; comprehensive logging improves forensic visibility but

UNDER PEER REVIEW

increases retention, privacy and cost obligations. Automated compliance checks reduce manual effort yet may overlook the substantive business legitimacy of granted access. Consequently, governance maturity should be assessed through the quality of decisions and exceptions rather than the number of deployed security features [18,19,25].

This review is limited by its selective, concept-oriented evidence base and absence of primary organisational data. It does not verify a live SAP BTP tenant, estimate return on investment or determine the legal applicability of individual Saudi controls to a named company. Future studies should compare anonymised configurations across sectors, measure entitlement-revocation delay, test emergency-access recovery and examine whether evidence-based governance reduces unresolved high-risk exceptions. Longitudinal designs could evaluate how policy drift changes during integration programmes and whether automated controls improve assurance without constraining legitimate operations [26,27].

9. Conclusion

Secure SAP BTP governance under Vision 2030 requires more than platform availability or successful single sign-on. Enterprises need clear account ownership, robust trust arrangements, least-privilege entitlements, governed service identities, protected integration routes and evidence that applicable cybersecurity and data-protection requirements operate in practice. This review proposes a layered, auditable relationship between business authority, technical enforcement and independent assurance. It identifies privilege persistence, unclear cloud responsibilities and weak evidence lineage as especially consequential failure modes. Adoption should proceed through discovery, ownership assignment, control testing and disciplined remediation, with legal applicability determined case by case. The framework provides a basis for implementation and future empirical evaluation rather than a claim of demonstrated organisational effectiveness.

References

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

[2] Joint Task Force (2020). Security and Privacy Controls for Information Systems and Organizations. NIST Special Publication 800-53, Revision 5. https://doi.org/10.6028/NIST.SP.800-53r5

[3] Joint Task Force (2020). Control Baselines for Information Systems and Organizations. NIST Special Publication 800-53B. https://doi.org/10.6028/NIST.SP.800-53B

UNDER PEER REVIEW IN IJAR

[4] Stine, K.; Quinn, S.; Witte, G.; Gardner, R. (2020). Integrating Cybersecurity and Enterprise Risk Management (ERM). NIST Interagency Report 8286. https://doi.org/10.6028/NIST.IR.8286

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

[6] Teerakanok, S.; Uehara, T.; Inomata, A. (2021). Migrating to Zero Trust Architecture: Reviews and Challenges. Security and Communication Networks, 2021, 9947347. https://doi.org/10.1155/2021/9947347

[7] Joint Task Force (2022). Assessing Security and Privacy Controls in Information Systems and Organizations. NIST Special Publication 800-53A, Revision 5. https://doi.org/10.6028/NIST.SP.800-53Ar5

[8] Syed, N. F.; Shah, S. W.; Shaghaghi, A.; Anwar, A.; Baig, Z. A.; Doss, R. (2022). Zero Trust Architecture (ZTA): A Comprehensive Survey. IEEE Access, 10, 57143–57179. https://doi.org/10.1109/ACCESS.2022.3174679

[9] He, Y.; Huang, D.; Chen, L.; Ni, Y.; Ma, X. (2022). A Survey on Zero Trust Architecture: Challenges and Future Trends. Wireless Communications and Mobile Computing, 2022, 6476274. https://doi.org/10.1155/2022/6476274

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

[11] Chandramouli, R. (2022). Implementation of DevSecOps for a Microservices-Based Application with Service Mesh. NIST Special Publication 800-204C. https://doi.org/10.6028/NIST.SP.800-204C

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

[13] Belal, M. M.; Sundaram, D. M. (2022). Comprehensive review on intelligent security defences in cloud: Taxonomy, security issues, ML/DL techniques, challenges and future trends. Journal of King Saud University – Computer and Information Sciences, 34, 9102–9131. https://doi.org/10.1016/j.jksuci.2022.08.035

UNDER PEER REVIEW IN IJAR

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

[15] Minna, F.; Massacci, F. (2023). SoK: Run-time security for cloud microservices. Are we there yet?. Computers & Security, 127, 103119. https://doi.org/10.1016/j.cose.2023.103119

[16] Zeng, Q.; Kavousi, M.; Luo, Y.; Jin, L.; Chen, Y. (2023). Full-stack vulnerability analysis of the cloud-native platform. Computers & Security, 129, 103173. https://doi.org/10.1016/j.cose.2023.103173

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

[18] Aboukadri, S.; Ouaddah, A.; Mezrioui, A. (2024). Machine learning in identity and access management systems: Survey and deep dive. Computers & Security, 139, 103729. https://doi.org/10.1016/j.cose.2024.103729

[19] Rizvi, S.; Williams, I. (2024). Analyzing transparency and malicious insiders prevention for cloud computing environment. Computers & Security, 137, 103622. https://doi.org/10.1016/j.cose.2023.103622

[20] Mei, Z.; Yu, J.; Zhang, C.; Wu, B.; Yao, S.; Shi, J.; Wu, Z. (2024). Secure multi-dimensional data retrieval with access control and range query in the cloud. Information Systems, 122, 102343. https://doi.org/10.1016/j.is.2024.102343

[21] Temoshok, D.; Fenton, J.; Choong, Y.-Y.; Lefkovitz, N.; Regenscheid, A.; Galluzzo, R.; Richer, J. (2025). Digital Identity Guidelines: Authentication and Authenticator Management. NIST Special Publication 800-63B-4. https://doi.org/10.6028/NIST.SP.800-63b-4

[22] Temoshok, D.; Richer, J.; Choong, Y.-Y.; Fenton, J.; Lefkovitz, N.; Regenscheid, A.; Galluzzo, R. (2025). Digital Identity Guidelines: Federation and Assertions. NIST Special Publication 800-63C-4. https://doi.org/10.6028/NIST.SP.800-63c-4

[23] Temoshok, D.; Abruzzi, C.; Choong, Y.-Y.; Fenton, J.; Galluzzo, R.; LaSalle, C.; Lefkovitz, N.; Regenscheid, A.; Vachino, M. (2025). Digital Identity Guidelines: Identity Proofing and Enrollment. NIST Special Publication 800-63A-4. https://doi.org/10.6028/NIST.SP.800-63a-4

UNDER PEER REVIEW IN IJAR

[24] Temoshok, D.; Proud-Madruga, D.; Choong, Y.-Y.; Galluzzo, R.; Gupta, S.; LaSalle, C.; Lefkovitz, N.; Regenscheid, A. (2025). Digital Identity Guidelines. NIST Special Publication 800-63-4. https://doi.org/10.6028/NIST.SP.800-63-4

[25] Chuprikov, P.; Eugster, P.; Mangipudi, S. (2025). Security Policy as Code. IEEE Security & Privacy, 23(2), 23–31. https://doi.org/10.1109/MSEC.2025.3535803

[26] Liang, X.; Xu, Y. (2025). A novel framework to identify cybersecurity challenges and opportunities for organizational digital transformation in the cloud. Computers & Security, 151, 104339. https://doi.org/10.1016/j.cose.2025.104339

[27] Ali, S.; Talpur, D. B.; Abro, A.; Alshudukhi, K. S.; Alwakid, G. N.; Humayun, M.; Bashir, F.; Wadho, S. A.; Shah, A. (2025). Security and privacy in multi-cloud and hybrid cloud environments: Challenges, strategies, and future directions. Computers & Security, 157, 104599. https://doi.org/10.1016/j.cose.2025.104599

[28] Walid, R.; Joshi, K. P.; Choi, S. (2024). Leveraging semantic context to establish access controls for secure cloud-based electronic health records. International Journal of Information Management Data Insights, 4(1), 100211. https://doi.org/10.1016/j.jjimei.2023.100211

[29] Trabelsi, R.; Fersi, G.; Jmaiel, M. (2023). Access control in Internet of Things: A survey. Computers & Security, 135, 103472. https://doi.org/10.1016/j.cose.2023.103472

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

UNDER PEER REVIEW IN IJAR