

From Cloud Security Alerts to Audit-Ready Evidence: A Risk-Based GRC Framework for Cyber Assurance in Australia

Parisasadat Shojaei*, Rezza Moieni

Cultural Infusion, Melbourne, Australia

Email: *Parisa.Shojaei@culturalinfusion.com

How to cite this paper: Shojaei, P., & Moieni, R. (2026). From Cloud Security Alerts to Audit-Ready Evidence: A Risk-Based GRC Framework for Cyber Assurance in Australia. *Open Journal of Social Sciences*, 14, 139-156.

<https://doi.org/10.4236/jss.2026.148009>

Received: June 7, 2026

Accepted: August 9, 2026

Published: August 12, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

Cloud security posture management (CSPM), vulnerability management, and cloud-native security services generate large volumes of findings about misconfiguration, vulnerable workloads, weak identity controls, public exposure, encryption gaps, and logging deficiencies. These findings are valuable for security operations, but they often remain difficult for governance, risk, compliance (GRC), audit, and management stakeholders to interpret, prioritise, and evidence. This paper proposes Cloud Alert-to-Governance Evidence (CAGE), a practical risk-based framework for converting cloud security findings into governance-ready artefacts. CAGE classifies each finding through five linked dimensions: technical severity, asset criticality, exposure context, business impact, and compliance relevance. It then connects prioritised risks to control mappings, treatment decisions, accountable owners, remediation evidence, exception records, residual risk status, and audit-ready reporting. The framework is developed through conceptual design and standards synthesis, drawing on cloud security literature, continuous monitoring guidance, risk assessment concepts, Australian cyber guidance, and recognised frameworks including the Australian Essential Eight, ISO/IEC 27001, ISO/IEC 27005, NIST SP 800-30, NIST SP 800-137, NIST Cybersecurity Framework 2.0, and the Cloud Security Alliance Cloud Controls Matrix. The paper strengthens the framework by adding a comparative validation table, an illustrative worked example, and an Australian workforce relevance section. The contribution is a structured bridge between cloud security operations and cyber assurance practice. The framework is conceptual and practice-oriented; it does not claim empirical validation. Future work should evaluate CAGE using real organisational cloud findings and practitioner review in Australian industry contexts.

Keywords

Cloud Security Posture Management, Governance, Risk and Compliance, Cyber Assurance, Audit Evidence, Cloud Security, Risk-Based Prioritisation, Australia, Continuous Monitoring

1. Introduction

Organisations increasingly depend on cloud infrastructure to deliver digital services, store sensitive data, and support distributed work. The operational characteristics that make cloud computing valuable—elasticity, rapid provisioning, managed services, automation, and distributed responsibility—also create security governance challenges. Cloud resources can be created, changed, exposed, or decommissioned quickly, while governance and assurance activities often rely on periodic reviews, manual evidence collection, spreadsheets, and retrospective audit preparation. This creates a timing and translation gap: the technical state of cloud security can change faster than governance artefacts can explain it. Prior research on health information systems and mHealth also shows that privacy, security, usability, trust, and user behaviour are interdependent concerns in digital services that handle sensitive information (Shojaei et al., 2024, 2025a, 2025b).

Cloud-native security services and CSPM platforms respond to this problem by generating findings about control failures, configuration drift, and risky exposure. AWS Security Hub CSPM, for example, runs checks against security controls and generates control findings, while Amazon Inspector automatically discovers supported workloads and continually scans for software vulnerabilities and unintended network exposure (Amazon Web Services, n.d.-a, n.d.-c). These tools provide useful technical visibility, but they do not automatically answer governance questions: Which findings are most important to the business? Which controls or obligations are affected? Who owns treatment? What evidence shows that the risk was remediated, mitigated, accepted, or formally excepted? How should the same finding be explained to engineers, risk managers, auditors, and executives?

This paper addresses that translation problem. It proposes a practical risk-based framework for converting cloud security findings into governance evidence, remediation actions, exception records, and compliance-ready reporting. The central argument is that a cloud alert should not be treated only as a technical event. It should be treated as the starting point of a traceable governance object that links the observed condition, affected asset, risk context, relevant controls, treatment decision, remediation record, validation evidence, and residual risk status.

The paper is positioned for Australian industry contexts, where organisations commonly need to communicate cyber security posture across technology, risk, audit, management, customer assurance, and third-party review functions. It uses the Australian Essential Eight as a practical mitigation baseline and aligns with broader frameworks including ISO/IEC 27001, ISO/IEC 27005, NIST SP 800-30,

NIST SP 800-137, NIST Cybersecurity Framework (CSF) 2.0, and the Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM). The framework remains cloud-provider neutral at the governance layer, although AWS Security Hub CSPM and Amazon Inspector are used as concrete examples because their finding structures illustrate how technical alert data can be enriched and mapped into evidence artefacts (Cloud Security Alliance, 2026).

The research objective is to develop a practical conceptual framework that helps cloud security and GRC teams transform technical findings into risk-based governance evidence. The paper addresses three research questions:

RQ1: What information is required to convert a cloud security finding into a governance-relevant risk object?

RQ2: How can cloud risks be mapped to control frameworks and treatment evidence without losing technical traceability?

RQ3: What governance artefacts are needed to make remediation status and residual risk understandable to technical and non-technical stakeholders?

2. Background and Related Work

2.1. Cloud Security Posture and Governance Translation

Cloud computing is commonly defined by on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service (Mell & Grance, 2011). These characteristics allow organisations to scale quickly, but they also make security state dynamic. Earlier cloud security research identified persistent issues including multi-tenancy, loss of direct control, trust, data protection, identity management, insecure interfaces, and the difficulty of demonstrating compliance when services are distributed across cloud providers and customers (Hashizume et al., 2013; Subashini & Kavitha, 2011; Takabi et al., 2010). Beyond security-specific studies, earlier cloud computing research has examined task scheduling and resource optimisation, which highlights the operational complexity of allocating, managing, and evidencing cloud workloads (Shojaei, 2016a, 2016b, 2016c, 2018; Shojaei & Naji, 2016).

Modern cloud security tools help detect these issues at scale, but alert generation is only the first step. Security teams must still triage findings, relate them to business context, prioritise remediation, and produce evidence that an identified risk was addressed. A high-severity technical issue on a non-production asset may be less urgent than a medium-severity issue on an internet-facing production system storing sensitive information. Similarly, a finding may be technically simple but governance-critical if it affects identity, encryption, logging, vulnerability management, customer assurance, or regulatory reporting.

2.2. Continuous Monitoring and Risk-Based Decision-Making

NIST SP 800-137 frames information security continuous monitoring as maintaining ongoing awareness of information security, vulnerabilities, and threats to support organisational risk management decisions (Dempsey et al., 2011). This is

directly relevant to cloud environments because configuration drift, changing network exposure, and new vulnerabilities may emerge between audit cycles. Continuous monitoring becomes valuable when it supports decisions rather than simply increasing alert volume.

Risk assessment guidance also supports the proposed approach. NIST SP 800-30 describes risk as a function of the likelihood of a threat event and the potential adverse impact should the event occur (Joint Task Force Transformation Initiative, 2012). ISO/IEC 27005 provides guidance for managing information security risks in support of an information security management system (International Organization for Standardization, 2022c). These sources justify a move from severity-only alert handling toward contextual prioritisation that considers asset value, exposure, business impact, likelihood, consequences, and control obligations. Related empirical work on distributed ledger technology and data privacy further supports the view that privacy risk cannot be assessed only as a technical condition; it also involves governance expectations, regulatory interpretation, and stakeholder trust (Shojaei & Moieni, 2025).

2.3. Control Frameworks and Audit Evidence

Control frameworks provide a common vocabulary for governance communication. ISO/IEC 27001 defines requirements for an information security management system (ISMS), and ISO/IEC 27002 provides guidance on information security controls (International Organization for Standardization, 2022a, 2022b). NIST CSF 2.0 provides a flexible structure for governing, identifying, protecting, detecting, responding to, and recovering from cybersecurity risks; its GOVERN function emphasises that cybersecurity risk management strategy, expectations, and policy are established, communicated, and monitored (National Institute of Standards and Technology, 2024). In Australia, the Essential Eight maturity model provides practical mitigation strategies and explicitly recommends risk-based implementation, minimising exceptions, documenting and approving exceptions, and regularly reviewing exceptions and compensating controls (Australian Signals Directorate, 2023). This is particularly important where AI-assisted analysis is used to classify, summarise, or report security evidence, because anonymity and identity protection remain governance concerns in AI-enabled digital environments (Shojaei et al., 2025c).

The practical challenge is that these frameworks usually operate at a management-system, control, or assurance level, while cloud security tools generate findings at a resource, vulnerability, account, image, policy, or configuration level. CAGE addresses this mismatch by specifying an intermediate evidence layer: each finding is normalised, enriched with risk context, mapped to control relevance, assigned to treatment, and stored with evidence artefacts that can support audit review, management reporting, and customer assurance responses.

3. Research Approach

This study uses a conceptual design and standards-synthesis approach. It does not

present a systematic literature review, statistical evaluation, or empirical case study. Instead, it develops a practice-oriented framework by synthesising four source streams: 1) cloud security literature that identifies recurring cloud risks and governance challenges; 2) cloud security tool documentation that shows how findings are produced and structured; 3) continuous monitoring and risk assessment guidance that explains how technical signals should support risk decisions; and 4) control frameworks that provide governance, audit, and assurance vocabulary. The design is also informed by the authors' prior privacy research, which emphasises that technical safeguards should be translated into terms that users, organisations, and assurance stakeholders can understand (Shojaei et al., 2024, 2025a, 2025b).

The design goal is practical utility. The framework must be specific enough to guide the handling of real cloud findings, but general enough to apply across providers and tools. AWS Security Hub CSPM and Amazon Inspector are used as illustrative inputs because they represent common categories of CSPM and vulnerability findings and because Security Hub findings can be represented using the AWS Security Finding Format (ASFF) (Amazon Web Services, n.d.-b). However, the governance layer can be applied to other sources such as Microsoft Defender for Cloud, Google Security Command Center, Wiz, Prisma Cloud, Lacework, Tenable, Qualys, or SIEM-integrated cloud detections, provided the findings contain or can be enriched with severity, affected asset, exposure, owner, and control mapping information.

The synthesis followed four steps. First, recurring cloud risk categories were identified from cloud security literature and tool outputs. Second, governance-relevant dimensions were derived from risk assessment and continuous monitoring concepts. Third, cloud risks were mapped to control areas in Essential Eight, ISO/IEC 27001, ISO/IEC 27005, NIST CSF, and CSA CCM. Fourth, the outputs were consolidated into a pipeline that produces governance artefacts: risk records, treatment actions, exception decisions, evidence packs, and reporting views.

4. The CAGE Framework

The proposed Cloud Alert-to-Governance Evidence (CAGE) framework transforms raw cloud findings into governance-ready evidence through five stages: finding normalisation, risk context enrichment, control and compliance mapping, treatment and ownership, and evidence packaging. **Figure 1** summarises the pipeline.




Figure 1. Cloud Alert-to-Governance Evidence (CAGE) pipeline.

4.1. Stage 1: Cloud Finding Normalisation

The first stage converts tool-specific findings into a common structure. A finding should include, at minimum, a finding identifier, source tool, affected resource, account or subscription, region, finding type, severity, first observed time, last observed time, current status, and a technical description. Where possible, the finding should also include resource tags, cloud service, vulnerability identifiers, exposed ports or protocols, network reachability context, identity principal context, and links to the source evidence.

Normalisation is essential because different tools describe risk differently. A vulnerability scanner may emphasise CVE severity and package versions; a CSPM platform may emphasise configuration state; a SIEM rule may emphasise suspicious activity; and a cloud-native service may use its own finding schema. Without normalisation, governance teams cannot compare findings consistently or aggregate them into meaningful risk and control reporting.

4.2. Stage 2: Risk Context Enrichment

The second stage adds business and exposure context. The framework evaluates each finding using five dimensions: technical severity, asset criticality, exposure context, business impact, and compliance relevance. These dimensions prevent over-reliance on technical severity alone. **Table 1** defines each dimension.

Table 1. Risk context dimensions used to enrich cloud findings.

Dimension	Meaning	Example inputs	Governance use
Technical severity	How serious the tool, vendor, or vulnerability source considers the condition.	Critical/high/medium/low severity; CVSS; exploitability; vendor score.	Initial triage and service-level expectation selection.
Asset criticality	How important the affected asset is to business operations or data protection.	Production status; data classification; customer impact; service tier; owner tags.	Prioritises business-critical systems over low-value assets.
Exposure context	Whether the condition is reachable, externally exposed, exploitable, or limited by compensating controls.	Internet exposure; network path; public bucket status; security group rules; IAM trust boundary.	Distinguishes theoretical weakness from reachable risk.
Business impact	Potential operational, financial, legal, privacy, safety, or reputational consequences.	Sensitive data; regulated service; outage consequence; customer-facing dependency.	Supports management-level risk communication.
Compliance relevance	Connection to a control, policy requirement, assurance criterion, or reporting obligation.	Essential Eight strategy; ISO/IEC 27001 control area; NIST CSF category; CSA CCM domain; internal policy.	Links findings to evidence, assurance, and audit reporting.

4.3. Stage 3: Control and Compliance Mapping

The third stage maps enriched findings to control areas. The purpose is not to claim automatic compliance with an entire standard. Instead, the goal is to identify which control objective, mitigation strategy, governance concern, or assurance criterion is implicated by the finding. Mapping should be treated as a traceability aid, not as a replacement for professional judgement or formal audit.

For example, a public storage bucket containing sensitive records may be mapped to access control, data protection, encryption, logging, and monitoring controls. A vulnerable public workload may be mapped to vulnerability management, patching, exposure management, and incident response readiness. A missing log configuration may be mapped to detection, monitoring, audit logging, and evidence retention. **Table 2** provides illustrative mappings. These examples are deliberately high-level because precise control references depend on organisational scope, chosen standard, implementation context, and auditor judgement.

Table 2. Illustrative mapping from cloud findings to governance evidence.

Cloud risk category	Typical finding	Relevant control area	Evidence artefact	Treatment example
Public exposure	Storage bucket, database, admin port, or workload exposed to the internet.	Access control; network security; secure configuration; NIST CSF Protect/Detect; CSA CCM infrastructure and network controls.	Resource configuration snapshot; exposure path; owner approval; remediation timestamp.	Restrict public access, implement least privilege, document exception if approved.
Weak access control	Overly permissive IAM policy, unused privileged access, or missing MFA for administrators.	Identity and access management; privileged access; Essential Eight MFA; ISO access control.	Policy diff; access review record; approval evidence; MFA status report.	Remove excessive permissions, enable MFA, schedule periodic access review.
Missing encryption	Storage volume, database, backup, or transit path not encrypted where required.	Cryptography; data protection; privacy and security control expectations.	Encryption setting; data classification; change ticket; validation screenshot or API output.	Enable encryption, rotate keys if required, document compensating control.
Vulnerability	Compute instance, container image, or serverless package with a high-risk CVE.	Vulnerability management; patching; secure operations; Essential Eight patching.	CVE details; asset criticality; patch ticket; retest result; closure evidence.	Patch, rebuild image, apply mitigation, or accept residual risk with expiry.
Logging gap	Audit logging, flow logs, object access logs, or threat detection not enabled.	Logging and monitoring; detection; auditability; NIST CSF Detect.	Logging configuration; retention policy; SIEM ingestion proof; exception record.	Enable logging, verify ingestion, define retention and alert owner.
Misconfiguration	Security control disabled, insecure default, or non-compliant parameter.	Secure configuration; change control; continuous monitoring.	Before/after configuration; control check result; change record.	Correct configuration and add preventive policy or guardrail.

4.4. Stage 4: Treatment and Ownership

The fourth stage converts prioritised risk objects into accountable treatment work. Each finding should receive an owner, treatment decision, target date, status, and

evidence requirement. Treatment decisions may include remediate, mitigate, transfer, accept, defer, or classify as false positive. The treatment path must be recorded because audit readiness depends not only on whether an issue was closed, but on whether closure was justified and evidenced.

Ownership should be tied to the affected system, platform team, application team, or risk owner. For industry practice, the framework recommends integrating treatment actions with existing workflow systems such as Jira, ServiceNow, GitLab Issues, Azure DevOps, or similar tools. This improves traceability between the cloud finding, engineering work, change approval, remediation implementation, and validation evidence.

4.5. Stage 5: Governance Evidence Pack

The final stage packages evidence into a reusable governance artefact. A CAGE evidence pack should be understandable to auditors and management while preserving enough technical detail for engineers to verify the state. **Table 3** specifies the minimum evidence pack fields. A suggested implementation structure for capturing and maintaining these evidence fields is provided in the **Appendix**.

Table 3. Minimum fields for a CAGE governance evidence pack.

Evidence pack field	Purpose	Example
Finding identity	Preserves traceability to source alert.	Security Hub finding ID, Inspector finding ARN, tool URL.
Affected asset	Identifies system, owner, and business context.	Account, region, resource ARN, application, data classification.
Risk statement	Converts technical issue into governance language.	Internet-exposed database may allow unauthorised access to customer records.
Control mapping	Links issue to framework, policy, assurance, or audit criterion.	ISO access control/logging; NIST CSF Protect/Detect; Essential Eight MFA or patching; CSA CCM domain.
Treatment record	Shows decision and accountability.	Remediate by date; owner; ticket; approval; exception expiry.
Validation evidence	Shows whether the issue was actually addressed.	After-state configuration, successful control check, vulnerability retest, log ingestion proof.
Residual risk decision	Documents acceptance, exception, or remaining exposure.	Risk accepted by owner until expiry date with compensating control.
Reporting status	Supports dashboards and audit packs.	Open, overdue, remediated, exception approved, false positive.

5. Risk Prioritisation Model

CAGE uses a semi-quantitative prioritisation model. The purpose is not to pro-

duce mathematically perfect risk scores; rather, it provides a consistent, explainable method for ranking findings and defending treatment decisions. The suggested model scores each dimension from 1 to 5 and calculates a contextual priority value. Organisations can adjust weights according to their risk appetite, sector, and regulatory environment.

A simple starting formula is:

$$\text{Priority Score} = (\text{Technical Severity} \times 0.30) + (\text{Asset Criticality} \times 0.25) + (\text{Exposure Context} \times 0.20) + (\text{Business Impact} \times 0.15) + (\text{Compliance Relevance} \times 0.10)$$

The weights intentionally give substantial value to technical severity while preventing severity from dominating the decision. For regulated or sensitive-data environments, organisations may increase the weight for business impact or compliance relevance. For internet-facing systems, exposure context may be weighted more heavily. The model should be calibrated through review by security, platform, risk, and business stakeholders (Table 4).

Table 4. Example priority interpretation for CAGE scores.

Priority score range	Suggested priority	Example treatment expectation
4.20 - 5.00	Critical	Immediate escalation; short remediation target; risk owner visibility; executive reporting if business impact is material.
3.40 - 4.19	High	Defined remediation window; accountable owner and ticket; validation evidence required.
2.40 - 3.39	Medium	Planned remediation cycle; grouped fixes where appropriate; monitor trends and overdue items.
1.00 - 2.39	Low	Track for hygiene improvement; accept, defer, or remediate based on cost, risk, and policy.

Accepted or deferred items should include explicit justification, compensating controls, approval, and an expiry date. This is consistent with risk-based implementation and exception handling guidance in the Essential Eight maturity model (Australian Signals Directorate, 2023).

6. Illustrative Worked Example

This section provides a synthetic example to demonstrate how CAGE can be applied. The example is not based on real organisation and is included only to show the logic of the framework (Table 5).

The worked example illustrates the central value of CAGE: the finding does not disappear into a technical dashboard, and it is not manually reinterpreted at audit time. It becomes a traceable object with a risk statement, control context, accountable treatment, validation evidence, and residual risk status.

Table 5. Synthetic example of CAGE applied to a cloud vulnerability finding.

Step	CAGE output for synthetic finding
Raw technical finding	Amazon Inspector reports a high-severity vulnerability on an EC2 instance. Security Hub shows the instance is in a production account and is associated with a security group allowing inbound HTTPS from the internet.
Normalised finding	Finding ID, account, region, instance ID, CVE identifier, severity, first observed date, current status, public exposure indicator, resource tags, application owner, and source URL are recorded.
Risk context enrichment	Technical severity = 4; asset criticality = 5 because the asset supports a customer-facing production service; exposure context = 4 because the service is internet-facing; business impact = 4 because compromise could affect customer trust and service availability; compliance relevance = 3 because vulnerability management and secure operations controls are implicated.
Priority score	$(4 \times 0.30) + (5 \times 0.25) + (4 \times 0.20) + (4 \times 0.15) + (3 \times 0.10) = 4.15$. The item is treated as High and close to Critical depending on organisational risk appetite.
Control mapping	Mapped to vulnerability management, patching, secure configuration, exposure management, incident readiness, and NIST CSF Identify/Protect/Detect categories. If the organisation uses Essential Eight, the item is linked to patching applications or operating systems, depending on the vulnerable component.
Treatment and ownership	Application owner and platform owner are assigned. A remediation ticket is created to patch or rebuild the instance or image. Target date is set according to the organisation's high-priority vulnerability SLA.
Evidence pack	Before-state finding, affected asset details, risk statement, change ticket, patch or rebuild evidence, after-state vulnerability scan, validation timestamp, owner approval, and residual risk decision are retained.
Management reporting	The issue appears in engineering as a patch ticket, in GRC as a high-priority vulnerability risk, and in executive reporting only if overdue, repeated across critical services, or accepted as residual risk.

7. Comparative Validation against Existing Approaches

Because this paper is conceptual, the validation presented here is a structured comparative assessment rather than empirical proof. The purpose is to show whether CAGE addresses practical gaps that are commonly present when organisations rely only on CSPM dashboards, manual risk registers, control-framework mapping, or CSA CCM control assessment. **Table 6** compares these approaches using design criteria derived from the research questions: technical traceability, business risk context, control mapping, ownership, remediation evidence, exception handling, reporting usability, and cloud-specific operational fit.

The comparison shows that CAGE is not intended to replace CSPM tooling, risk registers, ISO/IEC 27001, or CSA CCM. Instead, it acts as an operational translation layer between them. Its contribution is strongest where organisations

already have technical findings and control frameworks but lack a repeatable mechanism for producing risk-based, evidence-ready governance artefacts.

Table 6. Comparative validation of CAGE against existing approaches.

Approach	Strengths	Limitations when used alone	CAGE improvement
CSPM-only dashboards	Strong technical visibility; near-real-time or continuous checks; useful for engineers and security analysts; can show failed controls and affected resources.	Often prioritises by tool severity rather than business impact; may not capture risk owner, treatment decision, exception approval, residual risk, or audit evidence in a governance-ready format.	Preserves technical traceability while adding asset criticality, exposure, business impact, control relevance, ownership, validation evidence, and reporting status.
Manual risk registers	Familiar to risk and audit teams; useful for recording risk owners, treatments, ratings, and acceptance decisions.	Can become detached from live cloud findings; may rely on manual updates; often lacks resource-level evidence and source-tool traceability.	Connects register-style risk decisions to live or recent cloud findings, source evidence, remediation tickets, and validation results.
ISO/IEC 27001 control mapping	Provides recognised ISMS structure and audit vocabulary; supports policy, risk treatment, governance, and continual improvement.	Control-level mapping alone does not show which specific cloud resource failed, whether exposure is reachable, who remediated it, or what after-state evidence exists.	Uses ISO mapping as one evidence dimension while retaining finding identity, asset context, treatment status, and validation evidence.
CSA Cloud Controls Matrix	Cloud-specific control framework; useful for cloud assurance, supplier assessment, and mapping to cloud security domains.	A control matrix can identify relevant control domains, but it does not by itself prioritise live findings, assign owners, or package remediation evidence.	Uses CSA CCM as a cloud-control vocabulary while adding risk scoring, workflow integration, exception handling, and evidence packaging.
CAGE framework	Bridges technical alerts and governance artefacts; supports prioritisation, ownership, audit evidence, exceptions, and role-specific reporting.	Requires asset inventory quality, tagging discipline, maintained control mappings, and organisational agreement on scoring weights. It is not an automated certification method.	Provides a practical intermediate layer that can be implemented through spreadsheets, tickets, GRC tooling, or automation pipelines.

8. Governance and Reporting Outputs

The framework produces three reporting views. The first is an operational remediation view for engineers and security analysts. It lists affected resources, owners, technical details, recommended actions, due dates, and validation steps. The second is a risk and compliance view for GRC stakeholders. It groups findings by

control area, risk category, treatment status, overdue actions, exception expiry, and residual risk. The third is an executive posture view. It summarises trends, material risks, critical overdue items, risk acceptance decisions, and evidence readiness.

These views help avoid a common communication failure: technical teams may report hundreds of findings while management needs a smaller number of decision-ready risk messages. CAGE preserves technical depth but changes the reporting unit from alert count to risk and evidence object. This enables better questions: which business services are exposed, which control areas are repeatedly failing, which teams have unresolved treatment actions, where exceptions are accumulating, and which evidence packs are ready for assurance review? This translation is consistent with prior mHealth privacy research showing that adoption and trust are shaped not only by the existence of safeguards, but by whether privacy expectations, usability needs, and acceptable trade-offs are clearly communicated (Shojaei et al., 2025a, 2025b) (**Table 7** & **Table 8**).

Table 7. Reporting views produced by CAGE.

Reporting view	Primary audience	Typical questions answered
Operational remediation view	Security analysts, platform engineers, application teams.	What is affected? What action is required? Who owns it? What is the due date? What evidence confirms closure?
Risk and compliance view	GRC, audit, risk owners, security managers.	Which controls are affected? Which risks are accepted or overdue? Are exceptions approved and reviewed? Is evidence complete?
Executive posture view	Executives, boards, senior management, customer assurance leaders.	What are the material cloud risks? Are high-risk items reducing? Which business services remain exposed? Are assurance packs ready?

Table 8. Role-specific reporting views produced by CAGE.

Role	Relevant CAGE reporting view	Information needed	Governance value
Security analyst	Operational remediation view	Finding details, severity, affected resource, exposure context, recommended action, validation evidence.	Supports triage, investigation, remediation tracking, and technical closure.
Platform or cloud engineer	Operational remediation view	Resource ID, account or subscription, configuration state, owner, due date, change ticket, after-state evidence.	Links technical remediation activity to accountable evidence.
Application owner	Operational remediation view and risk/compliance view	Business service affected, asset criticality, treatment decision, due date, residual risk, exception status.	Connects cloud security findings to service ownership and business accountability.

Continued

GRC analyst	Risk and compliance view	Control mapping, treatment status, evidence completeness, exception approval, exception expiry, residual risk status.	Supports compliance tracking, audit preparation, and risk register updates.
Auditor or assurance reviewer	Risk and compliance view	Source finding, control relevance, before-and-after evidence, approval record, validation timestamp, residual risk decision.	Provides traceable evidence for assurance review.
Security manager	Risk and compliance view and executive posture view	Overdue high-risk items, recurring control gaps, exception trends, team ownership, evidence readiness.	Supports prioritisation, escalation, and governance oversight.
Executive or board-level stakeholder	Executive posture view	Material risks, trend summaries, exposed business services, accepted risks, assurance readiness.	Converts technical findings into decision-ready cyber risk information.
Customer assurance or third-party review lead	Executive posture view and risk/compliance view	Evidence pack status, control coverage, remediation progress, exception handling, assurance-ready summaries.	Supports customer assurance, supplier review, and external trust communication.

9. Alignment with Australian Industry and Employment Demand

The CAGE framework is relevant to Australian industry because it combines cloud security operations with risk management, compliance, evidence preparation, and stakeholder communication. These capabilities appear across several Australian role families, including cloud security engineer, security analyst, cyber GRC analyst, cyber assurance consultant, technology risk analyst, DevSecOps engineer, cloud engineer with security responsibilities, and security operations analyst.

Jobs and Skills Australia reported that approximately 70,900 people were employed as Database and Systems Administrators and ICT Security Specialists as of August 2025, and projected employment in this group to grow 14.2% from May 2024 to 2029, more than double the national average growth rate of 6.6% (Jobs and Skills Australia, 2025). This occupational grouping is broader than cloud security alone, but it is relevant because it includes ICT security specialists and roles responsible for system security, administration, reliability, and security policy implementation.

Australian employers increasingly value practitioners who can do more than operate a technical tool. CAGE demonstrates a hybrid capability profile: understanding cloud findings, assessing business risk, mapping issues to controls, coordinating remediation, documenting exceptions, preparing evidence, and explaining posture to non-technical stakeholders. This is particularly useful for organisations that need to answer customer security questionnaires, prepare for ISO/IEC

27001 audits, demonstrate Essential Eight uplift, manage third-party assurance, or provide risk reporting to management (**Table 9**).

Table 9. Australian employment relevance of the CAGE framework.

Employer need	How CAGE demonstrates capability	Relevant role examples
Cloud security visibility	Normalises CSPM and vulnerability findings into consistent records.	Cloud security analyst, security engineer, SOC analyst.
Risk-based prioritisation	Combines severity with asset criticality, exposure, business impact, and compliance relevance.	Cyber risk analyst, technology risk consultant, security manager.
GRC and audit evidence	Creates evidence packs that link findings, controls, treatment, validation, and residual risk.	GRC analyst, cyber assurance consultant, internal audit support.
Remediation coordination	Assigns owners, due dates, tickets, treatment decisions, and validation evidence.	DevSecOps engineer, platform engineer, cloud operations lead.
Executive communication	Converts alert volume into posture trends, material risks, exceptions, and evidence readiness.	Security lead, governance manager, cyber program manager.

The framework should therefore be positioned not only as a research artefact, but also as evidence of professional capability. A practitioner who can implement CAGE can show employers that they understand the full lifecycle from detection to risk decision, remediation, assurance evidence, and management reporting.

10. Discussion

10.1. Contribution to Practice in This Sense, CAGE Extends the Authors' Broader Privacy and Assurance Research Agenda from Health Information Systems and User Privacy Behaviour into Cloud Security Governance (Shojaei et al., 2024, 2025a, 2025b)

The main practical contribution is a repeatable translation mechanism between cloud operations and governance. Instead of asking GRC teams to interpret raw findings or asking engineers to produce ad hoc evidence near audit time, CAGE embeds evidence requirements into the finding lifecycle. Each finding becomes traceable from detection to risk context, control relevance, treatment decision, implementation, validation, and residual risk disposition.

This approach can improve prioritisation because it combines tool severity with business and compliance context. It can improve audit readiness because evidence is collected during remediation rather than reconstructed later. It can improve stakeholder communication because the same underlying finding can be viewed as an engineering task, a risk record, a control gap, and an executive posture indicator. It can also improve accountability because owners, due dates, decisions, and exceptions are explicit.

10.2. Alignment with Australian Assurance Expectations

Australian organisations operate in an environment where cyber resilience, third-party dependence, privacy expectations, customer assurance, and board-level accountability are increasingly important. APRA CPS 234 directly applies to APRA-regulated entities, but it illustrates a broader assurance expectation: regulated entities must maintain information security capability commensurate with threats and must consider information assets managed by related parties or third parties (Australian Prudential Regulation Authority, 2019). The Essential Eight provides a widely recognised mitigation baseline and encourages a risk-based approach to implementation and exception handling (Australian Signals Directorate, 2023).

CAGE is useful beyond narrow compliance. It helps organisations explain whether cloud controls are operating as intended, where drift has occurred, how remediation is progressing, and what evidence supports the current posture. This is valuable for regulated entities, public-sector suppliers, technology companies, managed service providers, SaaS providers, and organisations preparing for customer security questionnaires or certification activities.

10.3. Boundary of Claims

The framework should not be interpreted as an automated compliance certification method. Mapping a finding to ISO/IEC 27001, NIST CSF, Essential Eight, or CSA CCM does not prove full compliance with those frameworks. It creates traceability between a technical condition and a governance concern. Formal compliance still requires scope definition, management-system review, audit judgement, documented policies, operational evidence, leadership oversight, and organisational context.

The framework also does not claim that all cloud findings can be scored objectively. Risk scoring includes judgement, and different organisations may assign different weights to exposure, business impact, and compliance relevance. The value of the framework is that it makes that judgement explicit, documented, and reviewable.

11. Limitations and Future Research

This paper has limitations. First, the framework is conceptual and practice-oriented. It has not yet been empirically validated using real organisational cloud findings. Second, the worked example is synthetic and should not be treated as evidence of effectiveness in a production environment. Third, the framework uses AWS examples for clarity, which may underrepresent provider-specific differences in Azure, Google Cloud, hybrid cloud, and multi-cloud environments. Fourth, the proposed scoring model is illustrative and should be calibrated against organisational risk appetite, asset classification, threat context, and incident history. Fifth, control mapping can be subjective; future work should evaluate inter-rater consistency among cloud security, GRC, and audit practitioners.

Future research should apply CAGE in one or more organisational case studies.

A useful evaluation could compare alert handling before and after framework adoption using metrics such as time to triage, time to remediation, overdue risk count, evidence completeness, exception quality, and stakeholder satisfaction. Another direction is tool-supported implementation: findings could be ingested from cloud APIs, enriched using asset inventory and tagging, mapped to controls using a maintained ruleset, and exported into GRC dashboards or evidence registers. A third direction is validation in Australian job-market contexts, examining how cloud security and GRC teams collaborate when evidence is generated continuously rather than collected retrospectively.

12. Conclusion

Cloud security tools can identify misconfigurations, vulnerabilities, exposure, and control failures, but technical alerts do not automatically become governance evidence. Organisations need a practical mechanism for translating findings into risk-based decisions, treatment actions, exception records, and audit-ready artefacts. This paper proposed the Cloud Alert-to-Governance Evidence (CAGE) framework to address that need.

CAGE classifies cloud findings through severity, asset criticality, exposure, business impact, and compliance relevance; maps them to control frameworks; assigns accountable treatment; and packages the outcome as governance evidence. The framework contributes a clear bridge between cloud security operations and GRC practice. For industry, it supports better prioritisation, stronger remediation discipline, clearer communication, and more defensible assurance. For Australian employment and professional development, it demonstrates a valuable hybrid skill set that connects cloud security, risk management, compliance, audit evidence, and executive communication. For future research, it provides a conceptual foundation for empirical validation, automation, and refinement across cloud platforms and organisational contexts.

Acknowledgements

The authors acknowledge Cultural Infusion for academic and professional support. AI-assisted tools were used to support language refinement, structure, and consistency checking. The authors reviewed and edited the manuscript and take responsibility for the accuracy, integrity, and final content of the work.

Author Contributions

P.S. wrote and edited the manuscript. R.M. reviewed the manuscript. Both authors have read and approved the final version of the manuscript.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- Amazon Web Services (n.d.-a). *Introduction to AWS Security Hub CSPM*. AWS Documentation. <https://docs.aws.amazon.com/securityhub/latest/userguide/what-is-securityhub.html>
- Amazon Web Services (n.d.-b). *AWS Security Finding Format (ASFF)*. AWS Documentation. <https://docs.aws.amazon.com/securityhub/latest/userguide/securityhub-findings-format.html>
- Amazon Web Services (n.d.-c). *What Is Amazon Inspector?* AWS Documentation. <https://docs.aws.amazon.com/inspector/latest/user/what-is-inspector.html>
- Australian Prudential Regulation Authority (2019). *Prudential Standard CPS 234 Information Security*. APRA. https://www.apra.gov.au/sites/default/files/cps_234_july_2019_for_public_release.pdf
- Australian Signals Directorate (2023). *Essential Eight Maturity Model*. Australian Cyber Security Centre. <https://www.cyber.gov.au/business-government/asds-cyber-security-frameworks/essential-eight/essential-eight-maturity-model>
- Cloud Security Alliance (2026). *Cloud Controls Matrix and CAIQ v4.1*. CSA. <https://cloudsecurityalliance.org/artifacts/cloud-controls-matrix-v4-1>
- Dempsey, K., Chawla, N. S., Johnson, A., Johnston, R., Jones, A. C., Orebaugh, A., Scholl, M., & Stine, K. (2011). *Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations (NIST Special Publication 800-137)*. National Institute of Standards and Technology.
- Hashizume, K., Rosado, D. G., Fernández-Medina, E., & Fernandez, E. B. (2013). An Analysis of Security Issues for Cloud Computing. *Journal of Internet Services and Applications*, 4, Article 5. <https://doi.org/10.1186/1869-0238-4-5>
- International Organization for Standardization (2022a). *ISO/IEC 27001: 2022 Information Security, Cybersecurity and Privacy Protection—Information Security Management Systems—Requirements*. ISO. <https://www.iso.org/standard/27001>
- International Organization for Standardization (2022c). *ISO/IEC 27005: 2022 Information Security, Cybersecurity and Privacy Protection—Guidance on Managing Information Security Risks*. ISO. <https://www.iso.org/standard/80585.html>
- International Organization for Standardization. (2022b). *ISO/IEC 27002: 2022 Information Security, Cybersecurity and Privacy Protection—Information Security Controls*. ISO. <https://www.iso.org/standard/75652.html>
- Jobs and Skills Australia (2025). *Cyber Security Skills in Demand as Labour Market Evolves*. Australian Government. <https://www.jobsandskills.gov.au/news/cyber-security-skills-demand-labour-market-evolves>
- Joint Task Force Transformation Initiative (2012). *Guide for Conducting Risk Assessments (NIST Special Publication 800-30 Revision 1)*. National Institute of Standards and Technology.
- Mell, P., & Grance, T. (2011). *The NIST Definition of Cloud Computing (NIST Special Publication 800-145)*. National Institute of Standards and Technology.
- National Institute of Standards and Technology (2024). *The NIST Cybersecurity Framework (CSF) 2.0*. National Institute of Standards and Technology.
- Shojaei, P. (2016a). A Hybrid Algorithm for Task Scheduling Based on Ant Colony Opti-

- mization and Local Neighborhood Search. In *International Conference on Engineering and Applied Sciences* (pp. 95-101). Higher Education Forum.
- Shojaei, P. (2016b). An Approach to Optimized Genetic Algorithm for Task Scheduling in Cloud Computing. In *International Conference on Computer Engineering and Information Technology* (pp. 54-62). Higher Education Forum.
- Shojaei, P. (2016c). Threshold Acceptance Approach for Task Scheduling in Cloud Computing. In *International Conference on Engineering and Applied Sciences* (pp. 54-64). Higher Education Forum.
- Shojaei, P. (2018). An Approach to Optimized Imperialist Competitive Algorithm for Task Scheduling in Cloud Computing. In *5th International Conference on Electrical Engineering and Computer Science* (pp. 253-260). Francis Academic.
- Shojaei, P., & Moieni, R. (2025). Empirical Analysis of Data Privacy Concerns in Dei. *Open Journal of Social Sciences*, 13, 83-110. <https://doi.org/10.4236/jss.2025.136006>
- Shojaei, P., & Naji, H. R. (2016). A Hybrid Particle Swarm Optimization for Task Scheduling in Cloud Computing. In *3rd International Congress on Electrical Engineering, Computer Sciences and Information Technology* (pp. 88-95). IOP Publishing.
- Shojaei, P., Chow, Y. W., & Vlahu-Gjorgievska, E. (2025b). User Privacy Concerns and Preferences in mHealth Applications: Balancing Privacy and Usability. In *Studies in Health Technology and Informatics* (pp. 229-233). IOS Press. <https://doi.org/10.3233/shti250835>
- Shojaei, P., Vlahu-Gjorgievska, E., & Chow, Y. W. (2024). Security and Privacy of Technologies in Health Information Systems: A Systematic Literature Review. *Computers*, 13, Article 41. <https://doi.org/10.3390/computers13020041>
- Shojaei, P., Vlahu-Gjorgievska, E., & Chow, Y. W. (2025a). Enhancing Privacy in mHealth Applications: A User-Centric Model Identifying Key Factors Influencing Privacy-Related Behaviours. *International Journal of Medical Informatics*, 199, Article 105907. <https://doi.org/10.1016/j.ijmedinf.2025.105907>
- Shojaei, P., Zameni, N., & Moieni, R. (2025c). Anonymity in the Age of AI. *Open Journal of Social Sciences*, 13, 48-72. <https://doi.org/10.4236/jss.2025.138004>
- Subashini, S., & Kavitha, V. (2011). A Survey on Security Issues in Service Delivery Models of Cloud Computing. *Journal of Network and Computer Applications*, 34, 1-11. <https://doi.org/10.1016/j.jnca.2010.07.006>
- Takabi, H., Joshi, J. B. D., & Ahn, G. J. (2010). Security and Privacy Challenges in Cloud Computing Environments. *IEEE Security & Privacy Magazine*, 8, 24-31. <https://doi.org/10.1109/msp.2010.186>

Appendix: Suggested Implementation Artefacts

A practical implementation of CAGE can begin with a spreadsheet, ticket workflow, or GRC register containing the following columns: Finding ID, Source Tool, Resource ID, Account/Region, Finding Type, Technical Severity, Asset Owner, Asset Criticality, Exposure, Business Impact, Compliance Mapping, Risk Statement, Treatment Decision, Ticket Link, Due Date, Status, Validation Evidence, Residual Risk, Exception Expiry, Reviewer, Evidence Pack URL, and Last Reviewed Date.