International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Tracing Security and Data Governance Controls across Cloud Boundaries

Seema Sharma[1], Dr. Praveen Kumar Gupta[2], Dr. Veluru Chinnaiah[3]

[1]Assistant Professor, Birla Institute of Technology, Noida Off Campus, Uttar Pradesh, India, and Research Scholar, Department of Computer Application, Arni University, Himachal Pradesh, India. Registration No. APCS0012A/22.
[2]Research Supervisor, Department of Computer Science and Engineering, Arni University, Himachal Pradesh, India.
[3]Co-Research Supervisor, Dr. Veluru Chinnaiah, Department of Computer Science and Engineering Arni University, Kathgarh, Indora, Himachal Pradesh, India

Abstract
A cloud security policy may be documented and enforced at its point of origin while losing coverage when an identity, data copy or incident crosses an administrative boundary. This article develops a workload-level method for locating these breaks across multi-cloud, hybrid-cloud and combined estates. The method links each control obligation to a boundary crossing, a named owner, an enforcement point and retrievable evidence. A deterministic walkthrough compares policy-only, evidence-only and evidence-plus-ownership rules against four scenario expectations specified in advance. The three rules agree with one, three and four of the four expectations, respectively. These counts describe constructed cases and illustrate how missing ownership and unavailable evidence affect the classification. An explanatory sequential field protocol specifies how a planned survey of approximately 250 Indian cloud professionals and 20–25 follow-up interviews could examine whether these breaks recur in practice. The article contributes a traceable assessment unit, transparent decision rules and a field protocol for examining policy-to-practice gaps.

Keywords: multi-cloud security, hybrid cloud, data governance, identity federation, audit evidence, shared responsibility

1 Introduction

An employee can sign in through a central identity provider while a provider-native administrator account remains active after departure. Likewise, a production database can comply with an internal location rule while its recovery copy follows a different route. In both cases, the apparent success of the first control check conceals an unresolved obligation at the next boundary. A workload trace allows the reviewer to identify exactly where coverage ends and who can resolve it.

Mell and Grance (2011) define hybrid cloud in terms of distinct cloud infrastructures connected for data or application portability. This study uses multi-cloud to denote use of at least two independent public cloud providers. These definitions yield three mutually exclusive comparison groups when applied to the production estate supporting a focal workload. The central question is where evidence and ownership break along a workload path, and whether the observed pattern varies by estate configuration.

2 Prior Work and Research Gap

Chandramouli and Butcher (2023) describe identity-based enforcement for cloud-native applications across environments. The architecture motivates separate examination of human identities, workload credentials and local roles, because enforcement can occur at multiple points. Cloud Security Alliance (2026) provides a catalogue of control objectives and guidance on allocating responsibilities across the cloud supply chain. Joint Task Force (2022) provides assessment procedures that can be tailored to determine whether controls are implemented and operating as intended.

Together these sources establish the control and assessment foundations. The open question addressed here is how to compare gaps at specific workload handoffs using the same observation unit across different estate types. The framework adapts existing assessment principles to a record that preserves the relationship among the obligation, the exact crossing and the evidence from both sides. The research contribution is the operationalization of that comparison and the testable field protocol.

Table 1. From established guidance to workload-level observations

Foundation	Relevant mechanism	Observation in this study
NIST SP 800-207A	Application and service identities support granular enforcement.	Trace one role or service credential from issuance through enforcement and revocation.
CSA CCM v4.1	Control objectives and shared responsibilities structure assurance.	Record the accountable team at each provider or private-cloud handoff.
NIST SP 800-53A Rev. 5	Tailored procedures support control assessment.	Record whether evidence was retrieved and matched to the workload, obligation and review window.

3 Research Questions and Assessment Framework

RQ1. At which boundaries do documented identity, data governance, monitoring and incident-response controls lose operational support?

RQ2. How does the proportion of documented but unsupported obligations vary by estate configuration after accounting for workload and organizational characteristics?

RQ3. Which handoff mechanisms explain the gaps and the cases in which complex estates maintain strong evidence?

The unit of assessment is one control obligation at one workload boundary during a defined observation window. For every crossing, the register captures the workload ID, origin and destination, triggering action, obligation, accountable role, enforcement point, artifact type, observation date and exception status. A handoff may cross providers, the private–public boundary or an internal administrative boundary such as transfer from operations to incident response. The same obligation may generate multiple records when it applies at multiple crossings.

For a documented and applicable obligation, classify an observation as supported when a current artifact demonstrates the stated control at the relevant boundary and an accountable role is identified. Classify it as unsupported when a known defect is recorded, including explicitly unassigned ownership, a stale or mismatched artifact, or failed retrieval under a documented check. Record unknown separately when a

respondent cannot establish whether the check took place. Retain out-of-scope items separately and record approved exceptions with their scope, approval and expiry. Use a 90-day observation window for the survey and record artifact-specific freshness requirements before scoring.

Table 2. Control tracing prompts

Domain	Boundary to trace	Evidence to examine
Identity and access	Central identity provider to provider-native role or service identity	Entitlement review, revocation record, credential expiry or emergency-account review
Data governance	Primary store to replica, backup, telemetry or support-access path	Copy-level inventory and configuration evidence matched to the permitted destination
Monitoring	Provider event source to central investigation queue	End-to-end event trace with timestamps, coverage and retrieval record
Incident accountability	Internal triage to provider escalation or private-cloud containment	Acknowledgement, accountable role and dated exercise or action record

4 Exploratory Scenario Analysis

Four hypothetical cases were constructed to show how the rules behave when policy, evidence and ownership separate. Each case has an expected decision determined by a concrete scenario fact. The review compares three nested decision rules. A policy-only rule accepts any documented obligation. An evidence-only rule requires documentation and a current, matching artifact. The evidence-plus-ownership rule also requires an accountable role. A rule agrees with the expected outcome when both classify the case as supported or both classify it as unsupported.

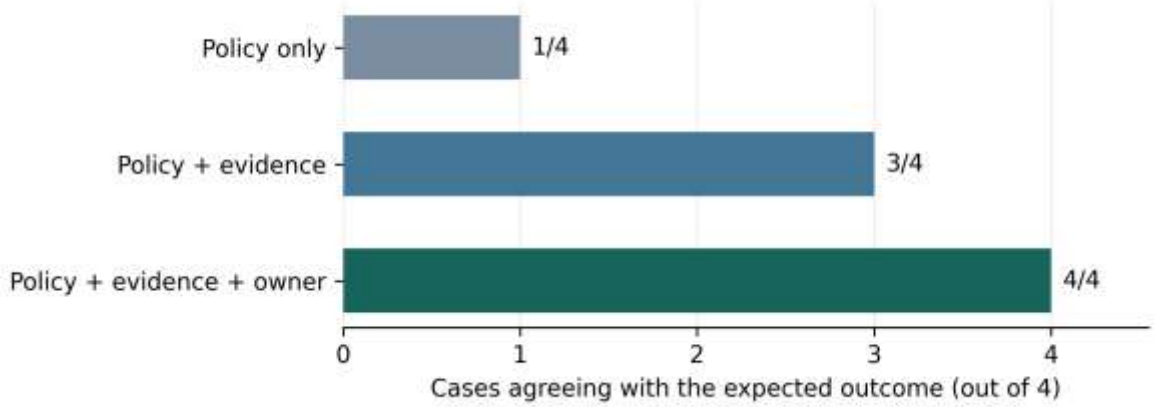
Table 3. Inputs and expected outcomes for the four constructed cases

Case and scope	Scenario fact	P / E / O	Expected
A Federated departure	Revocation is confirmed at the destination and assigned to the identity team.	1 / 1 / 1	Supported
B Recovery copy	The copy moves outside the approved internal destination and the artifact covers only the primary store.	1 / 0 / 1	Unsupported
C Incident handoff	A valid escalation record exists but no team owns the receiving queue.	1 / 1 / 0	Unsupported
D Log correlation	The current exercise cannot reconstruct the crossing. The available record concerns an	1 / 0 / 1	Unsupported

	older workload.		

P = documented policy. E = current, matching evidence. O = accountable owner. A value of 1 indicates presence, and 0 indicates absence.

Figure 1. Agreement between decision rules and expected outcomes in the constructed cases



Note. Counts are derived from Table 3. Cases were constructed to exercise distinct conditions and are illustrative.

The policy-only rule accepts all four cases. Requiring evidence changes the decisions for the recovery copy and the log-correlation exercise. Adding ownership also changes the incident handoff decision, where the artifact shows that escalation occurred but leaves responsibility for the next step unresolved. The apparent advantage of the fuller rule reflects the case design. The walkthrough establishes how the decision logic behaves under specified conditions and supplies no estimate of field accuracy or control effectiveness.

5 Field Study Protocol

5.1 Sampling and classification

The planned explanatory sequential design begins with purposive recruitment of approximately 250 adult cloud professionals in India. Eligibility requires at least one year of relevant work and direct knowledge of a production estate. Each respondent selects one focal production workload and reports on the estate supporting it. Multi-cloud only requires at least two independent public providers without an integrated private cloud. Hybrid only requires a private cloud integrated with one public provider. Combined estates meet both conditions. Screen out ambiguous and out-of-scope cases, record reasons, and monitor group sizes and sector coverage during recruitment. The target is a feasibility figure and group contrasts may remain imprecise.

5.2 Instrument and validation

For each eligible control, ask separately whether a policy exists, whether the respondent can identify the accountable role, what the most recent check showed and whether evidence was personally seen. Keep "unknown" and "not applicable" as distinct options. Ask independent cloud security, governance and methods reviewers to evaluate item relevance and clarity, then conduct cognitive interviews and a small pilot to check how respondents interpret the boundary and evidence questions. Version the instrument and freeze eligibility, scoring and exclusion rules before field analysis. Assess internal consistency only

for scales specified as reflective, since the individual control checks capture different operational conditions.

5.3 Quantitative analysis

Report counts for supported, unsupported, unknown, and out-of-scope obligations by domain and deployment group. The primary gap proportion is unsupported documented obligations divided by documented obligations with a determinable status. Publish the numerator, denominator and unknown count. Recompute the proportion with unknowns treated first as supported and then as unsupported to show how incomplete knowledge could shift the estimate. As an additional sensitivity check, restrict the analysis to responses where the respondent personally viewed evidence or personally observed the failed check.

Model support status among determinable documented obligations using logistic regression with adjustment for domain, sector, organization size, workload criticality and respondent role, subject to adequate overlap and sample size. Account for repeated control observations within respondents and within organizations where a privacy-preserving organization identifier is available. Report adjusted marginal differences with 95% confidence intervals and present governance maturity in a separate explanatory model using indicators that exclude the scored evidence and ownership items. If sparse cells prevent stable adjustment, report descriptive estimates and denominators. Interpret all estimates as exploratory associations within the sample.

5.4 Interviews and integration

Invite approximately 20–25 survey respondents for interviews, selecting cases across estate types and contrasting response profiles. In particular, include documented controls with missing evidence and combined estates with strong evidence. Use a recent non-sensitive example to trace the initiating event, handoff, actor and verification step. Code transcripts by domain and emergent mechanism, retain contradictory cases and keep a record of coding decisions. Following Fetters et al. (2013), integrate the phases in a joint display that links each quantitative pattern to a plausible mechanism, a disconfirming case and the resulting interpretation.

5.5 Ethics and data availability

Recruitment and recording are contingent on institutional ethics approval and informed consent. The protocol requests no credentials, customer records, sensitive topology or identifiable incident records. Store follow-up contacts separately and redact indirect identifiers before reporting quotations. For location questions, use a dated organization-approved requirement for the focal workload and record its source and scope. The manuscript includes the full inputs for the scenario analysis in Table 3. Field records will become available for analysis only after approved recruitment and collection.

6 Discussion

The scenarios show how three questions can yield different answers for the same control. Documentation identifies the obligation. Operational evidence establishes what happened at the selected crossing. Ownership establishes who can respond to a deviation. Collapsing these into a single yes-or-no question hides whether the obstacle is missing coverage, a retrieval problem or an unassigned role. Keeping them separate also makes an interview more specific because the interviewer can ask about the exact handoff where the trace ended.

Estate categories describe architecture, while workload routes and operating arrangements determine how many handoffs require verification. A combined estate can have strong evidence when enforcement

and incident ownership are standardized. A smaller estate can have weak evidence when local identities or recovery paths fall outside centralized review. The field study will record the number and type of crossings to explore these explanations and will retain high-performing complex estates as cases that challenge a simple complexity explanation.

7 Limitations and Conclusion

The four scenarios are intentionally selected and partly encode the logic being demonstrated. They establish neither predictive validity nor the completeness of the checklist. The survey will rely on practitioner reports, and respondents with greater visibility may report more gaps simply because they can detect them. Purposive recruitment, unequal group sizes and incomplete organizational identifiers may limit comparisons. The proposed interviews can clarify how gaps arose in particular settings, while independent artifact review would be needed to establish agreement with audited operating practice.

A workload-level trace makes a governance gap locatable. It records which obligation applied, where control passed between environments, what evidence could be retrieved and who owned the next action. The scenario walkthrough demonstrates the decision rules and the field protocol specifies how to examine their usefulness in Indian cloud estates.

References

1. Chandramouli, R., & Butcher, Z. (2023). *A zero trust architecture model for access control in cloud-native applications in multi-cloud environments* (NIST SP 800-207A). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207A
2. Cloud Security Alliance. (2026). *Cloud Controls Matrix and CAIQ v4.1*. https://cloudsecurityalliance.org/artifacts/cloud-controls-matrix-v4-1
3. Fetters, M. D., Curry, L. A., & Creswell, J. W. (2013). Achieving integration in mixed methods designs—Principles and practices. *Health Services Research, 48*(6 Pt 2), 2134–2156. https://doi.org/10.1111/1475-6773.12117
4. Joint Task Force. (2022). *Assessing security and privacy controls in information systems and organizations* (NIST SP 800-53A Rev. 5). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-53Ar5
5. Mell, P., & Grance, T. (2011). *The NIST definition of cloud computing* (NIST SP 800-145). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-145