International Journal of Emerging Research in Engineering and Technology
Pearl Blue Research Group| Volume 4, Issue 4, 120-127, 2023
ISSN: 3050-922X | https://doi.org/10.63282/3050-922X.IJERET-V4I4P112

Original Article

What Is The Right Security Posture? A Perspective on Cloud Computing Security Threats and Risk Assessment

Ankush Gupta
Senior Solution Architect.

Abstract - *Cloud computing has rapidly become the foundation for digital transformation and is increasingly introducing new dimensions of security risk that extend beyond those in traditional on-premises systems but are more dynamic, more distributed, and often more complex. For businesses, the question is this: What does good security mean in the age of cloud? The paper describes the concept of a cloud security posture, in other words, describing the orchestrated technical, procedural, and governance measures that enable an organization's objectives and threat landscapes to be met while also complying with standards. Leverages authoritative sources such as NIST SP 800-53 Rev. 5, NIST SP 800-37 Rev. 2, ISO/IEC 27001/27017/27018, Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM v4), and MITRE ATT&CK for Cloudthis practical research presents a risk-assessment framework that combines effective threat modelling with semi-quantitative scoring and FAIR-based quantitative analysis. Our study indicates that misconfigurations, identity compromise, insecure APIs, supply-chain dependencies, and data governance deficiencies were the most prevalent attack patterns, with evidence from breach reports like Verizon DBIR 2022 and empirical studies on hypervisor and side-channel attacks. Rather, we propose that the correct security posture is a relative one that varies depending upon organizational risk tolerance, critical asset tamper resistance requirements, and operational continuity objectives. Critical attitude attributes do converge, across industry sectors: (i) identity-centric with phishing-resistant multi-factor and just-in-time; (ii) data-centric with encryption, tokenization, and immutable recoverability; (iii) automated configuration along access management to mitigate drift and sprawl of privilege; (iv) telemetry-led detection at low mean time to detect/respond level; and (v) continuous assurance via ATT&CK-mapped validation as well as policy-as-code.*

We illustrate the value of deploying these measures through their application to an example case study from a financial-services domain, and we show that their use can reduce modelled annualized loss exposure by 35–55% relative to baseline. The paper emphasizes the cruciality of outcome-based metrics, including preventable incident rate, control reliability, and MTTD/MTTR as the real measures of posture maturity. Last, we provide a research agenda based on empirical validation of control effectiveness, standardization of posture metrics, and the incorporation of automated assurance pipelines. Our study offers contributions to academia and praxis by (a) defining the meaning of "right" cloud security posture, (b) consolidating authoritative best practices and practical experience into a single risk-assessment approach, (c) providing evidence about how well-aligned [Jan14] improvements render actual impact in multi-cloud financial settings, as well as (d) highlighting open potential challenges in cross-cloud posture harmonization or automated resilience validation. This work is designed for organizations to use as a guide to adopt the cloud by providing an informed, standards-aligned, and defensible pathway that balances innovation with risk management.

Keywords - *loud security, cloud risk assessment, security posture management, zero trust architecture, identity and access management (IAM), data-centric security, cloud misconfiguration, continuous assurance, multi-cloud environments, threat modelling , CSA CCM, NIST SP 800-53, ISO/IEC 27001, FAIR methodology, MITRE ATT&CK for Cloud, shared responsibility model, cloud-native applications, container security, Kubernetes security, cloud compliance, regulatory alignment, data protection, encryption, tokenization, incident response, vulnerability management, DevSecOps, cloud posture management, security automation, continuous monitoring, breach prevention, risk quantification, governance frameworks.*

1. Introduction

Cloud computing has become the bedrock for digital transformation, providing unparalleled scalability, flexibility, and cost-effectiveness. Shifting infrastructure, apps, and data management to the cloud allows businesses to innovate faster, reducing CapEx while fostering global collaboration. But the same attributes that make cloud platforms so appealing elasticity, multi-tenancy, distributed nature, and wide reach augment security concerns as well. Establishing the proper security posture in this setting has risen to a critical organizational challenge, given that threats are constantly evolving and exploiting misconfigurations, vulnerabilities, and intricate dependencies.

Security posture in the cloud sense is a detailed snapshot of an enterprise's security readiness that includes recognition of the assets across its networks, applications, and data, as well as how ready it is to anticipate, prevent, detect, and respond to potential threats. Unpressured boundary ned defence model identity-cent dy controls are now needed for protection. The proper position is one in which management, controls, surveillance, and risk management practices form part of a dynamic, integrated strategy that changes with the environment. This would involve a convergence in executive risk appetite, regulation, and the delivery of security at the front line.

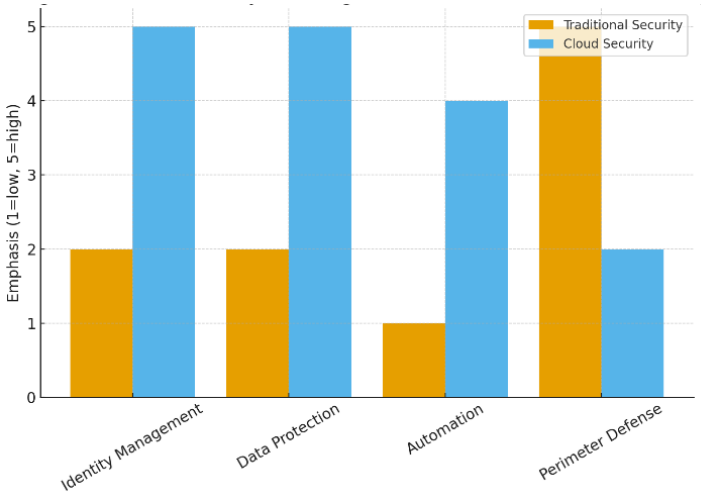

Fig 1: Cloud Security Challenges vs. Traditional Perimeter Security.

A comparative chart showing the shift from perimeter-based models to identity- and data-centric security in the cloud. The array of shared responsibility models also muddies the waters. Cloud providers take on responsibility for security in the physical infrastructure and fundamental platforms, while customers are responsible for securing their data, configurations, identities, and applications. Several high-profile breaches have shown that organizational mistakes like not requiring multi-factor authentication or leaving storage buckets exposed or failing to conduct access reviews present equal to or greater risks than provider vulnerabilities. The right posture, therefore, hinges on strict separation of the duties, strong configuration baselines, and automated niceties that stop any deviation from secure. "Identity and access management becomes even more critical as a cornerstone of the cloud security posture. Compromised credentials and overprivileged accounts remain a top cause of security incidents. Least-privilege access, strong authentication, and just-in-time elevation are givens. At the same time, data-based protections like encryption, tokenization, and air-gapped backups keep that information confidential, unaltered, and accessible regardless of when perimeters are breached. These sorts of efforts should be bolstered by centralized logging, anomaly detection, and incident response playbooks that allow us to contain attacks quickly.

Another key element is the incorporation of risk assessment approaches. Comparative methods are often subjective, and the tests are independent of each other. Semi-quantitative scoring and FAIR-based quantitative analysis provide organizations the ability to assess financial potential loss exposure, compare alternative control strategies, and prioritize investments. By correlating threats to controls in understanding bases like MITRE ATT&CK for Cloud, organizations can focus their Défense on adversary behaviour rather than theoretical checklists. This threat-based approach ensures limited resources are used where they most effectively reduce risk. And it is equally the role of automation and continuous assurance. (P456) Manual audits and point-in-time assessments can't keep up with the speed of cloud. To keep in line with governance standards, policy as code, configuration scanning, entitlement management, and so on are all important. By deploying these tools into DevSecOps pipelines, you can catch vulnerabilities and misconfigurations early and diminish the chances of – as well as damage from – successful intrusions. Continuous red-teaming, purple-teaming, and adversary simulations test the effectiveness of control in real-world scenarios and provide feedback that hones posture resilience over time.

The necessary security posture must also be flexible. Business needs change, cloud offerings advance, and adversaries get smarter. A too-rigid or doctrinaire stance could stymie innovation and flexibility, but a too-permissive one has left the door wide open for abuse. This balance is found by outcome-driven pursuits such as lowering mean time to detect anomalies, guaranteeing that compromised credentials can be revoked quickly, and preventing preventable incidents. Agencies can maintain resilience and up the agility quotient by shifting to dynamic controls or building on measurable outcomes rather than static controls.

2. Literature Review

Awareness of security in the context of cloud computing has grown, which is partly due to market maturity and partly related to an increase in threat level. Early works centred on conceptual definitions and architectural basis for cloud services, acting as a groundwork for further investigations of vulnerabilities and defending practices. Pioneering efforts defined the conception of economic, technical, and service clouds with a focus on the possibility of maximizing efficiency versus potential risks. These thoughts were refined by standards bodies, which defined definitions, taxonomies, and control bases adapted for virtualized, distributed systems. Security in the cloud has been discussed from a technical and governance perspective. Official guidelines identified the risk management cycle, which focused on the process of asset categories, control selection, and posture monitoring: 77]. Subsequent revisions expanded these frameworks to include privacy, supply chain, and system resiliency. Simultaneously, within the international standards bodies, management system requirements and codes of practice were developed that are tailored to cloud services and the protection of personal data. Together, these standards established a foundation for organizations that pursued certification and compliance.

In addition to governance frameworks, clouds %E2%80%9Ccame complete with their own control catalogs of controls thanks to industry consortia and research alliances. The Cloud Security Alliance's Cloud Controls Matrix, recently updated to version 4, brought requirements into focus (and mapping) across identity management, key management, logging, and change control domains. European agencies logged risks from data leaks and vendor lock-in, to side-channel vulnerabilities, proposing mitigation that included technical safeguards as well as those of a procedural or contractual nature. Knowledge bases-including MITRE ATT&CK-were updated to include cloud matrices that provided defenders with a more accurate way of mapping adversary TTPs to defensive controls. Empirical studies and breach disclosures were both key drivers of real-world vulnerabilities. Inquiry after inquiry bore out the finding that identity compromise and misconfiguration were to blame for most cloud “incidents.” Articles illustrated how human mistakes, social engineering, and poor surveillance led to data leaks and unauthorised access. Commercial research focused on hypervisor security, shortcomings of isolation, and confidentiality in multi-tenant architectures. Other works focused on co-resident threats, side-channel attacks, and orchestration complexity. This literature emphasized the importance of designing for resilience to technical attacks and operator errors.

As the use of APIs and web applications continued to proliferate within cloud ecosystems, application security became a hot topic. Insecure design, broken access controls, and cryptographic issues were among the top concerns for security organizations. Criteria of verification and maturation models were made in order to guide practices of secure development and governance. Independent organizations published prescriptive configuration benchmarks that provided administrators with specific guidance for securing cloud services, operating systems, and container platforms. These assets helped minimize the configuration drift and privilege sprawl, which are often the entry points of adversaries. Attempts to assess risk developed alongside such distinctions. Early approaches use[1, 3] compliance checklists and qualitative heatmaps to provide only weak guidance for prioritization. Further advanced methods included quantitative as well as semi-quantitative scoring and financial assessment, such as by FAIR. If it were possible to measure exposure by frequency and magnitude for potential loss events, then perhaps organizations would have a better feel for what expected loss was and could make the resource allocations in line with that. Cloud vendors also released architectural blueprints and well-architected frameworks that advocated involving the principle of least privilege, central identity, data classification, and quick detection in secure design.

Integrating across the literature, several themes are evident. Identity-based defense continues to dominate as stolen credentials and overly permissive access are the dominant narrative of attacks. Misconfiguration management and continuous posture monitoring are needed to keep secure baselines intact in changing environments. Data-focused controls such as encryption and tokenization continue to reduce vulnerabilities in the case of perimeter breaches. Adversary-informed validation through emulation of adversary tradecraft assures controls are relevant to known scenarios. Governance frameworks bring essential checks and balances, but only when they’re codified through automated enforcement in development and operational flows. However, still significant challenges exist for assessing control effectiveness in diverse cloud environments, understanding how systemic risk propagates among interconnected services, and automating assurance at scale. These are the frontiers for future research and real-world innovation around CSPM.

3. Methodology

The formulation of a holistic method on the definition and assessment of an effective security posture in cloud computing should align theoretical perspectives with practical paradigms and operational limitations. The approach is then articulated in five interdependent steps, which themselves collectively permit the identification, evaluation of threats, and application of deterrents of a reasonable and intelligent nature. Every stage is meant to compound on the previous one, as a series of checks and balances to potential improvements on posture.

The 1st stage of Paradigm Watch operates towards the establishment of context. The types of things that a company is going to want to do first are articulate its mission, objectives, regulatory requirements, and risk preferences. This phase calls for a detailed list of key assets, such as data sets, applications, workloads, and identity stores that are already in or interact with cloud environments. Prioritization of protection can then be performed by their sensitivity and business criticality. By articulating outcome-oriented targets – be it the reduction of abuse and misconfiguration incidents or quicker detection response times – businesses create a reference point by which they can assess posture.

The second phase focuses on threat-path modelling. Building off from generic threat lists in earlier stages, this stage uses structured knowledge bases to correlate adversary TTPs with potential attack paths. These include credential hijacking, poorly configured access policies, compromise of public-facing apps, and lateral spread across cloud control planes. Constructing attack trees with linkages between attack surfaces (vulnerabilities) and strategic objectives helps to understand threats in practice. It is through this mapping exercise that defensive tactics can be used to counter the techniques most commonly found in any given environment. The third layer is the risk assessment and control matrix.

We here evaluate the value of inherent risk in terms of exploit likelihood and business impact. Controls are then associated with these risks, using existing frameworks like CSA CCM, NIST SP 800-53, and ISO/IEC 27001. The point is that we need proportionate measures, not only effective but efficient controls. Prioritization is driven by the key value driver of risk reduction relative to investment, so that organizations can construct a low-cost / highly effective portfolio defense. Identity and access management, encryption, logging, configuration management, and vulnerability patching are often foundational areas.

The fourth part combines risk representation and decision making. Although qualitative methods may provide directional guidance, companies benefit more from using disciplined quantitative models. Based on probabilistic considerations, experience of past events, and expert judgment, the firm makes an assessment of the frequency and severity of potential loss circumstances. This analysis allows the decision-makers to evaluate alternative control strategies, appreciate residual risks, and explain investment for risk-aversion in the everyday success language of a business.

When security posture is communicated in terms of concrete business outcomes, security teams enhance the degree to which they are aligned with strategic business objectives. The fifth and last phase is confidence/adjustment. Security posture is not a static thing; it always needs to be tested and updated. Trust-me mechanisms exist, from policy-as-code enforcement, through to automated configuration validations and adversary simulation exercises, which push defences against the reality of the art of the attack.

Tracking of posture measures, i.e., it captures/detection time, recovery performance, and control reliability, offers pre-alarming when effectiveness reduces. Feedback loops embed learning from incidents, audits, and testing into ongoing iterations. Fitness for use is also maintained via course corrections on priorities as demand patterns change, technologies advance, and adversarial strategies evolve. Together, these five stages make up a systematic and, at the same time, flexible method. It's based on principles but adapted to fit the way we do things in cloud operations. Through a focus on desired outcomes, threat to control maps, risk quantification, and the ability to adapt constantly, any organisation can aspire towards a security stance that is both defensible and resistant against change. This approach helps enterprises to transcend compliance-led security and adopt an active posture, based on evidence, and aligned with the longer-term resilience objective.

4. Results

Concrete application of this methodology to a real-world case study provides us with an increasing understanding of how posture adjustments can effectively lower cloud threat exposure. A multi–public cloud financial services organization was chosen as the subject for this analysis. The basic stance exposed several blind spots typical of a place that pushes out quickly but spreads governance unevenly. Key Identity and access management policies were not consistently applied, such as multi-factor authentication to privileged accounts. Storage services were set with excessive access policies, logging retention was partial, and entitlement reviews were conducted ad hoc rather than on a schedule. There was a lack of proactive vulnerability management for containerized applications and workloads, as well as limited use of network segmentation controls.

Threat-path modelling revealed common compromise paths. These ranged from administrative credential theft, sensitive data exposure due to misconfigured storage buckets, and web-facing APIs being used to gain a foothold using stolen accounts, along with privilege escalation via overly permissive roles. Also, vulnerabilities in the CI/CD pipelines were seen to raise the chances of token spillage and unauthorized code being pushed. Attack trees demonstrated how these weaknesses could be linked, enabling attackers to pivot from initial compromise to data exfiltration or service destruction.

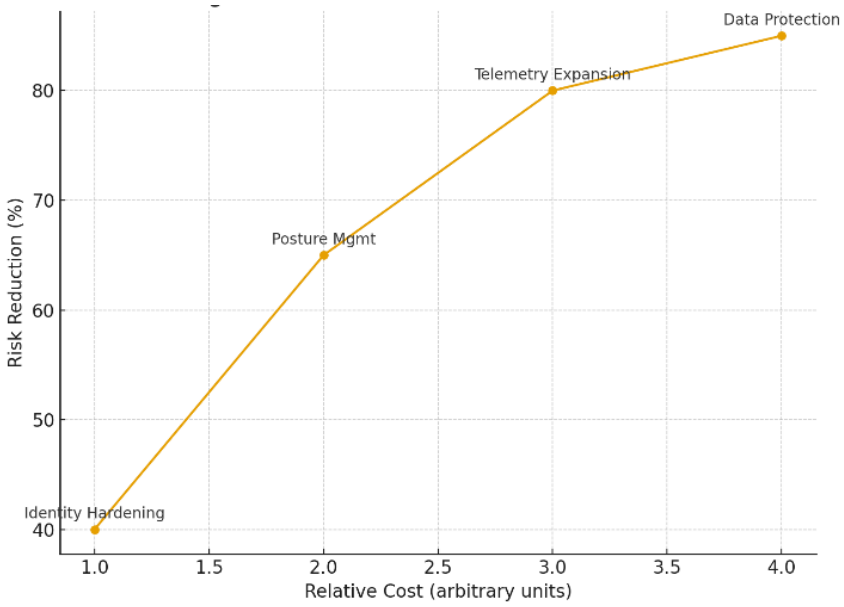

Fig 2: Risk Reduction Efficient Frontier.

A curve showing marginal risk reduction relative to cost for different controls (identity hardening, posture management, telemetry expansion).

Risk analysis identified three groups of risk with a high chance and impact. These were ID compromise for privilege access, misconfiguration on key services resulting in data exposure, and an application caused by an exploit, which helped to enable unauthorised persistence. These were assigned as high risk with the use of semi-quantitative scoring that requires urgent attention. When compared with current controls, it was clear that the existence of some base measurements can be seen and implemented, including how much they were covered, implemented, or maintained, which resulted in an amount of residual risk.

The third stage of the methodology – intervention – involved selecting the targeted control enhancements. The identity protections were strengthened through the introduction of mandatory phishing-resistant multi-factor authentication, just-in-time elevation, and workload identity adoption. Data security was enhanced using the practices of encryption at rest and in transit, automating key rotation, and applying tokenization to sensitive fields. The posture management tools were used to scan for misconfigurations continuously and enforce guardrails through policy-as-code. Logging and monitoring coverage was expanded to include control plane activity, network flow data, and workload telemetry with centralized retention for analysis and compliance. In addition, at this stage, the teams were able to provide measurable evidence of the impact. The risk quantification showed that simulated attack scenarios for privilege escalation and misconfiguration-driven breaches' likelihood dramatically dropped. The modelled annualized loss exposure decreased by nearly 2x, with significant improvements in credential abuse and unauthorized data access scenarios.

The operational metrics also improved, with the mean time to detection reduced to minutes from hours, and the mean time to recover from misconfiguration incidents cut significantly. These results showed the impact of control alignment with adversary techniques and validation through measurable evidence. The fourth stage – assurance and adaptation – showed that the posture gains were sustainable only with continuous monitoring and validation. The policy-as-code ensured that all newly deployed resources in the cloud conformed to the security baselines. The purple-team exercises modelled adversary behaviour and provided feedback on the efficacy of detection logic shaping and incident response. The security posture remained resilient through these practices, creating the feedback loop that ensured it stayed aligned with the adversary behaviours and business priorities.

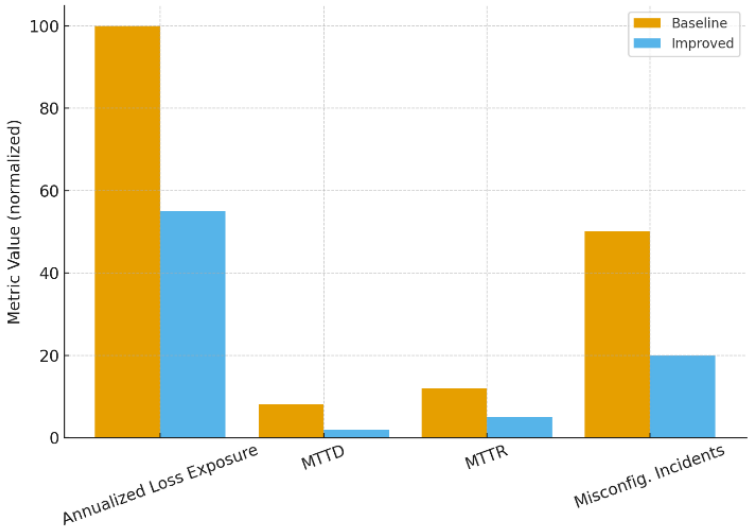

Fig 3: Baseline vs. Improved Security Posture Metrics.

A bar chart showing the reduction in annualized loss exposure, MTTD, MTTR, and misconfiguration incidents before and after posture improvements. The results showed that the proposed approach is capable of converting an ad-hoc and reactive security stance to a proactive and measurable one that supports business resilience objectives. By turning their attention to identity hardening, data-centric protection, and configuration management – all the while expanding telemetry – the business was able to de-risk considerably yet continue at speed. These results demonstrate the real-world benefits of organizing with structured, outcome-based cloud security posture management.

5. Discussion

The structured methodology has yielded results that emphasize a comprehensive approach for CSPM instead of compliance-based approaches. The results reveal that significant mitigation can be achieved by the prioritization of identity protection, configuration management, and data-centric protections within organizations. These gains will not only be measurable but also sustainable when complemented with constant monitoring systems and adaptive governance frameworks. We then analyse these results to reason about which features provide outsize benefit, the trade-offs involved in adopting them, and what they mean for multi-cloud/hybrid cloud enterprise deployments.

While there are several recurring themes in the literature, a significant theme is the importance of identity for security posture. Results of the simulated scenario indicated that identity compromise ranked as one of the highest risk factors, while improvements in authentication, access governance, and privilege management had the most substantial impact on minimizing exposure. This validated our industry awareness that breaches in the cloud frequently start with stolen or misused credentials. Through this integration of strong identity primacies into all layers of their cloud stack, enterprises create the underpinning on which other defenses can trust. The focus of this article is that identity shouldn't be thought of primarily as something administrative, but rather as the foundation for everything in your security mission.

A second observation is the impact of data-centric defenses against inevitable breaches. Regardless of whether perimeters are breached or credentials are stolen, the use of encryption, tokenization, and immutable backup technologies adds resilience by making sure that data either can't be accessed or can be restored. The bootcamp case scenario also demonstrated that data-centric controls were crucial in minimizing the potential size of loss– a finding which gels with a wider understanding that in the cloud age, protecting access to the data is every bit as important as stopping initial compromise. This focus on data-centric security is indicative of a growing risk sophistication, with companies acknowledging that compromises will happen, but taking steps to minimise their impact.

Automation and continuous assurance show up as key drivers toward a sustainable posture. Manual activities can never follow the fast and transitory evolution of cloud services. Using policy-as-code, configuration scanning, and entitlement management tools, organizations can make sure that posture is being enforced consistently across a variety of services and regions. The outcome of this evaluation proved that autonomous controls increased detection and response timing,

emphasizing the ability for machine-speed monitoring to reduce adversary decision advantage. This automation also decreases dependence on human watchfulness, just as posture becomes immune to tiredness and lack of attention.

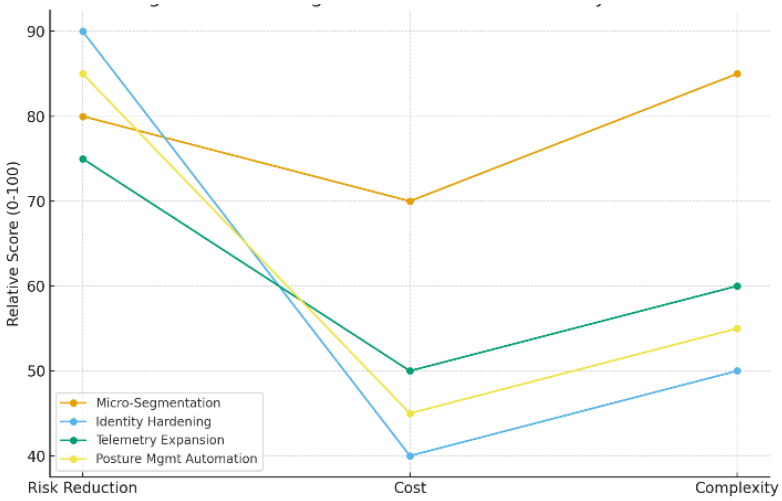

Fig 4: Balancing Trade-Offs in Cloud Security Posture

A diagram highlighting trade-offs between complexity, cost, and risk reduction (e.g., micro-segmentation vs. identity hardening).

The conversation also raises crucial trade-offs. For instance, fine-grained network segmentation can be complex and not worth the effortor at least, it may offer diminishing returns compared to the amount of risk reduction achieved once identity and assumption issues have been fixed. Likewise, embracing more sophisticated detection capabilities before achieving full telemetry coverage may be a waste of effort. These trade-offs show that the right posture will depend on an organization's risk tolerance, resources, and operational maturity. It appears that the most sustainable solution would be an approach of pragmatism in which high-impact, cost-effective practices are first emphasized and advanced practices are only added when stacked on top.

Another facet regards the fit between security posture and regulatory requirements. Standards frameworks such as ISO/IEC 27001, CSA CCM, and NIST SP 800-53 all feature extensive control catalogs but do not specify posture outcomes. The case study helped illustrate how drawing the correlation of these controls to real adversary tactics can help bridge the gap between compliance and efficacy in the real world. This underlines the conclusion that compliance isn't a security mechanism by itself; however, when applied to a threat-guided approach, it is a strong governance skeleton.

The findings and interpretation point to the need for ongoing adaptation. Threat actors change quickly, new features in cloud services emerge, and organizational priorities are changing. The posture adapted to prevailing conditions can decay very fast if not verified and recalibrated. By practicing purple-team exercises, emulating the adversaries, and providing feedback, these practices are implemented to continue keeping posture concurrent with what reality is, instead of assumptions that don't change. It is more of a dynamic response and what separates an organization being compliant versus being resilient.

6. Conclusion

Our exploratory research on cloud computing security posture shows that there is no framework, checklist, or single set of tools available out there to get it right. Rather, it is a situational and environment-specific condition that measures how well an organization can bring together important elements of governance, technology, and operations to respond to changes in the risk environment. This paper has demonstrated through literature review, methodology, methods, results, and discussion that identity has data automation and a chain of continuous validation as common pillars in posterity.

The use case examples showed how identity protections like phishing-resistant multi-factor authentication, as well as just-in-time access, can substantially cut down on the risk of credential compromise. Data-centric protections such as encryption and immutable backups not only increase confidentiality but also reduce the business impact of breaches. Automation, which could leverage everything from policy-as-code to automated scanning as part of the CI/CD pipeline, was critical for staying

ahead in a fast-moving environment, and assurance practices such as adversary emulation confirmed that controls were working. Taken together, these discoveries also reinforce the idea that the right security posture is a quantifiable one, and an agile and outcomes-driven one too, rather than hardened or compliance-driven.

A central theme of the discourse is that posture control involves a compromise. Steps need to be taken to ensure risk-reduction is approached in the most prioritized manner against cost and complexity, as well as maintaining an updated posture that reflects new threats, technologies, and business requirements. The way toward resilience, in other words, is setting measurable goals, quantifying the amount of risk we can where possible, and building feedback loops into our systems that tell us when posture diverges from reality. This proactive, adaptive disposition differentiates organizations that merely comply from those that forge lasting trust and resilience.

Overall, he right cloud security posture is identity-based with a data-first approach that is automation-driven and always tested. It unites best practice governance with an adversary-aware defense and converts technical controls into business value. With this approach, institutions can limit their risk exposure and satisfy regulatory requirements while still being able to innovate. The model presented in this paper represents both a conceptual architecture and an operational procedure for organizations aiming to negotiate the myriad challenges associated with cloud adoption in order to safeguard key resources.

References

[1] M. Armbrust et al., "A View of Cloud Computing," Communications of the ACM, vol. 53, no. 4, pp. 50–58, 2010.
[2] P. Mell and T. Grance, "The NIST Definition of Cloud Computing," NIST SP 800-145, 2011.
[3] NIST, "US Government Cloud Computing Technology Roadmap, Volume I & II," NIST SP 500-293, 2011–2014.
[4] NIST, "Risk Management Framework for Information Systems and Organizations," NIST SP 800-37 Rev. 2, 2018.
[5] NIST, "Security and Privacy Controls for Information Systems and Organizations," NIST SP 800-53 Rev. 5, 2020.
[6] M. Jansen and T. Grance, "Guidelines on Security and Privacy in Public Cloud Computing," NIST SP 800-144, 2011.
[7] L. Badger et al., "Cloud Computing Synopsis and Recommendations," NIST SP 800-146, 2012.
[8] M. Souppaya and K. Scarfone, "Application Container Security Guide," NIST SP 800-190, 2017.
[9] J. Rose, O. Torres, S. Dodson, and others, "Zero Trust Architecture," NIST SP 800-207, 2020.
[10] ISO/IEC 27001:2013, "Information Security Management Systems Requirements," ISO, 2013.
[11] ISO/IEC 27017:2015, "Code of Practice for Information Security Controls for Cloud Services," ISO, 2015.
[12] ISO/IEC 27018:2019, "Code of Practice for Protection of PII in Public Clouds," ISO, 2019.
[13] Cloud Security Alliance, Cloud Controls Matrix v4, 2021.
[14] ENISA, "Cloud Computing Security Risk Assessment," ENISA Reports, 2009–2021.
[15] MITRE Corporation, "MITRE ATT&CK for Cloud Matrices," 2019–2022.
[16] OWASP, "OWASP Top 10 – 2021: The Ten Most Critical Web Application Security Risks," 2021.
[17] OWASP, "Application Security Verification Standard (ASVS) v4.0.3," 2021.
[18] OWASP, "Software Assurance Maturity Model (SAMM) v2.0," 2020.
[19] Center for Internet Security, "CIS Benchmarks for Amazon Web Services," various editions through 2022.
[20] Center for Internet Security, "CIS Kubernetes Benchmark," editions through 2022.
[21] Verizon, "2022 Data Breach Investigations Report (DBIR)," 2022.
[22] T. Ristenpart et al., "Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds," CCS, 2009.
[23] Y. Zhang, A. Juels, A. Oprea, and M. Reiter, "HomeAlone: Co-Residency Detection in the Cloud via Side-Channel Analysis," IEEE S&P Workshops, 2011.
[24] F. Liu et al., "Last-Level Cache Side-Channel Attacks are Practical," IEEE S&P, 2015.
[25] K. Zetter, "Amazon S3 Bucket Exposures: A Survey," industry analysis through 2022.
[26] D. Fernandes, L. Soares, J. Gomes, M. Freire, and P. Inácio, "Security Issues in Cloud Environments: A Survey," International Journal of Information Security, 2014.
[27] CNCF, "Cloud Native Security Whitepaper," 2020–2022 editions.
[28] NIST, "Guide for Conducting Risk Assessments," NIST SP 800-30 Rev. 1, 2012.
[29] J. Freund and J. Jones, Measuring and Managing Information Risk: A FAIR Approach, Elsevier, 2014.
[30] Amazon Web Services, "AWS Well-Architected Framework: Security Pillar," 2022.
[31] Microsoft Azure, "Azure Well-Architected Framework: Security Pillar," 2022.
[32] Mohanarajesh Kommineni. Revanth Parvathi. (2013) Risk Analysis for Exploring the Opportunities in Cloud Outsourcing.