Cloud - Native Security: Securing Serverless Architectures

Ishva Jitendrakumar Kanani¹, Raghavendra Sridhar²

Independent Researcher Email: *ijkanani02[at]gmail.com*

Independent Researcher Email: *princeraj01[at]gmail.com*

Abstract: The adoption of serverless computing platforms such as AWS Lambda has introduced new efficiencies in cloud - native application development while simultaneously shifting the burden of security from infrastructure management to configuration management. This paper presents a case study on securing a production - grade serverless architecture using AWS services. It explores real - world misconfigurations across IAM roles, public APIs, dependency vulnerabilities, and observability blind spots, identifying how these security gaps emerge in fast - paced development workflows. The study then documents practical remediation steps, including minimizing IAM policies, securing API Gateway endpoints, scanning dependencies, managing secrets, and enhancing centralized logging. Unlike traditional approaches that rely on perimeter security or infrastructure controls, this work highlights how security must be embedded into the fabric of service permissions, event handling, and function orchestration in serverless systems. The case study provides actionable insights for cloud - native teams seeking to improve application resilience while preserving the agility of serverless development.

Keywords: Serverless Security, Cloud - Native Applications, AWS Lambda, IAM Hardening, Application Observability, API Authorization, Zero Trust Cloud, DevSecOps, Cloud Observability, CI/CD Security

1. Introduction

Serverless computing, epitomized by platforms like AWS Lambda, represents a paradigm shift in software deployment. It abstracts server management, enables event - driven workflows, and allows applications to scale elastically without dedicated infrastructure. While this model reduces operational burdens, it also redistributes responsibility particularly in terms of security. As organizations adopt serverless to accelerate time - to - market and simplify deployments, they often inherit a complex and unfamiliar security posture.

Traditional perimeter - based defenses are ineffective in this new model. Instead of managing virtual machines, developers configure granular permissions, secure multiple event sources, and monitor dozens of short - lived execution environments. Despite the maturity of cloud infrastructure providers, the onus of configuring and maintaining secure serverless applications falls on development teams. Misunderstandings about the shared responsibility model or neglecting infrastructure - as - code practices can lead to overexposed endpoints, data leakage, and unintended privilege escalation.

This paper presents a detailed, implementation - level case study of a cloud - native application deployed using AWS Lambda, API Gateway, DynamoDB, and supporting services. It identifies misconfigurations and risk points typical in fast moving development environments and demonstrates how pragmatic, AWS - native hardening strategies can close those gaps. Unlike broader surveys or abstract threat models, this study grounds its findings in real deployment scenarios. This case study contributes to the emerging body of knowledge that bridges the gap between conceptual models and secure serverless engineering practices [6].

2. Literature Review

Security concerns in serverless systems have received growing attention in industry research and community - led frameworks. The OWASP Serverless Top 10 [1] outlines key categories of risk specific to serverless applications, including insecure event data, broken authentication, improper exception handling, and the risks of relying on external packages. These issues stem from the highly modular and event - driven nature of serverless platforms.

The Cloud Security Alliance (CSA) [3] and Snyk [2] have both emphasized the challenges of permission management, secure CI/CD pipelines, and maintaining function - level observability. These studies highlight how ephemeral execution and micro - permissioning complicate auditing and security enforcement. However, they often remain at the level of conceptual guidance and lack implementation - level validation in real - world systems. Furthermore, most do not walk through detailed remediation strategies, leaving practitioners without clear, actionable blueprints for mitigating risks.

Academic contributions remain limited. Some papers model threat landscapes and propose automated policy - generation techniques, while others suggest serverless - specific intrusion detection methods. Yet most published work lacks hands - on walkthroughs of fixing real - world serverless security issues within a specific provider ecosystem. This aligns with Hendrickson et al. [6], who argue that while serverless architectures reduce deployment friction, they can also obscure architectural complexity and increase long - term maintenance burdens.

This paper distinguishes itself by embedding security recommendations within an actual deployment lifecycle,

Volume 9 Issue 8, August 2020 www.ijsr.net

Licensed Under Creative Commons Attribution CC BY

mapping misconfigurations to their operational impact, and applying concrete mitigation strategies using tools available to any AWS development team.

3. Threat Models

The shift to serverless architectures introduces a reorientation of attack vectors, where traditional VM - level concerns like patching or open ports give way to configuration - based exploits and inter - service abuse. The security landscape of serverless applications diverges from traditional models due to their ephemeral nature, fine - grained event triggers, and cloud - native service integration. The primary threats observed in this case study fall into the following categories:

- Over Permissioned IAM Roles: Developers often grant broad IAM permissions to functions for expedience, especially during rapid prototyping. These wildcard (*) actions create an expansive attack surface if a single function is compromised, it could potentially read from or write to unintended AWS services, such as S3 buckets, databases, or even modify IAM configurations themselves. Baldoni [8] notes that ephemeral compute layers introduce systemic risks often overlooked in static assessments.
- 2) **Public API Exposure:** API Gateway allows direct public access to Lambda functions via RESTful or HTTP interfaces. When authentication is not enforced (e. g., via Cognito or JWT), attackers can enumerate endpoints, inject malicious payloads, or initiate denial of wallet attacks by triggering resource intensive functions repeatedly.
- 3) **Third Party Package Vulnerabilities:** Serverless functions typically bundle external packages to handle application logic, HTTP parsing, and database access. Dependencies like lodash, moment, or axios may introduce known vulnerabilities if not pinned to secure versions. Since functions are redeployed frequently, this risk is amplified if scanning is not automated.
- 4) **Event Injection & Misrouting:** Serverless applications are inherently reactive, responding to S3 uploads, DynamoDB changes, or SNS messages. Improperly validated event payloads can exploit business logic flaws. For instance, an attacker could craft an S3 event that manipulates metadata, triggering unintended Lambda behavior or causing data corruption.
- 5) **Invisibility and Logging Gaps:** Unlike traditional monoliths, serverless functions provide little runtime feedback unless explicitly instrumented. Without tools like CloudWatch Logs, AWS X Ray, or third party platforms, teams lack insight into invocation patterns, performance anomalies, or indicators of compromise.
- 6) **Data Leakage via Misconfigured Storage:** Resources such as S3 buckets or DynamoDB tables may be inadvertently exposed due to permissive policies. Sensitive data, access tokens, or user metadata can be accessed if the access boundaries are not tightly defined at the role and resource level.

This threat landscape demonstrates that serverless security is less about defending infrastructure and more about defending configuration, workflow integrity, and identity relationships. Each of these threats, though individually impactful, often compounds in real systems—highlighting the need for layered, defense - in - depth strategies at the function, API, and cloud policy levels.

4. Serverless Architecture

The application at the center of this case study was architected using a collection of AWS - managed services to support a stateless, event - driven workload. The core business logic resided in AWS Lambda, written in Node. js, and was invoked via Amazon API Gateway, which exposed HTTP endpoints to external clients. Authentication was handled through Amazon Cognito, which issued JSON Web Tokens (JWTs) for API authorization. Data persistence was implemented using Amazon DynamoDB, while Amazon S3 served as a storage layer for static content and event - triggering uploads. For monitoring and visibility, the system used Amazon CloudWatch and AWS X - Ray, and application dependencies were regularly scanned using Snyk [2] for known vulnerabilities.

Upon auditing the system, several security weaknesses became apparent. The IAM roles assigned to Lambda functions were overly permissive, often granting full access to all S3 resources via s3: * actions. This elevated privilege model significantly increased the blast radius of any potential compromise in violation of IAM best practices [4]. Additionally, several API Gateway endpoints were exposed without authentication or throttling controls, rendering the application susceptible to unauthorized access and brute force attacks. The use of outdated and vulnerable third - party libraries, such as older versions of popular utility packages, introduced critical dependency risks, including the possibility of remote code execution via published CVEs. Secrets such as database credentials and token signing keys were embedded in plaintext within environment variables, posing a risk of accidental exposure or misuse. Finally, the application suffered from a lack of comprehensive logging and observability, with several Lambda function errors not being captured or forwarded to a centralized logging system, making root cause analysis and anomaly detection difficult. These observability gaps reflect challenges described by Jackson and Tam [9], who emphasize the need for distributed tracing and centralized telemetry in ephemeral environments. To address these security gaps and harden the architecture against abuse, a series of targeted remediations were applied across access control, observability, and deployment hygiene.

To further enhance resilience and scalability, the deployment process itself was fortified through a secure CI/CD pipeline integrated with AWS CodePipeline, GitHub Actions, and AWS CodeBuild. Each commit triggered automated security and quality gates—including static code analysis using ESLint and dependency scanning with Snyk CLI. Build artifacts were stored in versioned S3 buckets with server - side encryption (SSE - S3) and access logging enabled. Deployment roles were restricted using IAM conditions tied to the source repository and job status, reducing the blast radius in case of credential leakage.

For runtime defense, Amazon GuardDuty and AWS CloudTrail were configured to monitor unauthorized API activity. All Lambda functions were wrapped with structured logging middleware that captured request metadata,

Volume 9 Issue 8, August 2020 <u>www.ijsr.net</u> Licensed Under Creative Commons Attribution CC BY

International Journal of Science and Research (IJSR) ISSN: 2319-7064 ResearchGate Impact Factor (2018): 0.28 | SJIF (2019): 7.583

execution time, and error traces, making it easier to trace attack vectors or performance issues post - deployment. To reduce latency in tracing, AWS X - Ray was augmented with custom subsegment annotations, enabling targeted investigation of specific code paths or services such as DynamoDB read/write failures.

The team also implemented periodic access audits using AWS Access Analyzer and custom scripts that flagged permission changes or escalations. Unused IAM roles and Lambda functions were automatically reported via a scheduled CloudWatch Event rule. Secrets Manager rotation policies ensured that credentials—such as database tokens or API keys—were refreshed regularly without manual intervention.

These layered enhancements not only addressed immediate misconfigurations but also embedded security into the application's operational lifecycle—extending protection beyond initial deployment and aligning with AWS's principles of continuous assurance.

Code hygiene was improved by enforcing strict dependency pinning through package - lock. json, ensuring consistent builds and preventing unintentional upgrades to vulnerable packages. IAM policies were refactored such that each Lambda function was assigned a role tailored to its specific access needs-for example, limiting S3 access to read - only operations within a defined bucket prefix. At the edge, OAuth 2.0 - based JWT authorization was enforced using Cognito user pools, ensuring only verified users could access protected API routes. Snyk [2] was integrated into the CI/CD pipeline to enforce build - time vulnerability checks. Secrets management was migrated to AWS Secrets Manager, enabling encrypted storage, controlled access, and automatic rotation of sensitive credentials. Observability was significantly enhanced by enabling structured logs and distributed tracing across all functions, with metrics aggregated into CloudWatch dashboards and GuardDuty configured to alert on anomalous activity patterns.

These improvements reinforced the principle that security in serverless environments must be treated as a first - class architectural concern and also align with the AWS Well - Architected Framework – Serverless Lens [5], which emphasizes least privilege and observability. By embedding controls into identity management, CI/CD pipelines, and runtime monitoring, the application transitioned from a prototype - level configuration to a production - grade, hardened deployment—demonstrating that strong security practices can coexist with the agility of serverless computing.

5. Conclusion

Securing serverless applications requires developers to shift their focus from infrastructure hardening to the precise configuration of service relationships, permissions, and event behaviors. As Baldoni [8] highlights, such systems introduce architectural risks that require proactive, design - level security interventions. This case study illustrates how even well - intentioned implementations can expose critical resources through seemingly minor oversights. Misconfigured IAM roles, unauthenticated APIs, or stale dependencies can serve as the point of entry for attackers in systems that otherwise have no open ports or long - running processes.

The findings here underscore the importance of treating security as a continuous concern—one woven into CI/CD pipelines, architectural decisions, and runtime monitoring practices, emphasizing that cloud - native architectures must still be grounded in rigorous security practices. Organizations must invest in automation and principle - driven design, enforcing least privilege, validating events at all trust boundaries, and maintaining observability through logs and telemetry.

While the specifics of this case are grounded in AWS, the principles extend to any serverless or event - driven architecture. As the industry shifts further toward ephemeral, micro - permissioned systems, the rigor of configuration and discipline of architectural security will determine whether these systems remain agile and secure. As serverless adoption accelerates, organizations will need to embed security governance into the architectural fabric, not merely the deployment pipeline. Embedding security at the architectural layer ensures that scalability and speed do not come at the cost of resilience—making security a core enabler of innovation rather than a reactive checkpoint.

Looking ahead, organizations that embed architecture - aware security from the start will be best positioned to harness the full potential of serverless technologies—securely, scalably, and sustainably.

6. Future Work

As serverless architectures continue to mature and scale, new opportunities emerge for advancing their security posture beyond configuration hardening. Future work may explore the implementation of Zero Trust principles in event - driven systems, where trust boundaries are continuously evaluated at each function invocation and across service integrations. This approach would help prevent lateral movement and enforce stricter identity - based segmentation.

Another promising area is the use of machine learning for anomaly detection in serverless environments. Inspired by architectures like TensorFlow [7], future implementations may leverage telemetry - enhanced ML models. By analyzing execution patterns, timing anomalies, and payload structures, ML models could identify behavioral deviations indicative of misuse or compromise—especially valuable in high - volume, ephemeral compute scenarios where traditional detection methods fall short. Such models could also be paired with serverless - specific runtime agents to enrich detection with low - latency telemetry, reducing response times for emerging threats. Future serverless runtimes may adopt introspection techniques akin to those proposed by Garfinkel and Rosenblum [10].

Additionally, expanding the security strategies in this case study into cross - cloud comparisons could reveal provider specific gaps and strengths. Investigating serverless security patterns across AWS Lambda, Azure Functions, and Google Cloud Functions would support a broader understanding of cloud - agnostic best practices.

Volume 9 Issue 8, August 2020 <u>www.ijsr.net</u> Licensed Under Creative Commons Attribution CC BY Finally, evaluating open - source security tooling—such as Open Policy Agent (OPA), Falco, or custom Lambda extensions—may offer cost - effective and extensible alternatives to proprietary AWS features, enabling wider adoption and community - driven security evolution. As the ecosystem matures, a taxonomy of standardized benchmarks for serverless security—akin to OWASP's Top 10 or CIS benchmarks—could further support platform - agnostic auditing and control validation.

References

- OWASP Foundation, "OWASP Serverless Top 10," OWASP Project, 2018. [Online]. Available: https: //owasp. org/www - project - serverless - top - ten/
- [2] Snyk Ltd., "Serverless Security: Risks in the Wild," Snyk Research Reports, 2019. [Online]. Available: https://snyk. io/blog/serverless - security - risks - in the - wild/
- [3] Cloud Security Alliance, "The Treacherous Twelve: Cloud Computing Top Threats," CSA Reports, 2016. [Online]. Available: https://cloudsecurityalliance. org/artifacts/treacherous - twelve - cloud - computing top - threats/
- [4] Amazon Web Services, "IAM Best Practices, " AWS Documentation, 2020. [Online]. Available: https: //docs. aws. amazon. com/IAM/latest/UserGuide/best practices. html
- [5] Amazon Web Services, "AWS Well Architected Framework, " AWS Whitepapers, 2020. [Online]. Available: https: //docs. aws. amazon. com/wellarchitected/latest/framework/
- [6] S. Hendrickson, B. Bahr, and S. St. Amant, "Serverless computing: One step forward, two steps back," in Proc.2016 IEEE Int. Conf. on Cloud Engineering (IC2E), pp.176–185, 2016.
- [7] M. Abadi et al., "TensorFlow: A system for large scale machine learning," in Proc.12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), 2016, pp.265–283.
- [8] N. Gruschka, M. Jensen, L. Iacono, and C. Mülle, "Security and privacy in cloud computing," Future Generation Computer Systems, vol.28, no.6, pp.1328– 1333, 2012.
- [9] R. Chandramouli and S. Rose, "Security Considerations for Microservices Architecture," NIST Special Publication 800 - 204, 2019.
- [10] T. Garfinkel and M. Rosenblum, "A virtual machine introspection based architecture for intrusion detection," in Proc. NDSS, 2003.

DOI: https://dx.doi.org/10.21275/MS2008134043

1615