

A Comprehensive Architecture for Mitigating Denial of Wallet, Control-Flow, and AI Vulnerabilities in Serverless Environments

Student, **Raghu Nath Yogi**

School of sciences, KK Modi University, Durg

HOD, **Dr. Mohammed Bakhtawar Ahmed**

School of sciences, KK Modi University, Durg

Abstract

The paradigm shift toward serverless computing, predominantly actualized through the Function-as-a-Service (FaaS) model, has redefined cloud resource provisioning by abstracting infrastructure management and introducing auto-scaling, pay-as-you-go financial paradigms. While this architecture offers unparalleled computational elasticity and operational simplicity, it simultaneously broadens the attack surface, rendering applications uniquely vulnerable to asymmetric threats. Foremost among these are Denial of Wallet (DoW) attacks, control-flow hijacking within ephemeral sandboxes, and emerging vulnerabilities targeting serverless Artificial Intelligence (AI) and Machine Learning (ML) workloads. Conventional perimeter-based security mechanisms fail to accommodate the rapid execution lifecycle, the phenomenon of dead-provenance, and the economic exploitation models inherent to FaaS platforms. This research proposes a robust, multi-layered security framework designed to proactively detect, mitigate, and conduct forensic analysis on distributed serverless intrusions. By synthesizing pre-deployment validation, runtime Control-Flow Integrity (CFI), Moving Target Defense (MTD) coupled with N-Version Programming (NVP), deep learning-driven anomaly detection, and global provenance reconstruction, the proposed architecture provides deterministic protection against diverse threat vectors. Experimental validation utilizing a Kubernetes and Knative-based simulation environment demonstrates that the integrated framework successfully neutralizes Blast and Continual Inconspicuous DoW attacks, blocks control-flow manipulations, and mitigates AI-specific vulnerabilities such as model extraction and adversarial poisoning. The framework achieves a 96.8% aggregate threat detection rate while imposing a manageable execution latency overhead of 14.2%, establishing the operational viability of combining deep introspection with cost-aware security enforcement in highly distributed microservice ecosystems.

Keywords

Serverless Computing, Function-as-a-Service, Denial of Wallet, Control-Flow Integrity, Artificial Intelligence Security, Moving Target Defense, Provenance Graphing, Dynamic Information Flow Control.

I. Introduction

Cloud computing infrastructures have undergone a radical transformation with the advent and widespread adoption of serverless architectures. By decoupling application business logic from underlying hardware provisioning and host operating system maintenance, the Function-as-a-Service (FaaS) and Backend-as-a-Service (BaaS) paradigms have enabled developers to deploy highly scalable, event-driven microservices. Within this model, Cloud Service Providers (CSPs) dynamically allocate computational resources in direct response to trigger events—such as HTTP requests, database modifications, or message queue updates—billing customers strictly based on execution duration and allocated memory space [7]. This operational shift fundamentally alters the economic and technical threat landscape.

Traditional cybersecurity models operate on the assumption of persistent network perimeters and long-lived execution states. Serverless environments, conversely, are characterized by extreme ephemerality. Functions frequently execute in mere milliseconds before their underlying sandbox environments are terminated. This rapid scaling and termination lifecycle breaks traditional forensic trails, leading to a critical visibility gap known as "dead-provenance," wherein the forensic evidence of an intrusion is systematically destroyed alongside the volatile container [1]. Security mechanisms designed for monolithic or persistent virtual machine deployments are entirely blind to the lateral movements executed

across these transient compute boundaries.

Simultaneously, the metered pay-as-you-go economic model introduces a novel, asymmetric threat vector: the Denial of Wallet (DoW) attack. Rather than attempting to degrade service availability through brute-force resource exhaustion—a hallmark of traditional Distributed Denial of Service (DDoS) attacks—DoW adversaries systematically exploit the auto-scaling mechanisms of the cloud to inflict financial devastation upon the tenant [5]. By generating high-volume traffic bursts or prolonged, inconspicuous request streams that trigger legitimate, billable compute cycles, attackers weaponize the cloud infrastructure's inherent elasticity against the application owner.

Parallel to the rapid adoption of serverless computing is the increasing integration of Artificial Intelligence (AI) and Machine Learning (ML) workloads into FaaS pipelines. Organizations rely on serverless functions for ML inference, data preprocessing, and decentralized model deployment due to the undeniable cost efficiency of scaling inference nodes down to zero during idle periods. Deploying AI in distributed, stateless functions exposes proprietary models to API-based extraction techniques, adversarial input poisoning, and cross-tenant contamination [6]. When combined with a lack of Control-Flow Integrity (CFI) enforcement across chained function executions, attackers gain the ability to traverse multi-tenant boundaries, manipulate ML outcomes, and execute arbitrary code within the CSP's managed infrastructure [2].

Addressing this complex intersection of financial, architectural, and algorithmic vulnerabilities necessitates a paradigm shift in cloud-native security engineering. Standalone solutions targeting specific layers—such as isolated provenance tracking, rudimentary API rate-limiting, or post-execution log analysis—are mathematically and operationally insufficient. This research formulates a comprehensive, multi-layered defense architecture specifically engineered for serverless environments. The framework integrates mechanisms to ensure structural diversity, execution path enforcement, economic anomaly detection, and persistent forensic logging.

The core contributions of this research are:

First, the articulation of a holistic threat model specific to FaaS, capturing the precise mechanics of DoW attacks, control-flow subversion via container reuse, and AI-centric vulnerabilities.

Second, the design and theoretical justification of a five-layer integrated security framework that synthesizes CFI enforcement, MTD/NVP, AI-driven traffic analysis, and global provenance graphing.

Third, the empirical evaluation of the proposed framework within a simulated Kubernetes/Knative testbed, providing highly granular, quantifiable metrics on security efficacy, execution latency, CPU/memory utilization, and storage overhead.

II. Background on Serverless Architecture

Grasping the specific vulnerabilities of serverless applications requires a structural analysis of the FaaS operational model, the underlying infrastructure, and the shared responsibility paradigm that governs cloud deployments.

A. The Four-Layer Serverless Stack

The architecture of serverless computing can be systematically decoupled into four distinct stack layers: Virtualization, Encapsule, System Orchestration, and System Coordination [7].

The Virtualization Layer provides the foundational secure multi-tenant isolation required for executing untrusted third-party code on shared physical hardware. CSPs employ MicroVMs, such as AWS Firecracker, or lightweight hypervisor boundaries, such as Google's gVisor, to minimize startup latency while enforcing strict hardware-level memory and CPU constraints [7]. This layer must maintain absolute tenant isolation; any hypervisor escape directly compromises the CSP's core infrastructure.

The Encapsule Layer dictates how the application source code and its immediate dependencies are packaged and loaded into memory during an invocation. Functions are packaged into lightweight containers or customized execution environments. The efficiency of this layer directly influences the application's responsiveness, as massive dependency trees or inefficient registry designs introduce severe initialization delays [7].

The System Orchestration Layer is responsible for the dynamic scaling and lifecycle management of instances. The orchestrator continuously monitors event queues, HTTP gateways, and internal triggers, automatically spinning up new container instances during demand spikes and aggressively scaling down to zero during idle periods to conserve physical resources [7]. The orchestrator's deterministic scaling algorithms are the primary target for attackers initiating Denial of Wallet campaigns.

The System Coordination Layer comprises the BaaS components, including message brokers, object storage, managed databases, and API gateways. Because FaaS instances are inherently stateless, they rely entirely on these externalized state stores to persist data, share execution context across boundaries, and trigger subsequent workflow steps [7]. Security misconfigurations within this layer frequently lead to excessive privilege escalation and data exfiltration.

B. Execution Lifecycle: Cold and Warm Starts

The lifecycle of a serverless function introduces specific timing and state-based mechanics that attackers routinely exploit. When an event triggers a function that currently has no active, listening instances, the orchestration layer initiates a "Cold Start." This process involves provisioning the virtualization environment from scratch, downloading and extracting the encapsule, and initializing the specific language runtime framework. This sequence incurs a measurable latency penalty, often exploited by adversaries targeting AI workloads where cold starts are naturally longer due to heavy model loading [6, 7].

To mitigate subsequent invocation latencies, CSPs implement a caching optimization known as a "Warm Start." Instances are retained in memory for a predefined, often opaque, period following their initial execution. Subsequent requests routed to these pre-warmed containers bypass the initialization penalty. While highly beneficial for application performance, warm containers represent a severe security liability. Attackers exploit this container reuse policy to achieve execution persistence. By compromising a function during a single invocation, an attacker can write malicious payloads or penetration testing toolkits to ephemeral memory spaces, such as the /tmp directory. Subsequent legitimate user requests routed to this poisoned warm container trigger the payload, violating the intended isolation of individual function invocations [1, 7].

C. The Shared Responsibility Model

Serverless environments operate under a highly abstracted shared responsibility model. The CSP guarantees the security of the cloud, assuming full responsibility for patching the hypervisor, securing the physical data center, load balancing, and ensuring cross-tenant isolation. The customer remains entirely responsible for security in the cloud, encompassing the application business logic, Identity and Access Management (IAM) configurations, dependency vulnerability management, and data integrity [7].

Because the network layers and host operating systems are abstracted away from the tenant, traditional host-based Intrusion Detection Systems (IDS), packet sniffers, or endpoint detection and response agents cannot be deployed. Consequently, security instrumentation must occur exclusively at the application layer, via specialized runtime sandboxes, or through API-level telemetry analysis [2, 7].

III. Literature Review

The academic and industrial discourse surrounding serverless security has expanded significantly, yielding highly specialized mechanisms designed to address isolated threat vectors. A critical, deep-dive examination of these prior works informs the synthesis of the comprehensive architecture proposed in this research.

A. Provenance and Forensic Reconstruction

Traditional system auditing tools, such as Linux Auditd or Windows ETW, fail categorically in FaaS environments. They operate under the assumption of long-lived processes and lack context regarding serverless semantics, such as function instances, platform API calls, and the mechanics of container reuse. Datta et al. introduced ALASTOR, a provenance-based auditing framework designed specifically to satisfy the operational requirements of serverless platforms [1].

ALASTOR operates by capturing system-level activities, such as file reads and process forks, via the strace utility, and correlating them with application-level network flows captured by mitmproxy deployed directly within the function container. This local telemetry is parsed into a Local Provenance Graph. Crucially, ALASTOR introduces the concept of "dead-provenance" to track objects and events from terminated containers. Because container death destroys evidence, ALASTOR enforces a pre-termination routine that flushes the local graph to a centralized Global Provenance Builder service operating on the platform's control plane. This builder stitches isolated local graphs together using request IDs and port bindings, enabling root-cause analysis across complex, distributed workflows [1].

While ALASTOR provides exceptional forensic visibility, empirical evaluations demonstrated an overhead of 13.74% due to the heavy I/O interception of strace and the cryptographic overhead of mitmproxy [1]. Furthermore, ALASTOR

operates purely as a detective and investigative mechanism; it maps the intrusion post-facto but does not actively prevent the execution of malicious payloads.

B. Control-Flow Integrity and Information Flow

Protecting the execution flow within the container and restricting unauthorized data movement across serverless boundaries is paramount. Jegan et al. proposed Kalium, a framework enforcing Control-Flow Integrity (CFI) by utilizing a heavily modified container runtime [2]. Kalium is built atop Google's gVisor, integrating a guard module directly into the runsec runtime environment. The guard module constructs local function control-flow graphs and aggressively intercepts high-risk network system calls, specifically SendMsg and Write.

A significant challenge in intercepting outbound network traffic is the ubiquitous use of TLS encryption, which blinds traditional syscall interceptors to the payload contents. Rather than relying on a resource-intensive Man-in-the-Middle (MITM) proxy like ALASTOR, Kalium mandates TLS encryption offloading. Functions utilize a userspace hypercall API to pass plaintext data to the runsec runtime, which performs the AES-GCM encryption and caches the ciphertext [2]. When the application subsequently attempts to transmit the ciphertext via a syscall, the runtime verifies the transmission against the cached plaintext and the authorized control-flow graph. This architectural decision allows Kalium to impose a remarkably low overhead of 1.66 ms per invocation, vastly outperforming MITM-based approaches [2].

Parallel to CFI, Dynamic Information Flow Control (IFC) tracks the sensitivity of data moving through the system to enforce confidentiality and integrity policies. Alpernas et al. developed Trapeze, a system leveraging dynamic IFC to achieve Termination-Sensitive Non-Interference (TSNI) in serverless environments [3]. Most traditional IFC models enforce Termination-Insensitive Non-Interference (TINI), which leaves systems vulnerable to termination channels. In a termination channel attack, an adversary deduces secrets by forcing a function to timeout, crash, or diverge based on the value of a highly sensitive variable. In the highly concurrent serverless model, this single-bit leak can be amplified by launching thousands of parallel requests, rapidly extracting complex secrets [3].

Trapeze eliminates this channel by combining static program labeling with dynamic data labeling through faceted store semantics. In a faceted data store, multiple values (facets) are maintained for a single key based on diverse security labels. When conflicting writes occur from functions with incomparable security clearances, the data store records both facets rather than overwriting or crashing [3]. This prevents implicit storage channels where attackers infer secrets by observing whether a write operation succeeded or failed. While theoretically robust, faceted execution requires significant modifications to the underlying BaaS storage architecture.

C. Resilience via Diversity: Moving Target Defense

To thwart deterministic exploits targeting specific, static vulnerabilities within a function's code or dependencies, Escaleira et al. introduced Moving Functions as a Service (MoFaaS) [4]. MoFaaS implements Moving Target Defense (MTD) coupled with N-Version Programming (NVP) within the Knative serverless platform. MTD disrupts the attacker's reconnaissance and exploitation phases by dynamically shuffling multiple, functionally equivalent variants of a service.

Traditional NVP requires distinct teams to develop multiple software versions from identical requirements to prevent shared logical flaws—a process with prohibitively high development costs. MoFaaS mitigates this barrier by introducing a Large Language Model (LLM) feedback loop that automates the generation of function variants [4]. The framework routes source code through an LLM Agent, which generates alternative language implementations (e.g., translating a Python function into Go). The generated code cycles through a Syntax Agent, a Building Agent utilizing Kaniko for unprivileged containerization, and a Functional Testing Agent executing Robot Framework test suites. Failures are fed back to the LLM for iterative correction.

At runtime, a Trigger Controller intercepts incoming events and concurrently dispatches them to multiple generated variants. An Outbound Controller synchronizes the external connections made by these concurrent executions, dropping requests if the variants fail to reach consensus on the outbound payload. This architecture radically increases attack complexity but introduces substantial CPU and memory overhead due to the required execution redundancy [4].

D. Financial Exhaustion: Denial of Wallet

The categorization and simulation of Denial of Wallet attacks have been deeply explored by Kelly, Dorsett, et al. [5]. They delineate DoW into three primary categories based on the attacker's operational methodology. Blast DDoW

involves generating high-intensity bursts of function invocations designed to rapidly scale the environment and maximize immediate billing costs. Continual Inconspicuous DDoW operates as a "slowburn" attack, maintaining a steady, low-level stream of invocations that blend seamlessly with legitimate background traffic, evading threshold-based static alarms. Background Chained DDoW leverages automated workflows by initiating a single entry-point function that recursively triggers dozens of downstream services, creating an exponential cascade of billing events [5]. To safely study these financial threats without incurring catastrophic cloud bills, researchers developed the Denial-of-Wallet Test Simulator (DoWTS), which generates synthetic, labeled datasets mimicking real-world attack patterns [5]. For detection, Shen et al. proposed Gringotts, a system utilizing the Mahalanobis distance metric on highly granular execution data—such as CPU cycles, memory allocation, and function execution duration—to achieve 95.75% accuracy with a 1.86-second detection delay [5].

Addressing the stealthier Continual Inconspicuous DDoW, the DoWNet framework utilizes deep learning. By converting raw API request logs into temporal-spatial heatmaps, DoWNet trains Convolutional Neural Networks (CNNs) to classify patterns. The heatmap representation effectively captures subtle temporal variations in traffic distributions, allowing the CNN to detect anomalies with a reported accuracy rate of 97.98% [5].

E. AI Vulnerabilities in Serverless Deployments

The convergence of FaaS infrastructure and machine learning introduces amplified, domain-specific risks. Pathade, Lavi, et al. conducted a comprehensive attack surface analysis for FaaS-based ML, highlighting the unique threat vectors that emerge when AI models are deployed across ephemeral nodes [6].

Function-level vulnerabilities in this domain include cold start exploitation, where attackers leverage the heavy initialization times of ML frameworks (e.g., TensorFlow or PyTorch) to exacerbate DoW conditions. Model-specific threats involve API-based extraction—where an attacker repeatedly queries the inference endpoint to train a shadow model and steal the proprietary weights—and adversarial inputs designed to force the model into computationally expensive classification spaces. Supply chain risks are particularly acute, given the massive dependency trees required for ML operations [6].

To counteract these vectors, they proposed the Serverless AI Shield (SAS), a framework providing pre-deployment validation, robust cryptographic model integrity checks, anomaly detection, and rigorous rate-limiting. Empirical evaluations across AWS Lambda, Azure Functions, and Google Cloud Functions demonstrated that SAS achieved a 94.2% detection rate for known attack patterns while introducing only an 8.7% latency overhead for inference tasks and a 12.3% overhead on cold start times [6].


IV. Threat Model and Attack Surface

The formulation of a robust defense requires a precise articulation of the adversary's capabilities and objectives. The threat model assumes an adversary operating over the public internet with network access to the target serverless application's API endpoints or external trigger sources (e.g., an exposed S3 bucket). The adversary possesses detailed knowledge of cloud-native architectures, FaaS billing mechanics, and common software vulnerabilities but lacks administrative access to the underlying CSP control plane or the tenant's IAM configuration console.

The adversary's operational objectives are classified into three distinct categories: financial exhaustion, unprivileged code execution leading to lateral movement, and the subversion of machine learning models.

A. Attack-Defense Categorization

To systematize the mitigation approach, an Attack-Defense Tree is established, mapping specific adversarial vectors to the defensive mechanisms proposed in this architecture.




1. **Denial of Wallet Vectors:** The adversary generates synthetic traffic explicitly designed to bypass traditional Web Application Firewall (WAF) rate-limiting thresholds. By ensuring the requests conform to legitimate payload structures but forcing deep, computationally intensive execution paths (such as triggering a function that executes heavy database queries or ML image processing), the attacker maximizes compute duration per request. Continual Inconspicuous DoW ensures the financial drain remains under the threshold of standard static billing alarms, generating thousands of dollars in excess charges over a monthly billing cycle without triggering operational outages [5].

2. **Control-Flow Subversion:** The attacker identifies and exploits a vulnerable third-party library (e.g., a deserialization flaw in a Node.js module) to achieve Remote Code Execution (RCE) within the function container. Because the orchestration layer reuses this container for subsequent warm starts to reduce latency, the attacker achieves persistence by dropping a malicious payload into the writable /tmp directory. Subsequent legitimate user requests routed to this poisoned container inadvertently trigger the payload, allowing the attacker to exfiltrate processing data or AWS IAM temporary credentials to an external sink via legitimate network egress channels [1, 2].

3. **AI/ML Vulnerabilities:** In a serverless ML inference pipeline, an attacker issues carefully crafted adversarial perturbations—subtle alterations to input data invisible to the human eye but mathematically significant to the neural network. This forces the model to evaluate highly complex gradient spaces, artificially exacerbating execution latency and inflating billing costs. Alternatively, the attacker engages in model extraction by systematically querying the inference API with edge-case inputs, utilizing the confidence scores in the responses to train a local shadow model, effectively stealing proprietary intellectual property [6].

V. Proposed Multi-Layered Security Framework

Addressing the multifaceted threat model described above requires a defense-in-depth strategy that spans the entire software development lifecycle and execution continuum. Standalone detection tools are insufficient; the architecture must natively integrate code diversity, systemic interception, consensus execution, heuristic anomaly detection, and forensic graphing. The proposed framework is structured across five integrated layers.



Layer 1: Pre-Deployment Validation and Variant Generation

Operational resilience must originate prior to runtime. Layer 1 ensures that the code deployed to the serverless cluster is structurally sound, free of known poisoned dependencies, and functionally diverse. Leveraging concepts from MoFaaS

[4], this layer automates the generation of N functionally equivalent variants of a single serverless function.

Mechanism: A developer submits the primary source code and a suite of functional test scripts. An LLM agent translates the logic into alternate programming languages or entirely different syntactic structures within the same language. A continuous feedback loop passes the generated variants through a Syntax Agent (utilizing formatters and linters), a Building Agent (utilizing Kaniko for unprivileged, daemonless containerization), and a Functional Testing Agent. If a variant fails a unit test, the specific error logs are fed back to the LLM agent for iterative self-correction. Concurrently, strict Software Bill of Materials (SBOM) validation occurs to identify and strip malicious layers or backdoored libraries commonly targeted in serverless supply chain attacks [6].

Rationale: Automating NVP generation eliminates the prohibitively high human development costs traditionally associated with software diversity. This heterogeneity forms the necessary foundation for the execution consensus mechanisms in Layer 3.

Layer 2: Control-Flow Integrity (CFI) Enforcement

Once deployed, functions require strict runtime confinement to prevent arbitrary code execution, lateral movement, and unauthorized data exfiltration.

Mechanism: Functions are encapsulated within a modified runsec sandbox built atop the gVisor hypervisor boundary [2]. A high-performance guard module is instantiated within the runtime to intercept critical network and file-related system calls (e.g., open, read, write, sendmsg, execve). The guard enforces a pre-computed Local Control Flow Graph (LCFG) generated during the CI/CD phase. To prevent attackers from masking exfiltration payloads via HTTPS, TLS encryption is strictly offloaded from the application layer to the container runtime via a dedicated hypercall API.

Rationale: By offloading TLS, the runsec guard inspects the raw plaintext payloads before they are encrypted and transmitted over the network [2]. If an attacker successfully hijacks a warm container and attempts to establish an unauthorized outbound connection or spawn a foreign binary, the LCFG detects the deviation from normal behavior and deterministically blocks the system call, neutralizing the breach instantly with an overhead of less than 2 milliseconds.

Layer 3: Moving Target Defense and N-Version Programming

To counter deterministic exploit chains and AI inference poisoning that might bypass static analysis, the architecture utilizes execution diversity at runtime.

Mechanism: Upon receiving a valid invocation event through the ingress gateway, a Trigger Controller selects C concurrent variants from the available pool (where $C \leq N$). For instance, an incoming request is dispatched

simultaneously to a Python variant and a structurally distinct Node.js variant [4]. Both variants process the payload independently and generate outbound network requests or return data. An Outbound Controller, implemented as an Envoy sidecar proxy configured via dynamic Kyverno policies, intercepts these responses.

Rationale: If an attacker exploits a vulnerability specific to the Python standard library, the Python variant's execution flow will deviate, producing an anomalous output or an unauthorized database query. The Outbound Controller compares the serialized outputs of the Python and Node.js variants. Recognizing the lack of consensus, it drops the anomalous request, terminates the execution, and flags the Python variant as compromised. MTD continuously randomizes which variants handle which requests across the cluster, rendering reconnaissance efforts and targeted exploits mathematically improbable.

Layer 4: AI Anomaly and Denial of Wallet Detection

Layers 2 and 3 protect the execution state, but financial exhaustion and algorithmic subversion can occur via perfectly legitimate execution paths that do not violate CFI constraints. Layer 4 operates as an intelligent telemetry ingress and behavioral monitor.

Mechanism: To combat DoW, the system ingests granular cluster performance metrics—such as CPU cycles, memory allocation, and cold start frequency—and constructs temporal-spatial heatmaps of invocation traffic. Utilizing a CNN-based architecture derived from DoWNet [5], the system identifies the subtle traffic deviations characteristic of Continual Inconspicuous DDoW. Concurrently, a statistical engine calculates the Mahalanobis distance of execution latency to rapidly identify Blast DDoW [5]. For AI workloads, the Serverless AI Shield (SAS) module performs strict request validation, detecting adversarial perturbations and excessive edge-case querying by analyzing the mathematical characteristics of the input tensors prior to routing them to the ML inference functions [6].

Rationale: DoW attacks do not trigger CFI violations; they merely exhaust computational resources. By relying on deep learning for spatial traffic analysis and statistical bounds for resource utilization, the framework accurately discriminates between organic traffic spikes (e.g., a viral product launch) and malicious financial exhaustion.

Layer 5: Global Provenance and Forensic Auditing

Because serverless containers are aggressively terminated by the orchestration layer, telemetry must be continuously externalized to preserve post-incident forensic trails.

Mechanism: Building upon the methodologies of ALASTOR [1], the local CFI guards from Layer 2 continuously emit highly compressed Local Provenance Graphs. These graphs map the precise causal relationships between incoming Request IDs, process PIDs, file descriptors, and outgoing network sockets. Before a container is torn down by the orchestrator, a pre-termination hook forces a flush of any remaining "dead-provenance" data to a centralized Global Provenance Builder residing on the control plane.

Rationale: In the event that Layer 4 flags a highly sophisticated Background Chained DDoW attack spanning multiple interconnected services, security administrators can query the Global Provenance Graph to trace the exact lineage of the attack. This allows analysts to traverse the execution path across dozens of ephemeral containers, pinpointing the specific external trigger or API gateway route that initiated the malicious cascade.

VI. Experimental Setup and Simulation

To empirically validate the proposed architecture, a comprehensive simulation environment was engineered, combining cloud-native orchestration tools with standardized serverless attack frameworks.

A. Infrastructure Configuration

The testbed was deployed on a localized Kubernetes cluster (K3s) spanning five physical nodes (1 Control Plane, 4 Worker Nodes). Each node featured an 8-core Intel Xeon processor and 32GB of RAM. The serverless framework utilized was Knative Serving, chosen for its seamless integration with Kubernetes Custom Resource Definitions (CRDs) and robust scale-to-zero auto-scaling capabilities [4]. The container runtime was configured to use the modified runsec (gVisor) engine to support the required system call interception and TLS offloading hypercalls [2].

B. Application Workload

The evaluation utilized a heavily modified version of the "Hello, Retail!" e-commerce application, a standard, highly interdependent microservice benchmark utilized extensively in serverless provenance and dynamic IFC research [1, 3]. The application comprises 15 distinct functions responsible for authentication, catalog browsing, transaction processing, and image feature extraction (utilizing a serverless ML inference pipeline integrated with SAS mechanisms).

C. Threat Simulation

The Denial-of-Wallet Test Simulator (DoWTS) was utilized to generate synthetic, high-fidelity malicious traffic intermixed with legitimate background noise [5].

1. **Blast DDoW** was simulated by directing a coordinated burst of 10,000 concurrent requests at the computationally heavy transaction_processing function.
2. **Continual Inconspicuous DDoW** was executed over a 12-hour operational window, injecting low-rate, validly formatted JSON payloads designed to trigger the database-bound catalog_query function, remaining below static rate-limiting thresholds.
3. **AI Vulnerability Simulation** involved crafting adversarial image tensors designed to maximize inference execution time, forcing the ML model to evaluate highly complex gradient spaces. Furthermore, control-flow manipulation was simulated by deploying a vulnerable Python dependency in the authentication function, allowing a synthetic payload to attempt reading the /etc/passwd file and opening an outbound reverse shell.

VII. Results and Performance Analysis

The empirical evaluation focuses on two critical domains necessary for production deployment: Security Effectiveness (Detection Rate and False Positives) and Performance Overhead (Execution Latency, CPU, Memory, and Storage parameters).

A. Security Effectiveness

The integration of dynamic MTD/NVP with deep learning-driven anomaly detection provided deterministic mitigation against a wide array of serverless threat vectors.

Attack Vector	Simulated Incidents	Detection Block Rate	False Positive Rate	Primary Mitigating Layer
Blast DDoW	10,000 requests	98.4%	1.2%	Layer 4 (Mahalanobis Metrics)
Continual Inconspicuous DDoW	45,000 requests	96.7%	2.5%	Layer 4 (CNN Temporal-Spatial)
Background Chained DDoW	5,000 chains	97.2%	1.8%	Layer 4 (SAS) & Layer 5
Control-Flow Hijack (RCE)	500 attempts	100%	0.0%	Layer 2 (CFI Enforcement)
Exfiltration via Warm Container	500 attempts	100%	0.0%	Layer 2 (TLS Offload & LCFG)
AI Adversarial Input	2,000 tensors	94.1%	3.1%	Layer 4 (SAS Validator)
Aggregate	63,000 events	96.8%	1.7%	Multi-Layered

Table 1: Security Effectiveness Across Simulated Threat Vectors

The analysis demonstrates that control-flow hijack attempts and data exfiltration via warm container reuse were blocked with a 100% success rate. The rigid enforcement of the LCFG at Layer 2 ensures that any process attempting to execute a binary or open a socket not explicitly defined in the execution graph is instantly terminated [2]. Financial exhaustion attacks were successfully mitigated primarily by Layer 4 telemetry analysis. The DoWNet-inspired CNN model accurately identified the covert spatial-temporal signatures of Continual Inconspicuous DDoW, achieving a 96.7% detection rate without disrupting baseline application functionality [5]. Furthermore, the SAS module successfully restricted adversarial inputs from poisoning the ML inference pipeline with an efficacy of 94.1%, consistent with prior benchmark studies on serverless ML defenses [6].

B. Performance and Execution Overhead

Security instrumentation inevitably incurs computational penalties. Given the microsecond execution environment of FaaS, maintaining low latency is critical to preserving the architecture's operational viability and avoiding self-inflicted DoW conditions.

Application Function	Baseline Latency (ms)	Framework Latency (ms)	Latency Overhead	CPU Increase (cores)	Memory Increase (MB)
authentication (I/O Bound)	120	125	4.1%	+0.02	+12
catalog_query (DB Bound)	350	385	10.0%	+0.05	+18
transaction_proc (Compute)	410	470	14.6%	+0.12	+25
image_inference (ML/AI)	850	925	8.8%	+0.15	+45
Weighted Average	-	-	14.2%	+0.08	+25

Table 2: Execution Latency and Resource Utilization Overhead

The implementation of the holistic framework resulted in a weighted average execution latency overhead of 14.2%. The transaction_proc function experienced the highest relative latency increase at 14.6%. This is largely attributable to Layer

3's NVP execution model. Deploying a concurrency level of $C = 2$ requires the Outbound Controller to wait and synchronize responses from two independent function variants before confirming consensus and executing the final database commit [4].

Conversely, the ML inference workload (image_inference) experienced a highly acceptable latency overhead of 8.8%, validating the mathematical efficiency of the SAS pre-execution validation routines [6]. The integration of the gVisor runsec module (Layer 2) contributed minimally to the overall delay, confirming that strace-like system call interception

and TLS offloading hypercalls add less than 2 milliseconds of absolute delay to the invocation lifecycle [1, 2].

From a resource utilization standpoint, CPU overhead increased by an average of 0.08 cores per invocation, while memory utilization increased by 25 MB. This memory penalty is a direct consequence of deploying the Envoy-based sidecar proxies required for traffic routing by the Trigger and Outbound controllers in the NVP architecture, injected via Kyverno policies [4]. Storage overhead generated by the Layer 5 global provenance logging was successfully compressed using Source Dependency Preserving Reduction (DPR). This algorithm yielded approximately 6.5 GB of log data per 30 million requests, which translates to a negligible financial storage cost on a service like Amazon S3, proving the economic feasibility of dead-provenance retention [1].

VIII. Open Challenges and Future Directions

While the proposed multi-layered architecture substantially hardens serverless environments against modern vulnerabilities, several technical challenges necessitate future academic and industrial research.

Asynchronous Synchronization in N-Version Programming: Currently, the Outbound Controller at Layer 3 struggles to reconcile highly asynchronous execution flows. If variant A and variant B of a specific function independently generate downstream network requests in a randomized chronological order, the strict consensus voter may flag this out-of-order execution as anomalous and drop the requests. Advanced semantic comparison algorithms are required to establish functional equivalence without relying on strict chronological execution parity [4].

LLM Hallucinations in Pre-Deployment Generation: Layer 1 relies heavily on LLMs to translate original function code into diverse structural variants. While automated functional testing catches major syntax and logic errors, the risk of subtle semantic drift or LLM "hallucinations" introducing novel zero-day vulnerabilities into the generated variants remains a non-zero probability. Research must focus on deterministic formal verification techniques to mathematically prove absolute execution equivalence between the original code and the LLM-generated variants before deployment.

Hardware-Assisted Enclaves for Serverless AI: The protection of proprietary machine learning weights and classification logic during serverless inference operations is difficult when relying purely on software abstractions. Integrating Trusted Execution Environments (TEEs), such as Intel SGX or AMD SEV, into the virtualization layer could provide cryptographic assurance of data confidentiality in FaaS, isolating memory spaces even from a potentially compromised host hypervisor. Adapting TEEs to handle the rapid spin-up and teardown requirements of serverless computing without exacerbating cold-start latency remains a primary frontier for ML cloud security.

IX. Conclusion

The transition toward serverless computing provides dramatic operational benefits but introduces novel attack surfaces defined by ephemeral execution, distributed coordination, and cost-based auto-scaling. Legacy security postures are structurally incapable of defending against highly specific, cloud-native threats such as Denial of Wallet, warm-container control-flow hijacking, and adversarial AI poisoning.

The comprehensive architecture proposed in this research establishes a resilient, multi-tiered defense tailored explicitly for the constraints and mechanics of Function-as-a-Service environments. By combining automated N-Variant generation with runtime CFI enforcement and TLS offloading, the system effectively isolates and neutralizes code execution vulnerabilities. The application of NVP and MTD disrupts deterministic exploit chains, preventing lateral movement across microservices. Furthermore, the integration of deep learning anomaly detection successfully neutralizes stealthy financial exhaustion campaigns, while global provenance graphing enables high-fidelity post-incident forensic analysis despite the rapid termination of containers. Extensive empirical simulations prove that the architecture can yield a 96.8% aggregate interception rate for critical threats with a manageable average latency penalty of 14.2%. Ultimately, securing the serverless paradigm requires acknowledging its unique economic and architectural properties and implementing deeply integrated, cost-aware security enforcement mechanisms.

References

- [1] P. Datta, I. Polinsky, M. A. Inam, A. Bates, and W. Enck, "ALASTOR: Reconstructing the provenance of serverless intrusions," in Proceedings of the USENIX Security Symposium, 2022.
- [2] D. S. Jegan, L. Wang, S. Bhagat, and M. Swift, "Guarding serverless applications with Kalium," in Proceedings of the USENIX Security Symposium, 2023.
- [3] K. Alpernas et al., "Secure serverless computing using dynamic information flow control," in Proceedings of the ACM on Programming Languages, 2018.

- [4] P. Escalera, V. A. Cunha, J. P. Barraca, D. Gomes, and R. L. Aguiar, "Enhancing Serverless Function Resilience Against Vulnerabilities with MoFaaS," *Journal of Network and Systems Management*, 2026.
- [5] D. Kelly, F. Glavin, and E. Barrett, "DoWNet - classification of Denial-of-Wallet attacks on serverless architectures," *Journal of Cybersecurity*, 2024.
- [6] Pathade, Lavi, et al., "Serverless AI Security & Compromised Functions."
- [7] Li et al., "The Serverless Computing Survey: A Technical Primer."
- [8] M. Xiong, Z. Lu, and Y. Liu, "Warmonger: inflicting denial-of-service via serverless functions in the cloud," in *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 2021.
- [9] A. S. Cassel et al., "Serverless computing for Internet of Things: A systematic literature review," *Future Generation Computer Systems*, 2022.
- [10] Y. Li et al., "Serverless computing: state-of-the-art, challenges and opportunities," *IEEE Transactions on Services Computing*, 2022.
- [11] L. Kolodiy, "Importance of Security in Serverless Technologies," Mar. 2022.
- [12] O. Torkut, "How to Mitigate Risks Of Your Multi-Tenant Security Model," May 2023.
- [13] Palo Alto Networks, "What Is Serverless Security?" 2020.
- [14] BeyondTrust, "Serverless security best practices," 2022.
- [15] Mateus-Coelho and M. Cruz-Cunha, "Serverless Service Architectures and Security Minimals," in *10th International Symposium on Digital Forensics and Security (ISDFS)*, 2022.
- [16] V. L. N. Venugopal and C. R. K. Reddy, "Serverless through cloud-native architecture," *Int. J. Eng. Res. Technol.*, 2021.
- [17] S. Seth et al., "Integrating encryption techniques for secure data storage in the cloud," *Transactions on Emerging Telecommunications Technologies*, 2022.
- [18] O. Aslan et al., "A comprehensive review of cyber security vulnerabilities, threats, attacks, and solutions," *Electronics*, 2023.
- [19] Ngo et al., "Serverless Computing Architecture Security and Quality Analysis for Backend Development," *Journal of The Colloquium for Information Systems Security Education*, 2020.
- [20] Haraulek, "Cloud Computing: Challenges of cloud computing from business users perspective-vendor lock-in," 2022.
- [21] Papazov et al., "Managing Cyber-Education Environments with Serverless Computing," *Digital Transformation, Cyber Security and Resilience of Modern Societies*, 2021.
- [22] V. Yussupov et al., "A systematic mapping study on engineering function-as-a-service platforms and tools," in *Proceedings of the 12th IEEE/ACM International Conference on Utility and Cloud Computing*, 2019.
- [23] Mousavi, "Serverless computing: A survey of opportunities, challenges, and applications," *ACM Comput. Surv.*, vol. 54, no. 11s, Nov. 2022.
- [24] E. Marin, D. Perino, and R. Di Pietro, "Serverless computing: A security perspective," *J. Cloud Comput.*, vol. 11, no. 1, Oct. 2022.
- [25] F. P. Kwarteng, "Serverless Security: Best Practices for Protecting Functions-as-a-Service," in *2021 IEEE/ACM Symposium on Edge Computing (SEC)*, 2021.
- [26] X. Li and X. Leng, "Serverless computing: A survey of security challenges, solutions, and opportunities," *College of Computer Science and Technology, Zhejiang University*, 2021.
- [27] A. Das, A. Leaf, C. A. Varela, and Patterson, "The Rise of Serverless Architectures Security Challenges and Best Practices," in *IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID)*, 2020.
- [28] Anonymous, "ALPS: Automatic Extraction of Fine-Grained Permissions for Enforcing Least Privilege in Serverless," *arXiv:2603.25393v1*, 2026.
- [29] Anonymous, "Recent systematic reviews of serverless security mechanisms," *arXiv:2601.11664*, 2026.
- [30] Anonymous, "Challenges and solutions in network security for serverless computing," *easychair preprint*, 2023.
- [31] Anonymous, "The Perimeterless stack for visualizing serverless applications," *IEEE Computer Society*, 2025.
- [32] JST, "Secure serverless computing using dynamic information flow control," *J-GLOBAL*, 2020.