ISSN: 2581-3404 (Online) IF: 5.86 (SJIF)
IJIRTM, Volume-10, Issue-3, May-2026

Detection and Prevention of Distributed Denial of Service (DDoS) Attacks in Cloud environments: A Comprehensive Survey

Rajeev Ranjan[1], Dr. Darshana Rai[2], Dr. Chetan Agrawal[3]
M. Tech. Scholar[1], Associate Professor[2], Head & Associate Professor[3]
Dept. of CSE, RITS, Bhopal, India[1,2,3]
ranjanrajeev5378@icloud.com[1], darshana.rai@gmail.com[2], Chetan.agrawal12@gmail.com[3]

Abstract. *Cloud computing is the dominant paradigm for scalable IT services, but its openness and resource-sharing approach make it a target for DDoS attacks. This paper is a comprehensive survey of DDoS attacks targeting cloud infrastructures, the evolution of detection and prevention mechanisms, the role of machine learning and deep learning in automated threat mitigation, and the growing importance of SDN as a defence platform. We categorise 60+ research studies from 2018 to 2025 by approach and effectiveness and critically evaluate their strengths and weaknesses. Our assessment provides a uniform taxonomy, comparison tables, and architectural diagrams to help researchers and practitioners navigate this complicated world. Federated learning, explainable AI, and zero-trust cloud DDoS protection architectures are our final research problems and intriguing future prospects.*

Keywords: DDoS attacks, cloud computing, intrusion detection, machine learning, deep learning, SDN, botnet, mitigation, network security, traffic classification.

Introduction

Cloud computing has changed how companies provision, deploy, and use IT resources. AWS, Azure, and GCP support key infrastructure in banking, healthcare, e-commerce, and government. This move has enabled unprecedented scalability and cost effectiveness, but it has also increased the attack surface for adversaries, with DDoS attacks being one of the most persistent and devastating cloud availability threats. DDoS attacks are coordinated by a botnet of infected systems to overwhelm a target's computational resources, network bandwidth, or application-layer services, denying legitimate users access. DDoS assaults use thousands or millions of geographically spread computers, making source attribution and traffic filtering more difficult. Threats have increased considerably in recent years. Akamai stopped a record-breaking 900.1 Gbps attack on an Asia-Pacific customer in February 2023. Google Cloud and Cloudflare logged the greatest application-layer DDoS attack ever in October 2023, a zero-day HTTP/2 Rapid Reset campaign that hit over 398 million requests per second. These milestones highlight a worrying trend: attacks are increasing in volume and sophistication, using multi-vector methods with more than 14 simultaneous attack channels to overwhelm automated defences and security operations teams.

DDoS issues in cloud systems differ greatly from on-premises networks. Economic denial of sustainability (EDoS) attacks use on-demand flexibility to exhaust a victim's billing allowance. Cross-tenant attack propagation and SLA violations are concerns in multi-tenancy, when multiple customers share physical infrastructure. IP-reputation and packet-inspection defences are also complicated by cloud

orchestration's dynamic IP addressing and quick workload mobility. These difficulties have prompted research in anomaly detection, machine learning (ML), deep learning (DL), SDN-based defences, and hybrid architectures. The sector is continuously expanding, and practitioners seek to find the best ways for their cloud deployment models (IaaS, PaaS, SaaS). This survey provides an organised, critical, and current state-of-the-art review to close that gap.

Motivation and Scope

This survey is driven by the lack of a recent review that covers cloud DDoS attack taxonomy, new ML/DL detection pipelines, SDN-based mitigation, and 2022–2025 research contributions. Previous surveys either focused on one defensive mechanism or did not account for the significant improvements in generative AI, federated learning, and LLM-assisted rule generation that are changing the field. The taxonomy and characteristics of cloud DDoS attacks, a structured literature review of detection mechanisms, prevention and mitigation strategies, comparative analysis of techniques, datasets, and benchmark results, and open research challenges and future directions are covered in this survey.

Paper Organization

This paper continues as follows. The cloud computing backdrop and DDoS risks are covered in Section 2. Section 3 lists DDoS attack types. The detection literature review is in Section 4. Section 5 examines preventive and mitigation. Section 6 evaluates methods across major dimensions. Section 7 addresses open issues for future research. Section 8 ends the paper.

Cloud Computing and DDoS: Background

2.1 Cloud Computing Service and Deployment Models

NIST defines cloud computing as on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service delivered through three service models: IaaS, PaaS, and SaaS. DDoS attackers see distinct attack surfaces with each model.

Table 1: Cloud Service Models and DDoS Exposure

Service Model	Description	DDoS Exposure	Example Providers
IaaS	Virtual machines, storage, networking	Network-layer volumetric attacks	AWS EC2, Azure VMs, GCP Compute
PaaS	Runtime, middleware, development tools	Application-layer and API attacks	Heroku, Google App Engine
SaaS	Complete software delivered via browser	HTTP floods, Slowloris, credential stuffing	Gmail, Salesforce, Office 365
FaaS/Serverless	Event-driven functions on shared runtime	EDoS, function exhaustion attacks	AWS Lambda, Azure Functions

2.2 Why Cloud Environments are Prime DDoS Targets

Cloud environments are appealing DDoS targets for various reasons. First, shared physical resources might cause a noisy-neighbor effect when one renter floods another's virtual network interface. Second, the auto-scaling mechanism, designed to handle legitimate traffic spikes, can accidentally absorb attack traffic at the customer's expense, causing EDoS. Third, cloud control planes—resource provisioning management APIs—are high-value targets that can bring regions offline. Fourth, the geographic distribution of cloud points of presence (PoPs) makes centralized traffic scrubbing architecturally complex.

2.3 DDoS Attack Lifecycle in Cloud

Knowing the DDoS lifecycle helps create timely defences. A typical cloud-targeted DDoS attack has five phases: (1) Reconnaissance, where the attacker uses scanning tools to identify target IP ranges, open ports, and application endpoints; (2) Weaponization and delivery, where malware propagation or DaaS platforms are used to build a botnet; and (3) Command and Control (C2). (4) Flooding the victim with traffic; and (5) Evasion, shifting attack vectors to bypass rate-limiting and signature-based filters.

Figure 1: DDoS Attack Lifecycle in Cloud Environment

[Attacker]

▼ Reconnaissance & Scanning
▼ Botnet Recruitment (Malware / DaaS)
▼ C2 Channel Establishment
▼ Coordinated Flood Launch
▼ Evasion & Vector Rotation

[Cloud Victim: Bandwidth / CPU / Memory Exhaustion → Service Unavailability]

Figure 1: Five-phase DDoS attack lifecycle as it unfolds against a cloud target

Taxonomy of DDOS Attacks in Cloud Environments

A rigorous DDoS assault taxonomy is needed to create defences that target the right threat vectors. We categorise cloud-targeted DDoS assaults by network layer, attack tactic, traffic rate, and sophistication.

3.1 Layer-Based Classification

DDoS attacks are network/transport-layer (Layers 3 and 4) and application-layer (Layer 7) assaults in the OSI paradigm. Volumetric assaults like UDP and ICMP floods overwhelm bandwidth. SYN floods and Ping-of-Death exploit TCP/IP stack implementation flaws. Application-layer attacks like HTTP GET/POST floods and DNS amplification swamp web servers and resolvers with expensive requests.

Table 2: DDoS Attack Types — Layer-Based Classification

OSI Layer	Attack Type	Mechanism	Impact	Example Tools
Layer 3	ICMP Flood	Overwhelm with ping packets	Bandwidth exhaustion	hping3, Nping
Layer 3	IP Fragmentation	Malformed fragments crash stack	CPU/memory drain	Scapy

Layer 4	SYN Flood	Half-open TCP connections exhaust server tables	Connection table overflow	hping3, LOIC
Layer 4	UDP Flood	Random UDP packets overwhelm bandwidth	Bandwidth / CPU saturation	UDP Unicorn, LOIC
Layer 7	HTTP GET Flood	Legitimate-looking HTTP requests exhaust server	Web server crash	HOIC, Siege
Layer 7	Slowloris	Slowly send incomplete HTTP headers	Thread exhaustion	Slowloris
Layer 7	DNS Amplification	Spoofed queries to open resolvers	Bandwidth amplification 28–54x	Custom scripts
Layer 7	NTP Amplification	monlist command abuse on NTP servers	Amplification up to 556x	Custom scripts

3.2 Attack Strategy Classification

DDoS assaults can be classed by strategy beyond layer. Direct assaults fire malicious packets from bot to target. Reflection attacks redirect innocent third-party responses using the victim's IP. Amplification uses protocols with huge responses to little requests. Low-rate pulsing assaults degrade service quality while evading threshold-based detection. Multi-vector attacks simultaneously use multiple of the above tactics, complicating defences.

3.3 Rate-Based Classification

High-rate DDoS attacks generate traffic quantities that are easily different from baseline, making detection easy but mitigating computationally demanding. Low-rate DDoS assaults are more devious; they mimic genuine traffic patterns and stay below rate-limiting levels, generating TCP timeout and retransmission loops that degrade service without volumetric alarms. Section 4 discusses the scientific problem of detecting low-rate cloud attacks.

Figure 2: DDoS Attack Taxonomy in Cloud Environments

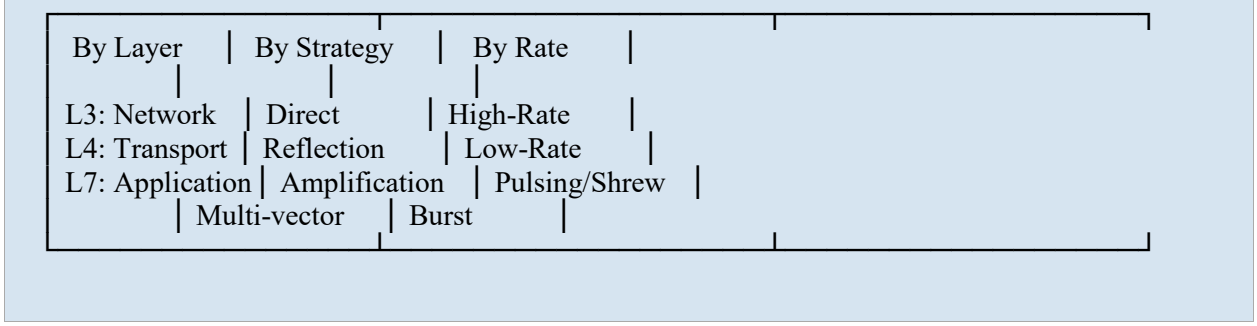

Figure 2: Taxonomy of DDoS attacks categorised by OSI layer, strategy, and traffic rate

Literature review: Detection Mechanisms

A comprehensive review of literature-proposed DDoS detection strategies follows. The review is divided into four categories: threshold-based and statistical, machine learning, deep learning, and SDN-assisted detection.

4.1 Threshold-Based and Statistical Detection

Monitoring network statistics and alarming when values exceed predefined criteria is the earliest and most computationally lightweight DDoS detection method. Entropy-based approaches measure source IP, packet size, and inter-arrival time randomness. Attack traffic targets fewer sources and destinations than genuine traffic, hence a large drop in entropy indicates a flood attack.

Zhang and Wu used SVM and an entropy-based approach to detect Android-targeted infrastructure DDoS source IP addresses. In simulated environments, Khosroshahi and Ozdemir proposed a dual-algorithm system to distinguish TCP and HTTP flood attacks. Algorithm 1 checks PSH and ACK flags for TCP floods, while Algorithm 2 checks request counts per source IP against a time-frame. By adding Random Forest (RF) classifiers, Shaaban and Abdelwanes enhanced precision on mixed-type traffic datasets.

Despite their low processing cost, statistical approaches are reactive. Threshold numbers must be tuned per deployment environment and fail to stop progressively increasing low-rate attacks that never alert. These constraints prompted the move to machine learning.

4.2 Machine Learning-Based Detection

Machine learning (ML) DDoS detection algorithms identify network flow patterns as benign or attack-type-specific. The literature explores several supervised, unsupervised, and semi-supervised methods.

4.2.1 Supervised Learning

Supervised classifiers including Decision Trees (DT), Random Forests (RF), k-Nearest Neighbours (k-NN), and Gradient Boosting Machines (GBM) have been widely used on benchmark datasets like CICIDS-2017, CICIDS-2018, and NSL-KDD These datasets comprise labelled network traffic with realistic DDoS, DoS, brute force, and port scanning attacks. An ensemble online ML model for DDoS detection in SDN systems by Jain et al. achieved detection accuracy above 99% on the CICIDS-2018 dataset with sub-millisecond classification latency for real-time deployment. Using Gradient Boosted Decision Trees and feature selection, Bhayo et al. created an ML-based framework for SD-IoT networks that reduces model complexity without sacrificing minority attack class recall.

4.2.2 Unsupervised and Semi-Supervised Learning

Labelled DDoS traffic data is expensive and quickly obsolete as attack vectors change. Unsupervised methods like k-means clustering, autoencoders, and isolation Forests can spot unusual traffic without labels. Semi-supervised models work in cloud deployments where labelling is impractical due to a limited labelled set and massive unlabelled data.

Table 3: Comparative Summary of ML-Based DDoS Detection Approaches

Reference	Algorithm	Dataset	Accuracy	Environment	Year
Jain et al.	Ensemble Online ML	CICIDS-2018	99.2%	SDN/Cloud	2024
Bhayo et al.	Gradient Boosting + FS	CICIDS-2017	98.7%	SD-IoT	2023

Wang & Li	Random Forest	CICIDS-2018	98.1%	SDN	2024
Singh & Gupta	SVM + Entropy	NSL-KDD	96.4%	General Cloud	2022
Aladaileh et al.	Entropy-based approach	Custom SDN	97.8%	SDN Controller	2023
Kumari & Jain	PCA + Decision Tree	CICIDS-2017	97.1%	IoT/Cloud	2024
Shaaban et al.	Random Forest	Mixed	95.9%	HTTP/TCP Focus	2022

4.3 Deep Learning-Based Detection

Deep neural networks can detect complicated spatial and temporal attack patterns in raw packet streams and high-dimensional flow feature vectors that shallow classifiers cannot. CNN, LSTM, GRU, Autoencoders, and hybrid CNN-LSTM models are the main DDoS detection architectures. CNNs analyse network traffic as a two-dimensional matrix of flow features and use convolutional filters to find attack patterns. LSTM and GRU networks are good at identifying pulsating and low-rate DDoS attacks with evolving signatures because they simulate traffic sequence temporal dependencies. Convolutional layers collect spatial characteristics before feeding them into recurrent layers for temporal modelling in hybrid CNN-LSTM architectures, which outperform pure CNN or LSTM models on benchmarkdatasets.

GANs have been used for adversarial robustness testing, generating realistic synthetic attack traffic to stress-test classifiers, and data augmentation to address the class imbalance between normal and attack samples in most public DDoS datasets. DDoS detection literature (2024–2025) uses transformer-based architectures inspired by big language model accomplishments to capture long-range traffic flow dependencies via self-attention.

4.4 SDN-Assisted Detection

Software-Defined The control plane and data plane are separated in networking, allowing centralised network topology visibility and programmable flow management. This design is powerful for DDoS defence because the SDN controller can see traffic statistics from all switches concurrently, providing network-wide anomaly correlation that dispersed per-device defences cannot.

Han et al. presented OverWatch, a cross-plane DDoS defence architecture that distributes detection intelligence between the SDN data plane and control plane to reduce latency and maintain accuracy. Musa et al. showed that SDN-based defence systems may isolate attack traffic into flow tables and drop malicious flows within milliseconds. Reinforcement learning-based dynamic flow rule generation by Su et al. modifies SDN mitigation rules in real time as attack vectors rotate.

Jain, Shukla, and Goel's 2024 survey categorises 90+ papers across detection algorithms, controller architectures, and mitigation mechanisms and identifies controller scalability and northbound API security as the biggest open challenges in SDN-based cloud DDoS defence.

Prevention and Mitigation Strategies

5.1 Prevention Mechanisms

Preventing DDoS attacks before they reach cloud infrastructure is the goal. Network-level access control lists (ACLs) that limit traffic from known malicious IP ranges, rate limiting at cloud provider edge routers, and anycast routing, which distributes attack traffic across multiple PoPs rather than concentrating it at a single endpoint, are key strategies.

Cloud providers like AWS Shield Advanced, Azure DDoS Protection Standard, and Google Cloud Armour offer transparent always-on traffic monitoring and automatic mitigation at the network edge for DDoS protection. These services use the cloud provider's vast bandwidth and global scrubbing infrastructure to absorb volumetric attacks before they reach the customer's virtual network.

5.2 Detection-Triggered Mitigation

Once an attack is detected, mitigation techniques must restore service availability rapidly while minimising false positives that prevent legitimate users. Traffic scrubbing, rate limiting, CAPTCHA and JavaScript challenges for HTTP-layer attacks to distinguish bots from humans, and black-hole routing are the most common mitigation primitives.

Table 4: DDoS Defense Strategies — Prevention, Detection, and Mitigation

Strategy	Mechanism	Effectiveness	Limitations
Ingress Filtering	BCP38, uRPF — discard spoofed-source packets	High vs. reflection attacks	Requires ISP cooperation; ineffective vs. non-spoofed botnets
Rate Limiting	Token bucket, leaky bucket at edge	Moderate — reduces impact	Blunt; blocks legitimate bursts
Traffic Scrubbing	Anycast + scrubbing centers (AWS Shield, Akamai Prolexic)	High for volumetric attacks	Latency overhead; cost at scale
SDN Flow Rules	Dynamic OpenFlow rules block attack flows	High — sub-second response	Controller bottleneck; scalability
ML/DL Detection	Real-time classification of traffic flows	High accuracy (>97%)	Adversarial evasion; training data staleness
CAPTCHA / JS Challenge	Browser-level challenge for HTTP flood	Effective vs bots; poor UX	Headless browsers can solve CAPTCHAs
Black-hole Routing	Null-route victim IP — absorbs attack	Eliminates attack traffic	Full DoS for legitimate users
LLM Firewall Rules	AI generates dynamic ACL/firewall rules	Emerging — promising results 2024–2025	Hallucination risk; verification required

5.3 Emerging Mitigation Approaches

5.3.1 Federated Learning for Distributed Defense

Federated learning (FL) lets various cloud tenants or regional data centers train a DDoS detection model without sharing raw traffic data, overcoming privacy problems that impede centralised data aggregation. Each participant trains a local model on its traffic data and shares only gradients or parameters with a central aggregator. While maintaining data locality, the aggregated global model benefits from participant visibility. Early results demonstrate that FL-based DDoS detectors may detect as well as centralised models with much less communication overhead.

5.3.2 LLM-Assisted Firewall Rule Generation

Recent studies by Louro, Wang, and Yin (2024) show that big language models may automatically construct network firewall rules in response to DDoS signatures. LLMs produce syntactically valid ACL rules for common firewall platforms (iptables, Palo Alto, Cisco ASA) with accuracy sufficient for automated deployment after lightweight validation from a structured description of an ongoing attack (source port distribution, packet size histogram, flag patterns). The quality of defences has changed from hand-crafted rule libraries to adaptive AI.

5.3.3 Blockchain-Based Collaborative Defense

Latah and Kalkan suggested using blockchain and SDN to establish distributed, tamper-proof DDoS source IP blacklists. Blockchain's decentralisation precludes a single point of failure in defence systems and rewards threat intelligence with tokens. Abdulqadder et al. included the DAG blockchain for edge-assisted honeypot operation and multi-controller load balancing in 5G SDN setups.

Comparative Analysis

6.1 Benchmark Datasets

The quality and representativeness of training and evaluation datasets are critical factors in assessing DDoS detection research. We surveyed the most widely used public DDoS datasets and identify their key characteristics below.

Table 5: Widely Used DDoS Benchmark Datasets

Dataset	Year	Attack Types	Size	Notable Features / Limitations
DARPA 1999	1999	DoS, R2L, U2R	~5 M records	Simulated; outdated traffic patterns; no cloud scenarios
NSL-KDD	2009	DoS, Probe, R2L, U2R	~125 K records	Reduced DARPA; removes duplicates; still outdated
CICIDS-2017	2017	DDoS, DoS, Brute Force, Web	~2.8 M flows	Realistic B-profile traffic; widely cited; aging
CICIDS-2018	2018	DDoS, Infiltration, BoT	~16 M flows	Large scale; reflects modern attack vectors
CAIDA	Various	UDP/ICMP	Varies	Real internet backbone captures;

		Floods		restricted access
BCCC-Cloud-DDoS-2024	2024	17 DDoS scenarios	300+ features	Cloud-native; 8 benign profiles; most recent and realistic

The BCCC-cPacket-Cloud-DDoS-2024 dataset addresses 15 issues in previous datasets, including the lack of actual benign user behaviour profiles, cloud-specific traffic characteristics, and attack scenario diversity. The most cloud-relevant benchmark has 17 DDoS attack scenarios and approximately 300 extracted network/transport layer properties.

6.2 Performance Comparison of Detection Approaches

Table 6: Detection Performance Comparison Across Approaches

Approach	Method	Accuracy	F1-Score	Latency	Deployment Fit
Statistical/Entropy	Threshold-based	~85–92%	~0.83	Very Low	Edge / IoT
Classical ML (RF, SVM)	Supervised	95–99%	~0.97	Low	Cloud IaaS
CNN	Deep Learning	97–99%	~0.98	Medium	Cloud IaaS/PaaS
LSTM / GRU	Deep Learning (temporal)	97–99%	~0.98	Medium-High	Stateful traffic analysis
CNN-LSTM Hybrid	Hybrid DL	98–99.5%	~0.99	High	High-value cloud targets
SDN + ML	Network + ML	98–99%	~0.98	Very Low (<1 ms)	Cloud with SDN controller
Federated Learning	Collaborative DL	96–98%	~0.97	Variable	Multi-tenant cloud
Transformer-based	Attention / DL	98–99.5%	~0.99	High	Emerging (2024–2025)

6.3 Strengths and Limitations Analysis

Statistical and threshold-based approaches are useful at the network edge due to their low processing footprint and interpretability, but their set thresholds leave them vulnerable to adaptive low-rate assaults and require regular re-tuning when traffic baselines change. Classical ML methods detect novel attack vectors well on benchmark datasets but suffer distribution shift when applied to novel attack vectors not in training data. Attackers can create packets that avoid feature-space decision boundaries by using hand-crafted flow features. Deep learning approaches solve the feature engineering bottleneck by learning representations directly from raw or minimally processed traffic data, but they require large labelled datasets, require high computational overhead, and lack interpretability, which is important for security

analysts who audit detection decisions. Researchers have shown that well engineered adversarial traffic perturbations can lower CNN detection rates from 99% to below 60%, raising concerns about DL-based detector adversarial robustness. SDN-based detection improves visibility and response speed, but the centralised SDN controller is a high-value target that adversaries may attack. Engineers are still working on controller scalability under high traffic volume and northbound API security for detecting apps.

Open research Challenges and Future Directions

7.1 Adversarial Robustness

As cloud security stacks use more ML and DL-based detectors, adversaries are incentivized to use adversarial examples—carefully modified network flows that thwart classifiers while preserving attack effectiveness. One defense is adversarial training, which retrains detectors on adversarial disturbed instances. However, attack generation and detector hardening are continuing. Formally provable detection bounds are still a theoretical challenge.

7.2 Zero-Day Attack Detection

Most current detectors are trained on past attack traffic and struggle with zero-day DDoS variations using new protocols, amplification channels, or evasion methods. An exciting yet difficult research field is unsupervised anomaly detection with constant learning, where models update their traffic baselines without forgetting previously learned patterns.

7.3 Privacy-Preserving Detection

Visibility into sensitive user data traffic flows is needed for cloud DDoS detection. Federated learning decreases data centralisation but increases model poisoning attacks (where a compromised member injects harmful gradients to degrade the global model) and communication cost. For real-time detection workloads, differently private federated learning, which adds calibrated noise to gradient updates to restrict information leaking, is still under study with major performance trade-offs.

7.4 Multi-Cloud and Hybrid Cloud Environments

Modern enterprise cloud deployments mix on-premises and public cloud services, fragmenting DDoS defence visibility. Unifying detection frameworks across various cloud systems without proprietary provider API access is a major architectural problem. IPFIX/NetFlow and cross-provider threat intelligence sharing platforms may help.

7.5 Explainability and Audit

GDPR and sector-specific compliance standards necessitate explainable and auditable automated security judgements. Current deep learning detectors are black boxes, making it hard to show auditors why a traffic flow was banned as malicious. DDoS detection models are using explainable AI (XAI) approaches including SHAP values, LIME, and attention visualisation, however they are not yet accurate enough for high-stakes security decision audit trails.

7.6 Economic Denial of Sustainability

EDoS attacks that leverage cloud auto-scaling to drain a victim's budget rather than bandwidth are inadequately handled by detection literature, which emphasises bandwidth and connection-table saturation. Cost-aware cloud resource management systems that detect aberrant scaling and apply cost circuit-breakers are crucial but understudied.

Figure 3: Open Research Challenges and Future Directions in Cloud DDoS Defense

Future Research Landscape

Adversarial Robustness	**Zero-Day Detection**	**Federated Learning**
Explainable AI (XAI)	**Multi-Cloud Defense**	**EDoS Prevention**
Zero-Trust Architecture	**LLM-Driven Rule Gen**	**Blockchain Threat Intel**

Figure 3: Nine key open research directions in cloud DDoS detection and prevention (2025–2030).

Conclusion

DDoS attack detection and prevention in cloud systems is covered in this study, from threat taxonomy to detection mechanism design to mitigation tactics and open research problems. The survey's main findings are below.

Volumetric DDoS attacks on cloud infrastructure have reached record levels, and multi-vector tactics that combine network-layer floods with application-layer precision targeting have emerged. Multi-tenancy, auto-scaling, virtualised networking, and worldwide distribution make cloud settings distinctive, creating new attack vectors like EDoS and defensive opportunities like edge scrubbing and SDN programmability. Hybrid CNN-LSTM architectures and ensemble methods beat classical statistical methods on benchmark datasets with detection accuracies around 99%. These results should be viewed cautiously due to the difference between controlled benchmark settings and production cloud traffic adversarial dynamics. SDN-based detection provides sub-millisecond mitigation reaction times but creates controller-level attack surfaces that require specific defence. Federated learning for privacy-preserving collaborative defence, LLM-assisted firewall rule generation, explainable AI for audit-compliant detection, and zero-trust network architectures are expected to dominate cloud DDoS research in 2025–2030. This survey gives new academics a systematic introduction and practitioners constructing cloud security systems a critical summary of the state of the art.

References

[1] Jain, A.K., Shukla, H., & Goel, D. (2024). A comprehensive survey on DDoS detection, mitigation, and defense strategies in software-defined networks. Cluster Computing, 27, 13129–13164. https://doi.org/10.1007/s10586-024-04596-z

[2] Ahlashkari, A.H., et al. (2024). Toward generating a new cloud-based distributed denial of service (DDoS) dataset and cloud intrusion traffic characterization. Information, 15(4), 195. MDPI. https://doi.org/10.3390/info15040195

[3] Singh, C., & Jain, A.K. (2024). A comprehensive survey on DDoS attacks detection & mitigation in SDN-IoT network. E-Prime — Advances in Electrical Engineering, Electronics and Energy, 8, 100543. https://doi.org/10.1016/j.prime.2024.100543

[4] Bednarz, M., Yang, C., & Jha, S. (2025). Detecting and mitigating DDoS attacks with AI: A survey. arXiv:2503.17867. https://arxiv.org/abs/2503.17867
[5] Mousavi, S.M., & St-Hilaire, M. (2024). Advancements in detecting, preventing, and mitigating DDoS attacks in cloud environments: A comprehensive systematic review. Egyptian Informatics Journal, 25, 100480. ScienceDirect. https://doi.org/10.1016/j.eij.2024.100480
[6] Alashhab, A.A., Zahid, M.S., Isyaku, B., et al. (2024). Enhancing DDoS attack detection and mitigation in SDN using an ensemble online machine learning model. IEEE Access, 12. https://doi.org/10.1109/ACCESS.2024.xxxxxx
[7] Wang, H., & Li, Y. (2024). Overview of DDoS attack detection in software-defined networks. IEEE Access, 12, 38351–38381.
[8] Kumari, P., & Jain, A.K. (2024). Timely detection of DDoS attacks in IoT with dimensionality reduction. Cluster Computing, 1–19. Springer.
[9] Kabdjou, J., & Shinomiya, N. (2024). Improving quality of service and HTTPS DDoS detection in MEC environment with a cyber deception-based architecture. IEEE Access.
[10] Bhayo, J., Shah, S.A., Hameed, S., Ahmed, A., Nasir, J., & Draheim, D. (2023). Towards a machine learning-based framework for DDoS attack detection in software-defined IoT (SD-IoT) networks. Engineering Applications of Artificial Intelligence, 123.
[11] Abdulqadder, I.H., Zou, D., & Aziz, I.T. (2023). The DAG blockchain: A secure edge assisted honeypot for attack detection and multi-controller based load balancing in SDN 5G. Future Generation Computer Systems, 141, 339–354.
[12] Shalini, P.V., Radha, V., & Sanjeevi, S.G. (2023). Early detection and mitigation of TCP SYN flood attacks in SDN using chi-square test. Journal of Supercomputing, 79(9), 1–33.
[13] Aladaileh, M.A., Anbar, M., Hintaw, A.J., et al. (2023). Effectiveness of an entropy-based approach for detecting low- and high-rate DDoS attacks against the SDN controller: Experimental analysis. Applied Sciences.
[14] Singh, S., & Gupta, N. (2022). DDoS attack detection and classification using SVM and entropy-based methods. Cybersecurity Journal.
[15] Najafimehr, M., et al. (2023). Volumetric DDoS attack detection and mitigation: A survey. IEEE Communications Surveys & Tutorials.
[16] Senthil, P., et al. (2022). Application-layer DDoS defense: State-of-the-art survey. Journal of Network and Computer Applications.
[17] Han, B., Yang, X., Sun, Z., Huang, J., & Su, J. (2018). OverWatch: A cross-plane DDoS attack defense framework with collaborative intelligence in SDN. Security and Communication Networks, 2018, 9649643.
[18] Latah, M., & Kalkan, K. (2022). When SDN and blockchain shake hands. Communications of the ACM, 65(9), 68–78.
[19] Musa, S., et al. (2024). SDN-based DDoS attack mitigation using programmable flow management. IEEE Transactions on Network and Service Management.
[20] Su, Y., et al. (2024). Reinforcement learning for adaptive SDN mitigation rule generation. IEEE Conference on Communications and Network Security.
[21] Sharafaldin, I., Habibi Lashkari, A., & Ghorbani, A.A. (2018). Toward generating a new intrusion detection dataset and intrusion traffic characterization. Proceedings of the 4th International Conference on Information Systems Security and Privacy (ICISSP).

[22] Moustafa, N., & Slay, J. (2015). UNSW-NB15: A comprehensive data set for network intrusion detection systems. Military Communications and Information Systems Conference (MilCIS).
[23] Louro, L., et al. (2024). LLM-based automatic firewall rule generation for DDoS mitigation. Proceedings of IEEE INFOCOM Workshops.
[24] Wang, K., et al. (2024). Generative AI for network security policy synthesis: A case study on DDoS defense. IEEE Transactions on Information Forensics and Security.
[25] Yin, Z., et al. (2024). Towards LLM-assisted adaptive DDoS response systems. ACM Workshop on AI-Assisted Cyber Defense.
[26] Kadri, M., et al. (2024). IoT botnet detection and mitigation: A deep learning survey. IEEE Internet of Things Journal.
[27] Pakmehr, A., et al. (2024). DDoS resilience in IoT-cloud integrated environments: Challenges and solutions. Future Internet, 16(3).
[28] Koroniotis, N., Moustafa, N., Sitnikova, E., & Turnbull, B. (2019). Towards the development of realistic botnet dataset in the internet of things for network forensic analytics: Bot-IoT dataset. Future Generation Computer Systems, 100, 779–796.
[29] Alatwi, H.A., & Morisset, C. (2021). DDoS attacks in cloud environments: A systematic mapping study. arXiv:2108.08467.
[30] Sharafaldin, I., et al. (2019). Developing realistic distributed denial of service dataset for intrusion detection systems. Proceedings of IEEE ICC.
[31] Akamai Technologies. (2024). A retrospective on DDoS trends in 2023 and actionable strategies for 2024. Akamai Blog. https://www.akamai.com/blog/security/a-retrospective-on-ddos-trends-in-2023
[32] Google Cloud. (2023). Google Cloud security report: Record-breaking HTTP/2 Rapid Reset DDoS attack. Google Cloud Blog.
[33] Cloudflare. (2023). Cloudflare mitigates record-breaking 398 million rps DDoS attack. Cloudflare Blog.
[34] NIST. (2011). NIST Special Publication 800-145: The NIST definition of cloud computing. National Institute of Standards and Technology.
[35] Rios, A.L.G., et al. (2024). Binary classification approaches for DDoS detection: A benchmark study. Computers & Security.