

Distributed Trust-Guided DDoS Defense for Information-Centric Networking Using a Name Resolution Service

Jiangfeng He* **, Yang Li* **

**National Network New Media Engineering Research Center,
Institute of Acoustics, Chinese Academy of Sciences, Beijing, China*

***University of Chinese Academy of Sciences, Beijing, China*

hejf@dsp.ac.cn, liyang@dsp.ac.cn

Abstract—Information-centric networking (ICN) uses name-based forwarding and in-network caching to reduce latency and backbone traffic and is therefore a key candidate for future networks. ICN architectures using a Name Resolution Service (NRS) improve scalability and flexibility, but also create new distributed denial-of-service (DDoS) threats against resolution and forwarding resources, such as cache pollution and forwarding exhaustion. We propose a trust-based defense for NRS-based ICN that binds entities to decentralized identifiers, derives multi-factor credit scores from resolution and caching behaviors, and uses lightweight credentials to steer rate limiting, caching, and forwarding policies. A federated-learning-based detector uses local observations to identify low-rate attacks that span multiple resolution domains without sharing raw traffic and with only modest extra coordination traffic. ns-3-based simulations under cache pollution, forwarding exhaustion, and hybrid attacks show that our scheme improves detection F1-score by up to about 10 percentage points and maintains request success and cache hit ratios around 88–90% and 70% during attacks, whereas representative ICN/NDN defenses can drop below 80% and 65%.

Index Terms—Information-centric networking, DDoS, distributed trust, decentralized identifier, federated learning

I. INTRODUCTION

With the rollout of 5G/B5G, the industrial Internet, and remote healthcare, future networks must deliver large amounts of content with strict latency, reliability, and availability requirements [1]. IP-centric architectures, which bind names to addresses and tightly couple control and data planes, are showing clear performance and security limitations [2]. ICN has therefore been proposed as a content-oriented paradigm that forwards on names or identifiers and uses in-network caching to reduce delay and backbone load.

ICN designs can be roughly split into two families. ICN architectures that use pure name-based routing, such as NDN/CCN, maintain PIT/FIB/CS tables and perform per-hop name-prefix matching [3]. ICN architectures using a Name Resolution Service (NRS), in contrast, delegate identifier-to-locator mapping to a name resolution system [4] and let the control plane distribute resolution results to forwarding devices in a protocol-agnostic way. This brings locality (edge resolution), elasticity (decoupled control/data planes), and

autonomy: deployments are typically partitioned into multiple resolution domains, where each domain groups a set of edge nodes, caches, and NRS nodes that share the same name-resolution and control policies [11], [12]. However, this also reshapes the attack surface: while it avoids per-hop state explosion, it introduces new weaknesses in cache management and in maintaining forwarding state along the resolution-control loop.

Most existing ICN DDoS defenses were built for ICN architectures that use pure name-based routing (e.g., NDN/CCN) [5]. They rely on per-hop state and static thresholds on metrics such as satisfaction ratios, PIT occupancy, or cache hits, and assume that each router can continuously measure per-name statistics. In ICN using a Name Resolution Service (NRS), however, the main bottlenecks move from per-hop PIT/FIB tables to the NRS cluster and the centralized control plane that installs and updates forwarding rules. Cache pollution and forwarding exhaustion attacks mainly manifest as bursts of resolution queries and rule updates along this resolution-control loop, overloading resolution servers and forwarding tables rather than individual routers. Simply porting NDN-style hop-by-hop defenses to this model is difficult, because forwarding nodes no longer maintain fine-grained per-name counters and global resolution state is concentrated in the NRS.

Under high concurrency, defenses that depend on centralized authentication or per-request cryptographic checks also add control overhead that conflicts with low-latency, high-availability requirements. This motivates a distributed, learning-assisted defense that is tightly integrated with the resolution and control planes of NRS-based ICN. In our design, the resolution domain is the basic unit for trust evaluation and collaborative defense. Our goal is to selectively admit and prioritize trustworthy requesters, suppress cache pollution and rule churn caused by low-trust entities, and still detect low-rate and collaborative attacks that span multiple resolution domains.

We propose a distributed trust-based anti-DDoS framework in which decentralized identifiers, multi-factor reputation, and lightweight credentials jointly steer resolution and forwarding,

allowing fast admission of trusted entities and elastic suppression of suspicious traffic, and limiting the impact of attacks on caches and forwarding state. A federated learning module helps detect cross-domain and collaborative attack patterns based on local observations from multiple domains.

The main contributions of this paper are summarized as follows:

- We propose a distributed trust-based defense framework for NRS-based ICN that binds traffic to Decentralized Identifiers (DIDs) and regulates access based on multi-factor behavioral credit scores.
- We design a lightweight gossip protocol to synchronize trust across resolution domains and a credit-aware policy enforcement mechanism to mitigate cache pollution and forwarding exhaustion.
- We integrate a federated learning module that enables collaborative detection of low-rate and cross-domain attacks without exposing raw traffic data.

The rest of this paper is organized as follows. Section II reviews related work on ICN security, DDoS mitigation, and trust management. Section III describes the attack and threat model in ICN using a Name Resolution Service (NRS). Section IV details the design and implementation of the proposed framework. Section V presents the evaluation methodology and ns-3-based simulations. Section VI concludes the paper.

II. RELATED WORK

Classic DDoS mitigation and name-resolution security on today's IP-based Internet rely on addresses, hierarchical name-to-address mappings, and centralized authorities. IP-layer defenses combine traffic filtering, rate limiting, and traceback, while DNS security hardens name servers and leverages authentication and access control. More recent systems incorporate reputation, machine-learning-based intrusion detection, and enhanced identity management, but they are still built around address-based forwarding and host- or server-centric trust anchors, and thus do not directly fit ICN's name-centric architecture and in-network caching.

Security in information-centric networking (ICN) is largely determined by names and in-network caches rather than IP addresses. Most work therefore focuses on ICN architectures that use pure name-based routing (e.g., NDN/CCN), where attacks such as Interest flooding, cache pollution, and content poisoning exploit name-based forwarding and caches [6], [7]. Existing defenses monitor traffic and cache statistics, apply rate limiting and pushback [8], adjust cache admission, and use data-centric signatures with ranking or reputation schemes, sometimes enhanced by machine learning [9], [10]. However, these mechanisms are tightly coupled to hop-by-hop name-prefix forwarding: each router keeps PIT/FIB/CS state and continuously measures per-name request and hit statistics. They also treat consumers as effectively anonymous—decisions are driven by aggregate traffic counters rather than long-lived identities or behavior histories—so forwarding, caching, and resolution cannot exploit per-requester trust information.

Beyond ICN architectures that rely purely on name-based routing (e.g., NDN/CCN), a line of work considers ICN using a Name Resolution Service (NRS), where a separate NRS translates object names or identifiers into locators or forwarding hints. RFCs 9138 and 9236 summarize the main design and architectural considerations for such NRS-based deployments [11], [12]. Concrete systems, including SDN-assisted NRSs and distributed or in-network schemes such as the Distributed Name Resolution System (DNRS) and the Distributed In-Network Name Resolution System (DIN-NRS), aim at scalable and low-latency resolution [13]–[15]. These proposals mainly optimize lookup performance, server placement, and control-plane efficiency, and they leave DDoS and cache-pollution resilience largely unaddressed, especially when attacks overload resolution servers or trigger excessive rule updates along the control path.

Meanwhile, trust and identity systems for ICN extend PKI or use blockchains [16], and more recent work adopts decentralized identifiers (DIDs), verifiable credentials (VCs) [17], and federated-learning-based intrusion detection [18], [19]. These schemes strengthen authentication and collaborative anomaly detection, but they are usually designed as overlays or generic security services, only loosely coupled to name resolution and caching. They do not directly feed trust or anomaly information into identifier resolution and cache or forwarding policies, particularly in multi-domain settings.

In contrast, this work targets ICN using a Name Resolution Service (NRS) and designs a distributed trust-based anti-DDoS framework in which DIDs and credentials provide low-overhead trust assessments for requesters, and multi-factor reputation is updated based on resolution, forwarding, and cache behavior. Trust results are fed back into the resolution and control planes to regulate which identifiers are resolved, at what rate, and with what forwarding constraints, while a federated-learning-based module lets domains collaboratively learn anomalous behavior without sharing raw traffic. By integrating distributed identity, trust evaluation, and anomaly detection into NRS-based ICN, the framework closes gaps in existing ICN/NDN defenses and strengthens the resilience of identifier-based networks to DDoS and cache pollution.

III. ATTACK MODEL

In this section we formalize two representative classes of attacks in ICN using a Name Resolution Service (NRS): cache pollution and forwarding exhaustion. We first state the attacker capabilities and goals, then detail each attack model and the corresponding metrics used later for evaluation.

A. Attacker Goals and Capabilities

We assume the following general capabilities:

- The attacker can control compromised hosts or bots at multiple access points or edge nodes, launching distributed request floods.
- The attacker can craft arbitrary identifier-based requests, including valid and invalid identifiers, semantically similar names, and short-lived dynamic names, and can freely

choose request rates and distributions (constant high-rate, bursty, mixed patterns, etc.).

- The attacker aims to degrade network availability and performance by increasing resolution and forwarding latency, lowering cache hit ratio, and causing rule churn in control/forwarding, rather than fully compromising controllers or resolution domains.

We further note that advanced threats such as Sybil identities, collusion, and federated-learning poisoning are possible in multi-domain settings. DT-NRS assumes consortium/domain-controlled DID enrollment (i.e., admission and revocation are enforced by the domain registry), and all trust reports and model updates are authenticated to prevent spoofing and replay. A full empirical study of Sybil/collusion/poisoning resilience (e.g., robust aggregation and outlier filtering) is left as future work.

B. Cache Pollution Attack Model

1) *Attack Motivation and Strategy*: The attacker manipulates request distributions so that edge and intermediate caches are occupied by low-reuse or non-popular objects, significantly reducing cache hit ratio for hot content. In ICN using a Name Resolution Service (NRS), cache pollution is tightly coupled to the resolution-cache loop: reduced hit ratio causes more resolution queries and back-end traffic, which in turn increases load on the NRS and SDN controller, exacerbating latency and cache replacement churn. This reinforcing loop is a key vulnerability.

2) *Attack Vectors*: Typical payloads include:

- **Long-tail/random requests**: near-uniform or high-entropy name distributions that continuously refresh caches with distinct objects.
- **Similar-name occupation**: names that are syntactically close to real hot names but map to useless or unreachable content, leading to wrong cache allocation.
- **Dynamic naming**: frequent name changes (e.g., time-stamped suffixes) for logically identical content, preventing cache reuse.

These strategies can be used individually or in combination, with tunable burstiness to evade static thresholds. Cache pollution attacks cause legitimate popular content to be pushed out by non-popular items and cache hits for hot content to collapse.

C. Forwarding Exhaustion Attack Model

1) *Attack Motivation and Strategy*: Forwarding exhaustion targets the resolution-control-forwarding chain and the limited rule capacity of forwarding nodes. By forcing the NRS and SDN controller to perform frequent resolution and rule installation/removal, attackers aim to exhaust forwarding table space or saturate control-plane bandwidth, delaying or blocking legitimate rule setups.

2) *Attack Vectors*: Typical vectors include:

- High-rate invalid-identifier requests triggering repeated resolution attempts and temporary rule installation.

- Dynamic short-lived sessions whose targets change rapidly, forcing frequent fine-grained rule updates and causing rule churn.
- Mixed requests that combine valid and invalid content to confuse detection and increase decision latency.

Attackers may choose centralized or distributed patterns to maximize desynchronization and delay. In a forwarding exhaustion scenario, malicious flows occupy forwarding resources and cause legitimate requests to be dropped.

D. Attack Scenarios and Evaluation Metrics

To facilitate repeatable experiments, we consider the following scenarios:

- **Single-domain persistent pollution**: high-entropy requests within one resolution domain with an attacker ratio of about 40%.
- **Distributed forwarding exhaustion**: multiple access points send high-rate invalid or short-lived name requests to stress the controller and forwarding tables.
- **Hybrid attacks**: combined cache pollution and forwarding exhaustion, including intermittent invalid-name bursts and dynamic naming.

As illustrated in Fig. 1, attackers located at multiple resolution domains send high-entropy requests, similar-name pollution traffic, and invalid identifiers toward edge nodes. In the single-domain scenario, malicious users are concentrated in one resolution domain and primarily stress a single NRS cluster, whereas in the distributed forwarding-exhaustion and hybrid scenarios the attackers are spread across several resolution domains and jointly overload multiple NRS clusters and forwarding tables. These malicious flows reduce the cache hit ratio for popular content and trigger bursts of resolution queries and rule updates at the NRS and SDN controller, increasing resolution and forwarding delay for legitimate users.




Figure 1. Cache pollution and forwarding exhaustion scenario

For each scenario we monitor three key metrics: F1-score of attack detection, cache hit ratio (CHR), and request success ratio (RSR), which are defined in Section V.

IV. TRUST-BASED DETECTION AND MITIGATION FRAMEWORK

This section details the design of the proposed framework, which integrates four synergistic components: (A) DID-based

identity management for accountability; (B) a multi-factor credit model to quantify requester behavior; (C) a gossip-based synchronization mechanism to share trust views across domains; and (D) a federated learning module for collaborative anomaly detection.

This section presents a distributed trust-based anti-DDoS framework for defending ICN using a Name Resolution Service (NRS) against cache pollution and forwarding exhaustion. As shown in Fig. 2, the architecture ties together four components: (i) DID-based identity and lightweight credentials for requesters, (ii) local credit evaluation from resolution and caching behavior, (iii) trust synchronization across resolution domains and credit-aware policy enforcement at edge nodes, and (iv) federated-learning-aided anomaly detection. Together, these components form a closed loop in which identities and behaviors are continuously assessed and fed back into resolution, forwarding, and caching decisions.

We assume that the deployment is partitioned into multiple resolution domains [11], [12]. A resolution domain is the set of edge nodes, caches, and NRS nodes that share the same NRS cluster and name-resolution and control policies, typically aligned with a hierarchical namespace or with the distributed NRS designs in [13]–[15]. The control plane maintains DID documents and public keys following the W3C DID specification, implemented over a consortium ledger or distributed key-value store. The data plane caches only verified key digests and validates lightweight credentials (e.g., JWTs) carried in request messages, avoiding full DID resolution on the fast path while anchoring each stream of requests to a stable identity. During enrollment, a domain authority issues or endorses the requester's DID document and records it in the domain registry; key rotation and revocation are handled by publishing an updated DID document or a revocation entry. Newly enrolled entities start from a neutral credit level and are subject to conservative rate limits until sufficient observations accumulate, mitigating cold-start abuse. Federated-learning update messages include a round identifier and sender signature, preventing unauthorized updates and replay.

A. Credit Evaluation Model

For each entity u , the credit evaluation layer summarizes its resolution and cache behavior at an edge node over a sliding window at time t . Let $H_u(t)$ denote long-term statistics (e.g., averages of request success and cache hits), and $R_u(t)$ denote recent observations such as resolution success ratio and cache hit ratio in the current window.

We first normalize these features:

$$\tilde{h}_u(t) = D_H^{-1}(H_u(t) - \mu_H), \quad \tilde{r}_u(t) = D_R^{-1}(R_u(t) - \mu_R),$$

where D_H and D_R are diagonal scaling matrices and μ_H, μ_R are mean vectors. We then form a combined feature vector

$$x_u(t) = [\tilde{h}_u^\top(t), \tilde{r}_u^\top(t)]^\top.$$

A linear scoring function produces an instantaneous credit score:

$$s_u(t) = w^\top x_u(t) + b, \quad \hat{C}_u(t) = \frac{1}{1 + e^{-s_u(t)}}.$$

To smooth short-term fluctuations and capture longer-term behavior, we apply exponential smoothing to obtain a reputation score

$$C_u(t) = \alpha C_u(t-1) + (1 - \alpha) \hat{C}_u(t), \quad \alpha \in [0, 1),$$

where $C_u(t)$ is the current (smoothed) credit or reputation. Larger α emphasizes history; smaller α reacts faster to recent anomalies. In practice, persistently low resolution success and low cache hit ratio for a given DID drive $\hat{C}_u(t)$ and hence $C_u(t)$ downward, reflecting the impact of cache pollution and forwarding exhaustion on that entity's behavior. In our implementation, w and b are obtained by training a logistic regression classifier on labeled ns-3 simulation traces containing both normal and attack traffic, and α is selected by grid search (we use $\alpha = 0.7$ as a trade-off between stability and responsiveness).

Based on $C_u(t)$, entities are classified into five levels: *trusted*, *normal*, *under observation*, *high-risk*, and *malicious*. Suspicious events (e.g., continuously low resolution success for a large number of distinct names) are encapsulated in signed reports that include the suspected DID, timestamps, and feature digests, and are disseminated to neighboring nodes as hints for collaborative defense.




Figure 2. Distributed trust- and federated-learning-based defense architecture. Solid arrows denote data/request/content flow; dashed arrows denote trust, control, and model flow.

B. Trust Synchronization and Policy Enforcement

The trust synchronization layer periodically exchanges credit information among neighboring edge nodes to build a near-consistent trust view across resolution domains. Lightweight gossip messages carry summarized credit updates; nodes merge these updates according to timestamps and signatures, and adjust their local $C_u(t)$ values accordingly. This avoids a single centralized trust database while ensuring that malicious entities are recognized network-wide.

Specifically, each node runs a lightweight gossip protocol every T_g seconds (we set $T_g = 2$ s in our experiments). In each round, a node randomly selects one or two neighbors and sends them a summary of the top- M DIDs whose credits have changed beyond a threshold ΔC since the last round.

Each summary contains tuples of the form $\langle \text{DID}, C_u(t), \tau, \sigma \rangle$, where τ is a logical timestamp and σ is a signature produced by the sender. Upon receiving summaries, a node merges records by keeping the entry with the latest timestamp τ and a valid signature, and updates its local $C_u(t)$ accordingly. Since each node exchanges at most $O(M)$ records per round with a constant number of neighbors, the communication overhead remains small even in large networks. More precisely, the per-node trust-synchronization bandwidth is bounded by $O(M/T_g)$ under bounded neighbor degree, independent of the total number of nodes N . In the prototype, gossip messages are signed with the sender's DID private key and verified using the corresponding public key in the DID document; we set $M = 50$ and $\Delta C = 0.1$, which keeps the additional gossip traffic below about 2% of the observed control-plane traffic in our simulations.

Policy enforcement directly maps credit to network behavior at edge nodes in each resolution domain:

- **Forwarding and rate control:** high-credit entities are given priority in resolution and forwarding and are subject to loose rate limits; low-credit and malicious entities are progressively throttled or blocked. As attacks degrade an entity's resolution success and cache hit behavior, its credit falls and its admissible request rate is reduced, protecting global request success.
- **Credit-aware caching:** content requested by high-credit entities is preferentially cached and retained longer; content requested by low-credit entities is generally not cached, and cache entries induced by them are evicted first under pressure. This biases cache resources toward trustworthy flows and preserves cache hit ratio for legitimate requests even under pollution.

When repeated anomalies cause a DID's credit to drop below the malicious threshold, the controller may broadcast a signed penalty notice within and across resolution domains. Other nodes then lower the corresponding entity's service priority or block its requests, forming a distributed collaborative mitigation against mobile or distributed attackers.

C. Federated Learning-Aided Detection

Local credit evaluation alone may miss low-rate or coordinated attacks whose impact on each individual node is weak but whose global effect is significant. By leveraging federated learning, edge nodes can jointly train a global anomaly detector without sharing raw traffic, thereby enhancing the detection of low-rate cross-domain attacks and feeding its decisions back into the credit update process.

Each edge node k maintains a local dataset $D_k = \{(x_i, y_i)\}$, where x_i contains compressed features derived from the resolution and caching behavior of entities (aligned with the variables used in the credit evaluation model), and $y_i \in \{0, 1\}$ indicates normal or attack labels derived from local rules and credit histories. A lightweight model $f_w(x)$ outputs an

anomaly score or attack probability. The local loss function is

$$L_k(w) = \frac{1}{|D_k|} \sum_{(x,y) \in D_k} [-y \log f_w(x) - (1-y) \log(1-f_w(x))],$$

and the global objective is a data-size-weighted sum:

$$F(w) = \sum_{k=1}^K \frac{n_k}{N} L_k(w), \quad N = \sum_{k=1}^K n_k.$$

We adopt FedAvg for aggregation [20]: in each communication round, a subset of nodes performs local updates on w , and the controller aggregates only model parameters. Raw traffic and detailed logs remain local, preserving data locality and privacy. In the experiments, f_w is implemented as a two-layer fully connected neural network with 32 hidden units and ReLU activations; each selected node runs 3 local epochs with a learning rate of 0.01 per round, and the controller performs 20 aggregation rounds per simulation.

The federated-learning output serves as an auxiliary signal to adjust credit. If the global model consistently assigns high anomaly scores to an entity across multiple resolution domains, the corresponding $\hat{C}_u(t)$ is scaled down before applying the smoothing update, accelerating the transition of that entity into the high-risk or malicious level. In our experiments, this federated-learning-aided adjustment notably improves detection F1-score and request success ratio under low-rate and collaborative attacks, while keeping false positives low.

V. PERFORMANCE EVALUATION

A. Simulation Setup

We implement the proposed framework on the ns-3 platform, extended to support ICN using a Name Resolution Service (NRS) with explicit name resolution and SDN-style control.

The topology has 20 edge nodes, 10 cache nodes, and 2 content producers. The 20 edge nodes are grouped into several resolution domains, each served by a local NRS cluster. A multi-level namespace and resolution path reflect this hierarchical resolution-domain structure. There are 300 legitimate users and 200 malicious users. Benign users generate requests according to a Zipf-like popularity distribution, while attackers follow the cache pollution and forwarding-exhaustion models in Section III, including high-entropy names, similar-name pollution, invalid identifiers, and dynamic short-lived names.

Each simulation runs for 250 s. Attacks start at 40 s and stop at 200 s so that both attack and recovery phases can be observed. The credit evaluation, gossip, and detection modules are implemented as extensions to the ns-3 forwarding strategy; unless otherwise stated, each link has a capacity of 10 Mb/s and 10 ms propagation delay, caches can store 1000 Data packets with LRU replacement, and the content namespace contains 10 000 objects.




Figure 3. F1-score, request success ratio (RSR), and cache hit ratio (CHR) under attacks

B. Metrics and Baselines

We use three metrics:

F1-score:

$$F1 = \frac{2TP}{2TP + FP + FN},$$

where TP , FP , and FN are true positives, false positives, and false negatives. Here, a higher F1-score means that the detector simultaneously achieves high recall (most attacks are detected) and high precision (most alarms correspond to actual attacks), whereas a low F1-score indicates many missed attacks and/or many false alarms, leaving the NRS and caches exposed or unnecessarily throttling legitimate users.

Request Success Ratio (RSR):

$$RSR = \frac{N_{resp}}{N_{req}} \times 100\%,$$

where N_{req} and N_{resp} are the numbers of total and successful requests. RSR measures the fraction of all requests that obtain a valid response; from the user perspective, a higher RSR corresponds to better service availability and fewer dropped or timed-out requests during attacks.

Cache Hit Ratio (CHR):

$$CHR = \frac{N_{hit}}{N_{req}} \times 100\%,$$

where N_{hit} is the number of requests served from in-network caches. CHR captures the efficiency of in-network caching: a higher CHR means that more requests are served from nearby caches, which reduces latency and lowers the load on origin servers and the backbone even when attacks are present.

We compare our framework, denoted DT-NRS (Distributed Trust for NRS-based ICN), with representative NDN defenses [5]–[8], [21]–[23] adapted to the same environment:

- token-based control-plane mitigation (Token) [6], [21];
- congestion-aware rate control (Cong) [5], [8];
- FITT-like feedback throttling (FITT) [8];
- cache-filtering (“cache shield”) against pollution (Cache) [7], [23];
- Isolation Forest (iForest) anomaly detection (IFOR) [25].

In Fig. 3 we use the abbreviations DT-NRS, Token, Cong, FITT, Cache, and IFOR to denote these schemes.

C. Results and Discussion

Figure 3 summarizes the F1-score, request success ratio (RSR), and cache hit ratio (CHR) under cache pollution, forwarding exhaustion, and hybrid attacks. In the figure, “DT-NRS” denotes the proposed framework, and the other curves correspond to the baseline schemes. In all three cases, the proposed distributed trust framework achieves the highest F1-score (about 0.94, 0.96, and 0.95), while the best baselines remain in the range 0.82–0.92, indicating fewer undetected attacks and fewer legitimate requests being misclassified.

The lower plots show that, during the attack window, baselines suffer a clear drop in RSR and CHR (often below 80% and 65%), whereas our framework degrades more slowly and recovers faster after the attack stops. RSR stays around 88–90% at the worst point and CHR remains close to 70%, meaning that roughly nine out of ten legitimate requests are still successfully served and that most of them can still be satisfied from caches rather than from the origin, so user-perceived availability and cache efficiency are largely preserved.

We also analyze the overhead of DT-NRS. On the fast path, credit updates and feature extraction require only constant-time bookkeeping and a constant-size metadata lookup per request (i.e., $O(1)$). On the control path, trust synchronization exchanges only top- M summaries every T_g seconds, yielding a per-node communication cost bounded by $O(M/T_g)$ under bounded neighbor degree; in our setup ($M=50$, $T_g=2$ s, $\Delta C=0.1$), the added gossip traffic stays below about 2% of the observed control-plane traffic. Federated learning communicates model parameters of size $|\theta|$ per round, avoiding raw-traffic sharing and keeping coordination lightweight.

VI. CONCLUSION

This paper investigated DDoS threats in ICN using a Name Resolution Service (NRS) and showed that cache pollution and forwarding exhaustion attacks are tightly coupled with name resolution, control, and in-network caching. Compared with ICN architectures that rely purely on name-based routing (e.g., NDN/CCN), ICN architectures using an NRS are more vulnerable to resolution overload and rule churn, and existing NDN defenses that rely on hop-by-hop state and static thresholds are difficult to apply directly in such environments.

We proposed a distributed trust-based anti-DDoS framework in which each entity is bound to a DID and a multi-factor credit model is derived from its resolution and cache behavior. Credit scores are synchronized through lightweight gossip and mapped to forwarding, rate-limiting, and cache policies, so that high-credit requests are prioritized while low-credit entities are progressively throttled and isolated. A federated learning module enables edge nodes to jointly learn anomaly patterns across multiple resolution domains without exposing raw traffic, enhancing the detection of low-rate and collaborative attacks.

Simulation results on an ns-3-based ICN topology using a Name Resolution Service (NRS) show that, under cache pollution, forwarding exhaustion, and hybrid attacks, the framework achieves higher detection accuracy and maintains higher request success and cache hit ratios than representative token-based, congestion-aware, feedback-based, cache-filtering, and machine-learning baselines.

In future work, we plan to validate the framework on a large-scale physical testbed and explore its deployment in emerging scenarios such as 6G edge computing and Industrial IoT, where lightweight, decentralized security is critical.

ACKNOWLEDGMENT

This work was supported by the self-deployed project of the Institute of Acoustics, Chinese Academy of Sciences: Research on Secure Transmission Techniques in Open Network Environments (Project No. E555180401).

REFERENCES

- [1] M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, "Toward 6G networks: Use cases and technologies," *IEEE Communications Magazine*, vol. 58, no. 3, pp. 55–61, 2020.
- [2] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *Proc. 5th Int. Conf. Emerging Networking Experiments and Technologies (CoNEXT)*, 2009, pp. 1–12.
- [3] L. Zhang *et al.*, "Named data networking," *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 3, pp. 66–73, 2014.
- [4] D. Raychaudhuri, K. Nagaraja, and A. Venkataramani, "MobilityFirst: A robust and trustworthy mobility-centric architecture for the future internet," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 16, no. 3, pp. 2–13, 2012.
- [5] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, "DoS and DDoS in named data networking," in *Proc. 2013 22nd Int. Conf. Computer Communication and Networks (ICCCN)*, 2013, pp. 1–7.
- [6] A. Compagno, M. Conti, P. Gasti, and G. Tsudik, "Poseidon: Mitigating interest flooding DDoS attacks in named data networking," in *Proc. 38th Annu. IEEE Conf. Local Computer Networks*, 2013, pp. 630–638.
- [7] M. Conti, P. Gasti, and M. Teoli, "A lightweight mechanism for detection of cache pollution attacks in Named Data Networking," *Computer Networks*, vol. 57, no. 16, pp. 3178–3191, 2013.
- [8] A. Afanasyev *et al.*, "Interest flooding attack and countermeasures in Named Data Networking," in *Proc. IFIP Networking Conf.*, 2013, pp. 1–9.
- [9] R. Tourani, S. Misra, T. Mick, and G. Panwar, "Security, privacy, and access control in information-centric networking: A survey," *IEEE Commun. Surveys Tuts.*, vol. 20, no. 1, pp. 566–600, 2018.
- [10] Z. Zhi, Y. Liu, J. Wang, and H. Zhang, "Resist interest flooding attacks via entropy-SVM and Jensen-Shannon divergence in information-centric networking," *IEEE Syst. J.*, vol. 14, no. 2, pp. 1776–1787, Jun. 2020.
- [11] J. Hong, T. You, L. Dong, C. Westphal, and B. Ohlman, "Design considerations for name resolution service in information-centric networking (ICN)," *RFC 9138*, Internet Research Task Force, Dec. 2021.
- [12] J. Hong, T. You, and V. Kaffe, "Architectural considerations of information-centric networking (ICN) using a name resolution service," *RFC 9236*, Internet Research Task Force, Apr. 2022.
- [13] J. Hong, T. Kim, J. Choi, and J. Park, "SDN-assisted name resolution service for information-centric networking," *IEEE Commun. Lett.*, vol. 22, no. 4, pp. 740–743, 2018.
- [14] W. Elbreiki, M. A. Maarof, and A. B. Hassan, "A distributed name resolution system in information centric networks," in *Proc. AIP Conf.*, vol. 1761, no. 1, art. no. 020038, 2016.
- [15] Z. Ma, J. You, and H. Deng, "DINNRS: A distributed in-network name resolution system for information-centric networks," *Comput. Commun.*, vol. 213, pp. 188–198, 2024.
- [16] Y. Zhang, C. Xu, H. Li, K. Yang, and J. Zhu, "Blockchain-based secure content caching in information-centric networking," *IEEE Access*, vol. 8, pp. 173543–173552, 2020.
- [17] M. Sporny, D. Longley, and M. Sabadello, "Decentralized Identifiers (DIDs) v1.0," *W3C Recommendation*, 2022. [Online]. Available: <https://www.w3.org/TR/did-core/>
- [18] R. Doriguzzi-Corin, S. Millar, S. Scott-Hayward, J. Martinez-del-Rincon, and D. Siracusa, "Lucid: A practical, lightweight deep learning solution for DDoS attack detection," *IEEE Trans. Netw. Service Manag.*, vol. 17, no. 2, pp. 876–889, 2020.
- [19] T. D. Nguyen, S. Marchal, M. Miettinen, H. Fereidooni, N. Asokan, and A. R. Sadeghi, "DfIoT: A federated self-learning anomaly detection system for IoT," in *Proc. 2019 IEEE 39th Int. Conf. Distributed Computing Systems (ICDCS)*, 2019, pp. 756–767.
- [20] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Proc. Int. Conf. Artificial Intelligence and Statistics (AISTATS)*, 2017, pp. 1273–1282.
- [21] H. Wang, Y. Liu, L. Wang, and K. Ren, "Detecting and mitigating interest flooding attacks in named data networking: A reinforcement learning approach," *IEEE Trans. Dependable Secure Comput.*, vol. 19, no. 1, pp. 199–213, Jan./Mar. 2022.
- [22] S. Bilgili, A. K. Demir, and S. Alam, "IfNot: An approach towards mitigating interest flooding attacks in Named Data Networking of Things," *Internet Things*, vol. 25, art. no. 101076, Apr. 2024.
- [23] Z. Rezaeifar, J. Wang, and H. Oh, "A trust-based method for mitigating cache poisoning in Name Data Networking," *J. Netw. Comput. Appl.*, vol. 104, pp. 117–132, Feb. 2018.
- [24] P. Kumar, R. Kumar, G. Srivastava, G. P. Gupta, R. Tripathi, and T. R. Gadekallu, "PPSF: A privacy-preserving and secure framework using blockchain-based machine-learning for IoT-driven smart cities," *IEEE Trans. Netw. Sci. Eng.*, vol. 8, no. 3, pp. 2326–2341, 2021.
- [25] F. T. Liu, K. M. Ting, and Z.-H. Zhou, "Isolation forest," in *Proc. 2008 IEEE Int. Conf. Data Mining (ICDM)*, 2008, pp. 413–422.