Lightweight Coordinated Defence Against Interest Flooding Attacks in NDN

Hani Salah Julian Wulfheide
Computer Science Department, TU Darmstadt
Hochschul Str. 10, 64289 Darmstadt, Germany
hsalah@cs.tu-darmstadt.de julian.wulfheide@ps.tu-darmstadt.de

Thorsten Strufe
Computer Science Department, TU Dresden
Mommsen Str. 8, 01187 Dresden, Germany
thorsten.strufe@tu-dresden.de

***Abstract*—Named-Data Networking (NDN) is a promising architecture for future Internet. However, routers and content providers in NDN can be targets for a new DDoS attack called the Interest Flooding Attack (IFA). As a consequence, affected routers drop legitimate interest packets.**

We argue that IFA can be defended effectively when it is detected and mitigated, at early stage, based on timely and aggregated information of exchanged packets and forwarding states. Towards this end, we adapt *CoMon*, a framework that we developed formerly to coordinate caching-related decisions in NDN. This choice is motivated by CoMon's proven ability to realize efficient, yet lightweight, coordination. A preliminary evaluation confirms the effectiveness of our solution against IFAs.

I. Introduction

The Internet, which was designed for reliable host-to-host communications, is heavily used today by content distribution applications. These applications generate massive traffic volumes, mainly due to continual distribution of popular content [1]. Aiming to narrow the gap between the Internet's design and usage, Named-Data Networking (NDN) [2] was proposed as an architecture for future Internet. It shifts the current host-centric communication model to a content-centric one.

Motivated by the importance of securing a prospective future Internet architecture, this paper deals with an NDN-tailored DDoS attack called the Interest Flooding Attack (IFA). In IFA, the adversary exploits two properties of NDN: (*i*) routing based on longest name-prefix match, and (*ii*) storing a forwarding state per interest packet in routers' Pending Interest Tables (PITs). More precisely, the adversary sends interest packets using unique non-existent content names (usually targeting a certain name-space). Consequently, one PIT entry is created per interest packet in each router on the path. These entries consume routers' memory till they eventually expire, and legitimate interest packets passing through these routers are dropped. Prior defence mechanisms against IFA (e.g. see [3], [4] and the references therein) are not highly effective, especially against low-rate IFAs, and/or incur high overhead.

In this paper, we propose an enhanced defence mechanism against IFA. It performs detection and mitigation of attacks based on *timely* and *aggregated* information of both: (*i*) exchanged packets and (*ii*) corresponding forwarding states. Practically, the new mechanism is based on CoMon, a framework for **Co**ordination with lightweight **Mon**itoring. CoMon was developed in our prior work [5] to coordinate caching-related decisions in NDN at a domain-wide scale, and succeeded to realize effective coordination with low overhead.

The proposed defence mechanism (Section II), aiming to achieve accurate and robust detection of IFAs, assigns the detection task to a small group of routers through which majority of network traffic is expected or enforced to pass. All packets pass through these routers before they reach the targeted victims of potential IFAs. Reactions to detected attacks are also performed at early stage, either by the same group or by ingress routers. Our preliminary results (Section III) show that the proposed mechanism can encounter IFAs effectively.

II. Overview of the Proposed Defence Mechanism

System architecture: The mechanism is currently designed to work within an autonomous domain. The system architecture of such a domain (Fig. 1) has three principal components. In the following, we introduce them and describe how they work together to defend against IFAs:

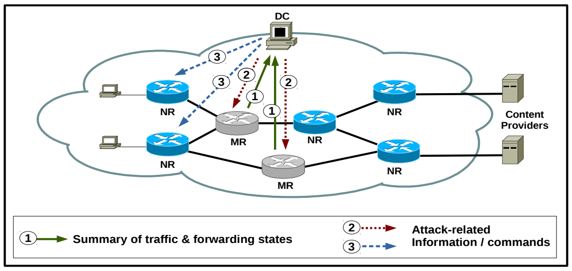
DC
NR
NR
MR
MR
NR
NR
NR
Content Providers
1
2
3
1 Summary of traffic & forwarding states
2 3 Attack-related Information / commands

Fig. 1: System architecture: illustrative autonomous domain (adapted from [5])

1) *Domain Controller (DC):* Each domain has a controller that collects monitored information about exchanged packets and corresponding forwarding states from a predetermined set of routers (hereafter, Monitoring Routers (MRs)). The DC uses this information to detect attacks. It then decides which routers should mitigate detected attacks, and shares attack-related information with them.
2) *NDN Routers (NRs):* These are similar to standard NDN routers [2]. However, part of them can be also commanded to mitigate detected attacks.

3) *Monitoring Routers (MRs):* A set of routers are selected as MRs. In addition to the functions of regular NDN routers, MRs monitor interest packets passing through them, and check whether these packets are satisfied (by data packets) or not. MRs also compute interests' satisfaction values per name-space and/or name-prefix. They periodically report summaries of their observations and results to the DC, which in turn sends them a feedback summarising information of attacks ongoing over the entire domain. Unless otherwise specified by the DC, MRs (by default) are also responsible for mitigating detected IFAs.

Detection and mitigation of IFAs: Attacks are detected at two levels: (*i*) at a domain-wide scale, performed by the DC on name-prefixes according to their (un)satisfaction ratios, and (*ii*) at each MR based on its own observations of (un)satisfaction per name-prefix per incoming interface. MRs can also consider the domain-wide state which they receive from the DC.

As for mitigation of IFAs, it is performed by default by MRs. However, in case the volume of attacks is very large, the DC can command ingress routers to react against detected attacks by specifying the corresponding (i.e. malicious) name-prefix(es). In the current implementation, defending routers use the satisfaction ratio as a direct probability for accepting or dropping interest packets. This approach was proposed in [3] and performed by each router. Instead, as mentioned above, our mechanism assigns reaction tasks to a small number of routers based on aggregated information. Furthermore, defending routers mark monitored packets to avoid overreactions.

Placement of MRs: Intercepting all or majority of network traffic by MRs is essential to detect IFAs accurately. Furthermore, intercepting traffic close from clients enables faster detection and mitigation. To this end, we propose a greedy algorithm called *PRCS* (Placement based on covered Routes and Closeness to Sources). As the name implies, PRCS weights a set S of routers based on the number of routes on which at least one router $r \in S$ is located, but also takes into the account r's distance from the sources of interest packets.

In particular, for each route t of length $L(t)$, the algorithm calculates the weight of each router r located on t as: $w(r,t) = 1 + \frac{P(r,t)}{L(t)}$. Here, $P(r,t)$ represents r's position on t: $P(r,t) = 0$ for the egress router (i.e. closest from the target), and incremented by 1 for each hop towards the ingress router.

In [5], we proposed two techniques (namely: MR-Aware Routing (MAR) and Forward-Till-Be-Monitored (FTBM)) to maximize traffic coverage achieved by MRs. With these techniques, all interest packets (thus the corresponding data packets) are *enforced* to pass through at least one MR.

III. Preliminary Results

We performed a simulation study to evaluate both the effectiveness and the messaging overhead of the proposed defence mechanism. In particular, we implemented the mechanism in ndnSIM [6], and fed the simulator with a real AS topology from [7]. We used the AS-3967 topology, consisting of 79 routers: 10% are selected (using the PRCS algorithm) as MRs. At the beginning of each simulation run, we randomly chose 70% and 3 routers as ingress routers and egress routers, respectively. Also, 25% of clients are selected randomly as attackers. Each attacker issues 1000 malicious interest packets per seconds, whereas each legitimate client issues 100 interest packets per second. Each simulation run lasts for nine minutes: the attack starts at the beginning of minute 2 and stops at end of minute 6. Uniform sizes of PITs (5000 entries) and caches (100 contents) are configured in all routers.

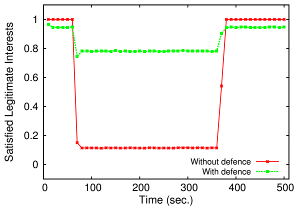

Fig. 2: Satisfied Legitimate Interests

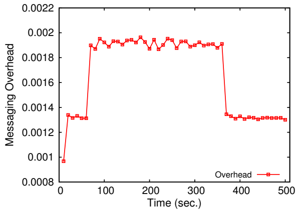

Fig. 3: Messaging Overhead

We evaluate the effectiveness of our scheme against IFAs by measuring the satisfaction ratio of legitimate interest packets. As can be seen in Fig. 2, the satisfaction ratio without defence was about 12%. In contrast, when the defence mechanism was enabled, this ratio increased to about 78%. However, one drawback of the defence mechanism is that false attack detections cause legitimate interest packets being dropped. As a result, satisfaction ratio was always below 100% even when no attack exists.

As for the messaging overhead, we measure it by normalizing the number of bytes used by defence-related messages (Fig. 1) over the number of bytes used by regular data packets. We used a uniform data packet's size of 1100 bytes. Fig. 3 plots the messaging overhead over time: We can see that it is very marginal even during the attack's period.

The results above suggest that the proposed solution is both effective and feasible.

IV. Ongoing and Future Work

We are currently working on reducing the overhead of the proposed defence mechanism, improving the monitors' placement and IFA's mitigation algorithms, and performing more extensive evaluation. We next plan to extend the mechanism to detect and mitigate IFAs over multiple domains, and employ it in the detection and mitigation of cache pollution attacks.

References

[1] "Cisco Visual Networking Index: Forecast and Methodology, 2013–2018," *CISCO White paper*, 2014.

[2] V. Jacobson *et al.*, "Networking named content," in *ACM CoNEXT*, 2009.

[3] A. Afanasyev *et al.*, "Interest flooding attack and countermeasures in Named Data Networking," in *IEEE IFIP Networking*, 2013.

[4] H. Dai *et al.*, "Mitigate DDoS Attacks in NDN by Interest Traceback," in *INFOCOM Workshops*, 2013.

[5] H. Salah and T. Strufe, "CoMon: An Architecture for Coordinated Caching and Cache-Aware Routing in CCN," in *IEEE CCNC*, 2015.

[6] A. Afanasyev *et al.*, "ndnSIM: NDN simulator for NS-3," *University of California, Los Angeles, Tech. Rep*, 2012.

[7] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," in *ACM SIGCOMM*, 2002.