NDN Interest Flooding Attacks and Countermeasures

Alberto Compagno*, Mauro Conti*, Paolo Gasti†, Gene Tsudik‡
*University of Padua, Italy — acompagn@studenti.math.unipd.it
†University of Padua, Italy — conti@math.unipd.it
‡New York Institute of Technology, USA — pgasti@nyit.edu
§University of California, Irvine, USA — gts@uci.edu

I. Introduction and motivation

While the Internet has been a great success story, the emerging usage models and new access methods expose some limitations of the current architecture, that was conceived back in 1970-s. To this end, there have been recent research efforts (e.g., [9], [6], [8], [2]) with the long-term goal of designing and deploying a next-generation Internet architecture. One such effort is Named Data Networking (NDN) [7].

NDN is an instance of Content-Centric Networking, where content – rather than hosts – is named and occupies the central role in the communication architecture. NDN is one of five NSF-sponsored Future Internet Architectures (FIA) [3] projects, and is an on-going research effort.

NDN is primarily oriented towards efficient large-scale content distribution. Consumers in NDN directly request (i.e., express *interest* in) pieces of content by name; the network is in charge of finding the closest copy of the content, and of retrieving it as efficiently as possible.

One of the key goals of NDN is *"security by design"*. In contrast to today's Internet, where security problems were (and are still being) identified, the NSF FIA program stresses both awareness of issues and support for features and counter-measures from the outset. To this end, our work investigates distributed denial of service (DDoS) attacks in NDN.

NDN supports two types of messages: *interests* and *content* [1]. Interests implement consumer requests and carry a (human-readable) name that identifies the desired content; content messages include a name, a payload and a digital signature computed by the content producer. Names are composed of one or more components, which have a hierarchical structure.

Content is delivered to consumers only upon explicit request. Each request corresponds to an *interest* message and causes NDN routers to store a small amount of transient state in a structure called PIT (Pending Interest Table). This information is used to route content back to consumers.

If the PIT is completely full, new interests are dropped. Flooding a router with interests allows the adversary to saturate the PIT. This has been identified in previous work under the name of interest flooding attack (IFA) [4]. However, there has been no evidence to support real-world plausibility of IFA. In particular, there is no experimental work that estimates the amount of resources (i.e., bandwidth, number of compromised hosts available to the adversary, etc.) required to successfully perform IFA.

We believe that IFA and its countermeasures deserve an in-depth investigation before NDN can be considered ready for large-scale deployment.

II. Contribution

Our contribution is two-fold: first, we show that implementing IFA using limited resources is indeed possible; then we design and evaluate a countermeasure that limits the effect of the attack. Our experiments are based on the official NDN implementation codebase; we argue that this setup provides reliable results, and closely mimics the behavior of physical (non-simulated) networks.

Effects of IFA. Our simulations show that the most effective way to implement IFA is using interests requesting non-existing content (hereby we use the name *fake* interests to identify them). In fact, state corresponding to such interests is not removed from PITs of intervening routers until it expires. This allows the adversary to quickly and efficiently fill up its victim's PIT.

We run simulations on two topologies: a simple architecture and the German research network DFN [5]. Due to lack of space, we only discuss our results on the (more realistic) DFN topology. We connected two producers, P0 and P1, at opposite sides of the DFN topology. The adversary controls three consumers, which issue only fake interests. While honest users ask for content from P0 and P1, malicious users generate interests that are forwarded only to P0.

Figure 1(b) shows the effectiveness of the attack. In particular, the adversary is able to significantly reduce the bandwidth allocated to content by one of the routers under attack (cf. baseline in Figure 1(a)).

Countermeasure. There are several parameters that routers can monitor to determine whether they have been targeted by a (successful) IFA. For example, a completely filled up PIT or a very low bandwidth available to forward content are very good indicators of an in-progress attack. However, routers that are carrying traffic from the attacker may not be able to identify such traffic as malicious.

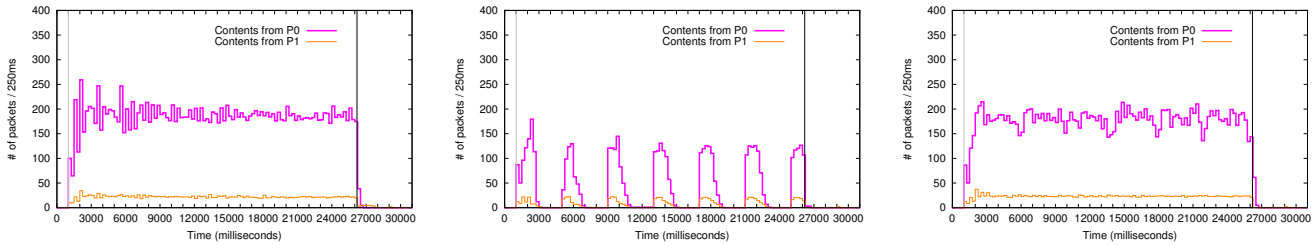

(a) Baseline (b) Attack (c) Countermeasure

Figure 1. IFA and countermeasure effectiveness compared to baseline. Values shown refer to bandwidth allocated for content by one of the routers forwarding fake interests.

For this reason, we focus on collaborative detection mechanisms that allow routers to exchange information about their state.

We suggest to periodically check some router's parameters (PIT usage and rate of unsatisfied interests). If they exceed a threshold, the router assumes that there is an IFA in progress. Routers react by rate-limiting interests from attacked interface(s) and by sending "alert" packets, containing information about the attack, to their neighbors. When a router receives an "alert" packet, it lowers the thresholds used to detect IFA. This allows routers that are not direct targets of the attack to detect (and mitigate) IFA.

Figure 1(c) shows the effects of our countermeasure. Due to the limited space, we only report results for one router in the DFN topology (figures 1(a), 1(b), and 1(c)). In particular, this router receives interests (fake and legitimate) directed to the producer under attack. Figure 1(a) shows throughput for content forwarded by the router without IFA. The thick line represents throughput for content coming from P0, while the thin line refers to the throughput of content from P1. Figure 1(b) shows how IFA reduces the bandwidth allocated to content, while Figure 1(c) shows the throughput of the incoming content on routers running our countermeasure. As shown in Figure 2 for different routers in the considered network (on x-axis), our solution is able to provide almost the same throughput as in the baseline scenario (y-axis).

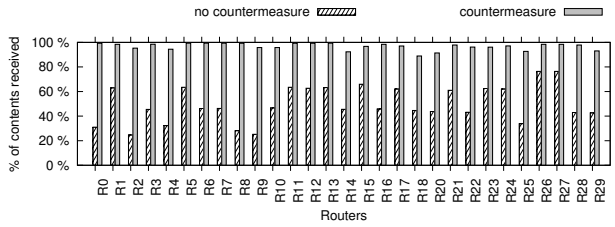

Figure 2. Effects of IFA and of our countermeasure with respect to the baseline.

III. Conclusion and Future Work

This work identifies and characterizes IFA, an effective DDoS attack on NDN. We provide, to the best of our knowledge, the first experimental evaluation of the attack.

We demonstrate that IFA is a realistic threat; in particular, we show that an adversary with limited resources can significantly reduce the amount of bandwidth allocated to content.

We introduce our techniques to detecting and mitigate IFA. We show that their benefits are significant, allowing routers to forward over 80% of the legitimate traffic when under attack.

As part of future work, we intend to further improve the effectiveness of our techniques introducing new metrics for early detection of IFA. Additionally, we plan to design more sophisticated reaction algorithms that explicitly take into account the topology of the network. We also plan to investigate new techniques for identifying and removing fake interests from PITs before they expire.

References

[1] CCNx protocol. http://www.ccnx.org/releases/latest/doc/technical/CCNxProtocol.html. Retrieved Jul. 2012.

[2] ChoiceNet - Evolution through Choice. https://code.renci.org/gf/project/choicenet/. Retrieved Jul. 2012.

[3] National science foundation (NSF) future of internet architecture (FIA) program. http://www.nets-fia.net/.

[4] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang. DoS & DDoS in named-data networking. Technical report, University of California, Irvine, 2012.

[5] O. Heckmann, M. Piringer, J. Schmitt, and R. Steinmetz. On realistic network topologies for simulation. In *ACM SIGCOMM MoMeTools*, pages 28–32. ACM Press, 2003.

[6] MobilityFirst FIA Overview. http://mobilityfirst.winlab.rutgers.edu. Retrieved Jul. 2012.

[7] Named data networking project (NDN). http://named-data.org.

[8] Nebula. http://nebula.cis.upenn.edu. Retrieved Jul. 2012.

[9] XIA - eXpressive Internet Architecture. http://www.cs.cmu.edu/~xia/. Retrieved Jul. 2012.