Analyzing the Producer-Consumer Collusion Attack in Content-Centric Networks

André Nasserala*† and Igor Monteiro Moraes†
*Universidade Federal do Acre
Centro de Ciências Exatas e Tecnológicas, Rio Branco, Brazil
Email: nasserala@ufac.br
† Universidade Federal Fluminense
Instituto de Computação, Niterói, Brazil
Email: igor@ic.uff.br

***Abstract*—This paper evaluates a denial-of-service attack in Content Centric Networks (CCN) that aims at increasing the content retrieval time. In this attack, malicious consumers and producers collude by generating content and changing content popularity. Malicious contents are stored by nodes and occupy the cache space that should be occupied by legitimate contents. Thus, the probability of legitimate consumers retrieves contents directly from the producer increases as well as the content retrieval time. We evaluate the impact of the attack by varying the number of consumers in collusion, the interest packets rate, and the way contents are requested. Results show if 20% of consumers are malicious and send 500 interests/s each, the content retrieval time experienced by legitimate users increases by 20 times.**

I. Introduction

Information-centric networking (ICN) is a new communication paradigm for the Internet [1]. ICN aims at delivering contents to users regardless of the location of these contents, as opposed to the TCP/IP architecture that aims at interconnecting end systems. One of the most cited architectures ICN in the literature is the Content Centric Networking (CCN) [2], [3]. One of the most advantages of CCN is the indirect content retrieval brought by the in-network caching technique. With CCN, any node in the network that receives a content request and has that content stored in cache is able to send back this content to the requesting node. The consumer is the node that requests and the producer is the source of this content. In-network caching allows the consumer to retrieve contents from nodes that are closer than the producer and thus the content retrieval time is reduced. In-network caching also increases the content availability and can reduce bandwidth usage, because the content traverses less number of hops towards the consumer. CCN, however, is not able to handle a particular Denial-of-Service (DoS) attack [4], [5] called producer-consumer collusion attack, which aims at increasing the content retrieval time. In this attack, malicious consumers request contents that are only available in malicious producers at a high rate. The retrieval time of legitimate contents increases because the cache miss ratio increases and thus legitimate nodes have to retrieve contents directly from the producer more often. In addition, the standard signature-verification mechanism of CCN does not detect the collusion attack because requests and contents sent by malicious nodes are legitimate from the network point of view.

In this paper, we demonstrate the impact of the producer-consumer collusion attack in CCN. To the best of our knowledge, none of the studies in the literature evaluates this attack. We perform simulation experiments considering different configurations. We vary the number of consumers and producers acting in collusion, the rate and the pattern of malicious-content requests. We consider the following metrics: legitimate-content retrieval time, cache miss ratio of legitimate contents, and the percentage of legitimate contents retrieved from the producer. Results show the attack compromises the performance of in-network caching employed by CCN. We conclude that if 20% of the consumer nodes are malicious and send 500 requests per second each, the legitimate-content retrieval time increases twenty times for the analyzed topology. Furthermore, we observe that up to 67% of the contents are retrieved directly from the producer for the analyzed configurations.

II. The CCN Architecture

CCN employs two types of packets only: interest and data packets. Consumers send interest packets to request a content. Producers and intermediate nodes store the content itself or pieces of it [1]. CCN nodes forward both interest and data packets based on the name of the content, instead of the destination address of the node that has the content. Each CCN node has two data structures used in packet forwarding: the Pending Interest Table (PIT) and the Forwarding Information Base (FIB). The PIT stores the state of each interest packet forwarded by a node and not already satisfied by a data packet. Each PIT entry also records the arrival interface of an interest packet. The PIT has a limited number of entries, therefore, new incoming interests are not forwarded if this table is full. This fact is explored by malicious users, as detailed in the next paragraphs. The FIB is used to forward interest packets to potential sources, i.e., producers and intermediate nodes. Each FIB entry has a name prefix and a list of output interfaces. Packets with names that match a given prefix must be forwarded to all these listed interfaces.

When a interest packet arrives, a CCN node first verifies its CS in order to find a copy of the content requested. The content name is indicated in the interest packet header. If the content is stored in the CS, the node sends the corresponding data packet to the consumer. Otherwise, the node verifies its PIT. If there is a PIT entry for the requested content, the node updates the corresponding entry by adding the arrival

978-1-4673-9292-1/16/$31.00 ©2016 IEEE

interface to the interface list and drops the interest packet. This procedure is called aggregation of interest packets and makes the CCN more robust against current DoS attacks. If there is no PIT entry, the node creates a new one and, then, looks up the FIB to determine the output interface to forward the interest packet. If there is no match in the FIB associated with the content name, the interest packet is discarded. CCN nodes repeat this forwarding process for each interest packet received. Data packets follow the reverse path traversed by interest packets, as the PIT stores the list of interfaces with pending interests. [6].

III. Related Work

The attacks in ICN are classified into four main categories: naming, routing, caching, and other miscellaneous related attacks [7]. Denial of service attacks in CCN are classified in two types: interest flooding and cache poisoning attacks [3]. The goal of interest flooding attacks is to overload the PIT with content requests sent by malicious nodes at a high rate [8]. Malicious interest packets, in general, contain names of contents that do not exist but nodes have to keep track of these pending malicious interests as they do for legitimate interests. The state of a malicious interest is not removed from the PIT until the timer defined for each entry expires. During this time, new interests for other contents that do not exist arrive at the node. In the worst case, the PIT is full, and the node under attack will not respond to legitimate interest packets, which compromises network performance. Gasti *et al.* [9] define the interest flooding attack and propose a push-back mechanism as a countermeasure.

The goal of the cache poisoning attack is to store polluted contents in the CS of nodes. Polluted contents are corrupted versions of legitimate contents modified by malicious nodes and requested by both legitimate and malicious consumers. The goals of the attack are twofold. First, the objective is to reduce the space available in cache to store legitimate contents and deliver polluted contents. Second, if malicious consumers request polluted contents, the objective is to remove legitimate contents from CS by assuming that nodes employ a cache replacement policy based on content popularity, such as LRU or LFU. Signature verification is the standard countermeasure employed by CCN against cache poisoning attack [6]. By default, signature verification is mandatory for consumers but not for intermediate nodes. Signature verification guarantees that consumers do not receive data packets containing polluted content. In this case, the CCN service can be denied if the consumers receive polluted contents very often. To enforce the signature verification of every content at each node implies processing overload and, for that reason, it is hard to deploy in practice [9], [5].

IV. The Producer-Consumer Collusion Attack in CCN

The collusion attack in CCN has at least two actors: the malicious producer and the malicious consumer. The former generates malicious contents according to the malicious consumer demand. These contents are similar to legitimate ones and thus are forwarded by CCN nodes with no differentiation, i.e., nodes cache malicious contents as they do for the legitimate ones. Malicious contents have also names that follow the CCN specification. The malicious consumer, in turn, requests malicious content at a high rate.

The collusion attack, aims at increasing the retrieval time of legitimate contents by forcing the consumer to retrieve the desired content from the producer more often. This goal is achieved by compromising the in-network caching technique through the manipulation of cached-contents popularity. Consumers send interest packets to contents that are only available in malicious producers at a high rate. Thus, a malicious node can increase the content popularity even if this content has not been requested by legitimate users. This is the reason to classify the content as malicious. The collusion attack is possible because CCN nodes employ cache replacement policies that are based on content popularity. A given content is considered less popular, for example, if it is not often requested or has not been requested recently by consumers. Consequently, this content is the first dropped to accommodate a new content in cache. By requesting a set of specific contents and at high rate, malicious nodes manipulate the cache. The more the number of malicious contents stored in cache, the lower the hit ratio for legitimate contents. Consequently, nodes have to retrieve contents directly from the producer more often, compromising in-network caching benefits. Even if the legitimate consumers do not have to retrieve contents directly from the producer their interest packets are probably forwarded for more hops until reach a node with the content desired in cache. From the network point of view, interest and data packets used in the collusion attack are legitimate, and, therefore, are not detected by signature verification mechanisms. Malicious data packets have a valid digital signature and carry the public key of the producer (or indicate how to obtain this key, which is not the focus of our work). Consequently, the verification of integrity and authenticity is successful and nodes are neither able to identify nor drop malicious data packets.

V. Evaluation Scenario

Our goal is to evaluate the impact of the collusion attack in the network performance. To achieve this goal, we consider the network topology illustrated by Figure 1 in our experiments. There are 32 nodes and the 24 leaf nodes are consumers. The number of legitimate consumers (LC) is the same for all the configurations analyzed and is equal to 16. The number of malicious consumers (MC) varies from 0 to 8. The position of the LCs and MCs is randomly chosen in each simulation run. The legitimate producer (LP) is always the root node. The malicious producer (MP) is always the child node of the LP. The other six nodes are routers (RTR), i.e., intermediate nodes that are able to cache contents in CCN. The transmission rate and delay are equal to 100 Mb/s and 1 ms, respectively, for all links. We consider the topology tree to the MP is in the path between the LC.

The malicious producer makes 12 contents available. Each malicious consumer sends interest packets to these contents at 10, 100, and 500 interest/s. Each malicious content has 100 chunks of 1024 bytes and a different name prefix. Malicious chunks are requested in two ways: requests based on the popularity of content, following a Zipf distribution with parameter $\alpha = 0.7$ [10], and sequential requests, called CBR [11], in which the consumer sends interest packets ordered by name in a cyclical manner. Legitimate consumers

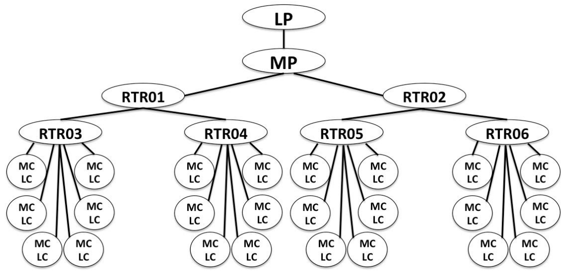

Fig. 1. The network topology used in simulation.

always send 10 interests/s to 12 other contents available in the legitimate producer. Each legitimate content has 100 chunks of 1024 bytes and a different name prefix. Legitimate contents are always requested according to a Zipf distribution with parameter $\alpha = 0.7$. The cache size of legitimate users and routers is equal to 1000 chunks of 1024 bytes. Malicious consumers do not have cache in order to always send interest packets regardless of previous requests. The PIT has unlimited size in order the avoid the side-effect discussed in the previous section and thus we observe only the effect of the increasing malicious occupation in nodes' cache. We use the ndnSIM module for NS-3 in simulation [11]. For each configuration, we perform 50 simulation runs of 180 s each. For every point of the curves, we calculated the confidence interval for a 95% confidence level.

VI. Results

The results presented in this section demonstrate the impact of the producer-consumer collusion attack on CCN performance. Figure 2 shows the behavior of the average retrieval time of legitimate contents as a function of the number of malicious consumers. For both configurations, when malicious consumers send interest packets sequentially, CBR, (Figure 2(a)) or when they request contents based on popularity by following the Zipf distribution (Figure 2(b)), the behavior observed is the same: the more the number of malicious consumers, the higher the content retrieval time. Similarly, the higher the rate of malicious interests, the higher the time to retrieve legitimate contents. For the configuration represented in Figure 2(a), for instance, if only legitimate consumers request contents (MC=0), the average retrieval time for legitimate content is 0.34 ms. On the other hand, if 4 malicious consumers request malicious contents, the time to retrieve legitimate contents increases to 1.92 ms and 7.21 ms, if they send 10 and 500 interests/s respectively. If there are 8 malicious consumers, the average retrieval time is 2.93 ms and 7.99 ms for rates of 10 and 500 interests/s, respectively. These results show the average retrieval time of legitimate contents increases by 23.5 times in the worst case for the configurations analyzed. Therefore, legitimate contents experience both lower cache miss rate (Figure 3) and lower retrieval time. The content retrieval time less than 1 ms if there is no attack is explained by the fact that legitimate consumers have cache and always request chunks based on its popularity. Thus, these nodes are able to retrieve contents from its own cache several times.

Figure 4 shows the legitimate producer's load, i.e., the percentage of legitimate contents retrieved from the producer, as a function of the number of malicious consumers for different interest transmission rates. These results corroborate the collusion attack reduces the efficiency of the in-network caching technique used by CCN. Figure 4(a) shows that only 0.5% of the contents is retrieved directly from the producer if none of the malicious users request contents. In this case, each legitimate content is retrieved from the producer twice at most, until it is stored by RTR1 and RTR2 (Figure 1). On the other hand, 4 malicious consumers operating at a rate of 10 interests/s increase the producer's load to approximately 12%. In the worst case, legitimate consumers retrieve about 67% of legitimate contents directly from the producer.

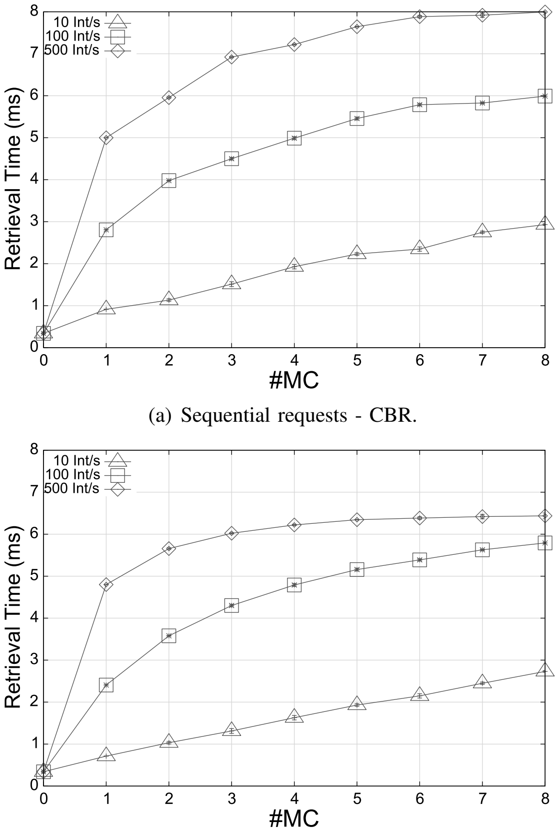

(a) Sequential requests - CBR.

(b) Popularity-based requests - Zipf.

Fig. 2. Average retrieval time of legitimate contents for sequential and popularity-based malicious requests.

Results also show that the geographical distribution of malicious consumers is more effective than the increase of the aggregated transmission rate of malicious interests. For example, Figure 2(a) shows that the average content retrieval time is in the order of 5 ms if 4 malicious consumers send 100 interests/s each (aggregated rate of 400 interests/s) or if only one malicious consumer sends 500 interests/s.

VII. Conclusion

We have evaluated a particular DoS attack for the CCN architecture, referred to as the producer-consumer collusion attack. This attack aims at increasing the retrieval time of legitimate contents by increasing the cache miss ratio of legitimate contents at intermediate nodes.

We have considered different configurations for the exper-

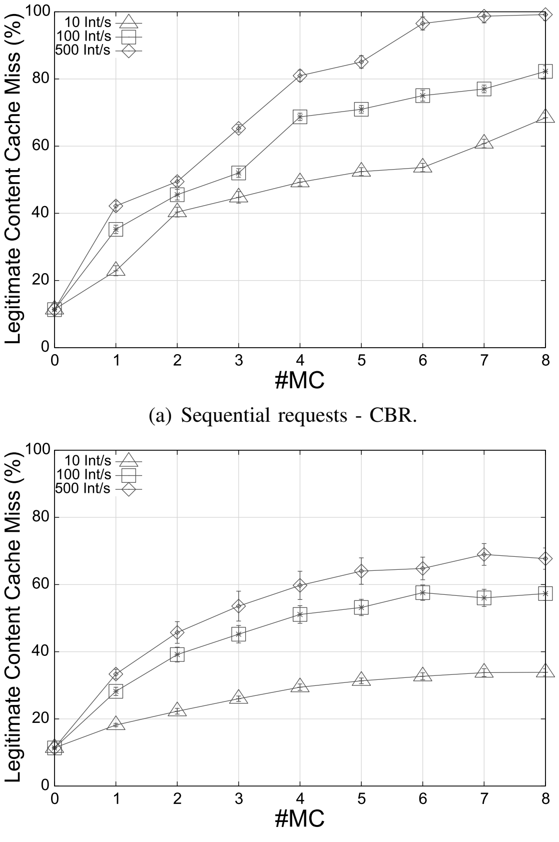

(a) Sequential requests - CBR.

(b) Popularity-based requests - Zipf.

Fig. 3. Average cache miss ratio of legitimate contents for sequential and popularity-based malicious requests.

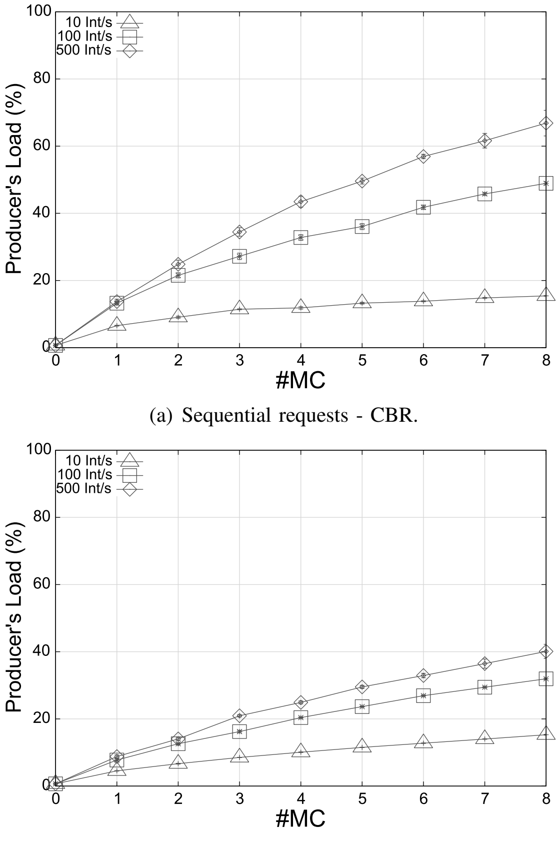

(a) Sequential requests - CBR.

(b) Popularity-based requests - Zipf.

Fig. 4. Average load of the legitimate producer for sequential and popularity-based malicious requests.

iments, ranging the number of malicious consumers and the transmission rate of malicious interests. We have concluded the collusion attack is effective because it compromises the in-network caching technique used by CCN. The average content retrieval time increases by 23.5 times if the network is under attack for the worst-case configuration analyzed. This is a consequence of the high cache miss ratio equal to 99% for this worst-case configuration. Legitimate consumers in this configuration retrieve 67% of the requested contents directly from the producer. We have also observed the geographical distribution of malicious consumers is more effective than the increase of the aggregated transmission rate of malicious interests. We intend to perform simulation for real network topologies, and then propose a countermeasure to the attack in future works.

Acknowledgments

This work is supported by CCET/UFAC, CNPq, CAPES, FAPERJ, Proppi/UFF, TBE/ANEEL, and CELESC/ANEEL.

References

[1] G. M. Brito, P. B. Velloso, and I. M. Moraes, *Information-Centric Networks, A New Paradigm for the Internet*, 1st ed., ser. FOCUS - Networks and Telecommunications Series. Wiley-ISTE, 2013.

[2] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R. Braynard, "Networking named content," in *International Conference on emerging Networking EXperiments and Technologies - CoNEXT*, Dec. 2009, pp. 1–12.

[3] D. Smetters and V. Jacobson, "Securing network content," Xerox Palo Alto Research Center - PARC, Tech. Rep. TR-2009-1, 2009.

[4] M. Xie, I. Widjaja, and H. Wang, "Enhancing cache robustness for content-centric networking," in *IEEE Conference on Computer Communications - INFOCOM*, Mar. 2012, pp. 2426–2434.

[5] M. Conti, P. Gasti, and M. Teoli, "A lightweight mechanism for detection of cache pollution attacks in named data networking," *Computer Networks - Elsevier*, vol. 57, no. 1, pp. 3178–3191, Aug. 2013.

[6] L. Zhang, D. Estrin, J. Burke, V. Jacobson, J. Thornton, D. K. Smetters, B. Zhang, G. Tsudik, K. Claffy, D. Krioukov, D. Massey, C. Papadopoulos, T. Abdelzaher, L. Wang, P. Crowley, and E. Yeh, "Named Data Networking (NDN) project," Xerox Palo Alto Research Center - PARC, Tech. Rep. NDN-0001, 2010.

[7] E. AbdAllah, H. Hassanein, and M. Zulkernine, "A survey of security attacks in information-centric networking," *Communications Surveys Tutorials, IEEE*, vol. 17, no. 3, pp. 1441–1454, thirdquarter 2015.

[8] A. Afanasyev, P. Mahadevan, I. Moiseenko, E. Uzun, and L. Zhang, "Interest flooding attack and countermeasures in named data networking," in *IFIP Networking*, May 2013, pp. 1–9.

[9] P. Gasti, G. Tsudik, E. Uzun, and L. Zhang, "DoS and DDoS in named-data networking," in *International Conference on Computer Communications and Networks - ICCCN*, Aug. 2013, pp. 1–7.

[10] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: Evidence and implications," in *IEEE Conference on Computer Communications - INFOCOM*, Mar. 1999, pp. 126–134.

[11] A. Afanasyev, I. Moiseenko, and L. Zhang, "ndnSIM: NDN simulator for NS-3," NDN, Technical Report NDN-0005, October 2012. [Online]. Available: http://named-data.net/techreports.html