

Traffic Analysis Resistant Network (TARN) Anonymity Analysis

Nathan Tusing, Jonathan Oakley, Geddings Barrineau,
Lu Yu, Kuang-Ching Wang, and Richard R. Brooks
Department of Electrical and Computer Engineering
Clemson University, Clemson, SC, USA
{ntusing,joakley,cbarrin,lyu,kwang,rrb}@g.clemson.edu

Abstract—We proposed a Traffic Analysis Resistant Network (TARN) that randomizes IP addresses in a fashion similar to Frequency Hop Spread Spectrum (FHSS), allowing users to blend into background traffic. IP hopping alone is not enough. TARN may still be susceptible to side-channel analysis. To remove the vulnerabilities, we introduce a SDX-based solution. In this work, we describe the design and implementation of TARN and experimental environment used to test TARN.

Index Terms—Traffic Analysis, SDX, TARN, NMTD

I. INTRODUCTION

Every day millions of individuals access their preferred news source. The supposition is that using HTTPS provides sufficient anonymity. Unfortunately, even when additional countermeasures are used, it is possible to identify specific web sites using packet sizes, traffic volume, and other features [1]. Worse, it is trivially easy to re-route this traffic globally [2]. This means that, with no particular expertise, it is possible to learn the political leanings of an individual, or to identify groups of people with similar political views.

In order to mitigate the effects of this malicious traffic analysis, we proposed a Traffic Analysis Resistant Network (TARN) that reduces the effectiveness of traffic analysis by randomizing IP addresses [3]. Our network-based moving target defense (NMTD) randomizes IP addresses to disrupt traditional flow-based traffic analyses. This novel approach to traffic analysis circumvention utilizes the ideas behind BGP hijacking and Frequency Hop Spread Spectrum (FHSS) [4] in order to hide the true destination. Although IP hopping removes the most effective feature that can be used to associate a communication session with the users, TARN users may still be vulnerable to side-channel analysis. We propose a SDX-based solution to eliminate the side channels of TARN.

TARN is similar to the Tor’s onion routing approach, which solves the anonymity problem by creating three hop encryption circuits. Tor requires end-users to install software on their machines, which adds overhead and deters potential users. TARN operates at the infrastructure level. Users are not required to install or configure software. TARN utilizes an SDN’s ability to redirect packets at the SDX layer by controlling the flow tables on Openflow Switches - effectively altering routing.

The core idea behind TARN is SDX nodes dynamically announcing IP prefixes and randomly forwarding traffic through



Fig. 1. TARN implementation strategies for the SDX-based solution.

other TARN SDXs. The connection between TARN SDXs uses link-layer encryption, along with the connection between an autonomous system (AS) and a TARN SDX node. Traffic between TARN SDX nodes will appear encrypted and destined for random IP addresses. This IP address hopping is similar to frequency-hopping spread spectrum (FHSS) [4] for wireless communications, which was used to solve a similar issue.

BGP announcements are used to ensure the traffic is routed to the TARN server, where the traffic is decrypted, the IP address is mapped back to its original value, and the traffic is sent to its final destination.

II. SDX-BASED TARN

Traditional VPNs leak side channel information, which can be used by attackers to associate packets with a target user. To prevent side-channel leakage, we propose the SDX-based TARN that provides traffic analysis resistant communication between any two SDX centers. TARN users connect to a nearby SDX center via a secure link layer connection. A user’s traffic is then randomly forwarded at least one hop according to a particular probability distribution before it is sent over the Internet to its final destination. This effectively removes side-channels like packet size and inter-packet delays. To demonstrate the performance of the SDX-based TARN, we compare the false positive rate (FPR) and true positive rate (TPR) of an observer attempting to associate packets with a given user. A low FPR and a high TPR indicate an observer can effectively de-anonymize a given user, while an FPR and TPR of 0.5 indicate the observer has no additional information about the user.

A. Baseline Experiment

The objective of the baseline experiment is to demonstrate that encrypted channels are vulnerable to side-channel analysis. The results are used as a reference to show the extra layer of anonymity offered by TARN.

In the baseline experiment, an *attacker* attempts to associate packets within an encrypted stream using side channel information. Specifically, we consider two side channels – packet size and interpacket delays. It is assumed three users tunnel their traffic through the same encrypted channel. The purpose of this experiment is to show that the attacker is able to group the packets based on side channels.

Packets will be generated according to six different normal probability distributions. The means of the packet size are 1kB, 2kB, and 3 kB, respectively. The standard deviation of the packet size of each distribution is 250B. The means of interpacket delays of the three sessions are 50ms, 90ms, and 130ms. The standard deviation for each interpacket delay distribution is 10ms.

We will show that we are able to correlate each packet with a particular session based on their size and interpacket delays using the Hidden Markov Models (HMMs)-based approach described in [5]. The observation window is the number of packets used to determine the prior probability of packet association. The trade-offs between several different window sizes along with different distributions of packet sizes and interpacket delays will be considered. Receiver operating characteristic (ROC) curves will be generated to assess the ability (FP rate and FN rate) of the attacker, which in other words, is the anonymity level of simple encrypted tunnel.

B. TARN Anonymity Analysis

In the second experiment, we will show that the side-channels that are used to associate packets in the baseline experiment are eliminated with the adoption of TARN. The probability of making a Type I error and that of making a Type II error are used as the measurement of user anonymity offered by TARN.

The attacker tries to associate packets in the encrypted communications between two TARN SDX nodes. A simple TARN configuration will be used, where three TARN SDX nodes are connected with each other via an encrypted link. Traffic is sent to the ingress TARN node through a trusted gateway using link layer encryption. The traffic is then randomly forwarded to the next hop according to a uniform distribution. When the traffic arrives at the next TARN node, it is sent over the Internet to its final destination. Since a user’s traffic is split between two different TARN SDX nodes, the attacker will not have a complete view of the traffic. The attacker is assumed to be able to observe traffic on a particular link. The same set of probability distributions from the baseline experiment will be used to generate the traffic, but the attacker will only be able to view half of the traffic. Various window sizes will be tested to determine trade-offs between accuracy and performance. As with the previous experiment, different mean values of packet size and interpacket delay will be considered.

III. PRELIMINARY RESULTS

TPR and FPR rates are considered for a preliminary analysis of the results to provide a tangible comparison between the user anonymity offered by a regular VPN and TARN, as shown in Table I. Figure 2a shows that some information can be




Fig. 2. Packet size side-channel comparison between VPN and TARN.

TABLE I
TRUE POSITIVE AND FALSE POSITIVE RATES OF PACKET ASSOCIATION USING PACKET SIZE SIDE-CHANNEL.

Anonymity Method	TPR	FPR
VPN	0.725	0.275
TARN	0.333	0.666

gained by looking at the packet size distribution, and the TPR and FPR rates reflect this. Conversely, no information can be gained by observing the TARN packet size distribution, shown in Figure 2b, and the TPR and FPR in Table I show that an observer has no additional information regarding packet association.

IV. CONCLUSION AND FUTURE WORK

TARN provides a novel infrastructure-level approach to circumventing traditional traffic analysis. TARN is based on emerging SDN technologies and the GENI and PEERING testbeds which ensure the test conditions are representative of the Internet.

Future work will (1) change the data collection methods to use a browser instead of wget, (2) investigate possible reasons for slow speeds, (3) continue to compare the traffic features of TARN to unprotected traffic, and (4) conduct an in-depth security analysis. Our future security analysis will extend beyond the proof-of-concept we present here and focus on deep packet inspection, side-channel attacks, and full statistical comparison between TARN and unprotected traffic.

ACKNOWLEDGMENT

This material is based upon work sponsored by the National Science Foundation under Grant No. 1643020. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] M. Yang, X. Gu, Z. Ling, C. Yin, and J. Luo, “An active de-anonymizing attack against tor web traffic,” *Tsinghua Science and Technology*, vol. 22, no. 6, pp. 702–713, 2017.
- [2] A. Gavrichenkov, “Breaking https with bgp hijacking,” *Black Hat. Briefings*, 2015.
- [3] L. Yu, Q. Wang, G. Barrineau, J. Oakley, R. R. Brooks, and K.-C. Wang, “TARN: A SDN-based traffic analysis resistant network architecture,” *arXiv preprint arXiv:1709.00782*, 2017.
- [4] S. M. Schwartz, “Frequency hopping spread spectrum (FHSS) vs. direct sequence spread spectrum (DSSS) in the broadband wireless access and WLAN arenas,” *white paper*, 2001.
- [5] J. M. Schwieter, R. R. Brooks, C. Griffin, and S. Bukkapatnam, “Zero knowledge hidden markov model inference,” *Pattern Recognition Letters*, vol. 30, no. 14, pp. 1273–1280, 2009.