Enhancing Traffic Analysis Resistance for Tor Hidden Services with Multipath Routing

Lei Yang and Fengjun Li
The University of Kansas, Lawrence, Kansas, US, 66045
Email: {lei.yang, fli}@ku.edu

***Abstract*—Hidden service is a very important feature of Tor, which supports server operators to provide a variety of Internet services without revealing their locations. A large number of users rely on Tor hidden services to protect their anonymity. Around 30,000 servers are running hidden services every day. However, hidden services are particularly vulnerable to traffic analysis attacks especially when an entry guard of a hidden server is compromised by an adversary. In this paper, we propose a multipath routing scheme for Tor hidden servers (*m*TorHS) to defend against traffic analysis attacks. By transferring data through multiple circuits between hidden server and a special server rendezvous point, *m*TorHS is able to exploit flow splitting and flow merging to eliminate inter-cell correlations of the original flow. Experiments on the Shadow simulator [1] show that our scheme can effectively mitigate the risk of traffic analysis even when the most robust watermarking technique is applied.**

I. Introduction

Tor hidden service is a very important factor that makes Tor [2] be the most popular and widely deployed low-latency anonymous communication system today. The hidden service allows server operators to hide their locations while providing a variety of Internet services via the rendezvous point. This is a very appealing feature that makes Tor stand out, because anonymous publishing is of great importance especially for people in countries with strict censorship, however, other popular low-latency anonymity systems such as Anonymizer and Java Anon Proxy (JAP) do not support such hidden service since it is out of the scope of their initial designs. Therefore, a large number of users with strong anonymity needs deploy their services on the Tor network for its practical support to location-hidden services and low latency.

However, Tor hidden services are still under the risk of de-anonymization due to specialized traffic analysis attacks [3]. It is argued that the current Tor design is vulnerable to traffic analysis attacks if the adversary can monitor a user's traffic entering and leaving the anonymity network at both the sender side and the receiver side. Since the malicious client is always at one end of the anonymous path in hidden services, she can successfully perform the attack if she is able to observe the traffic at the hidden server end. Øverlier et al. proposed the first documented attack against Tor hidden services by exploiting traffic analysis. The effectiveness of such attacks is mainly caused by the low latency in anonymized paths, which unwilling preserve the inter-cell timing correlation between the original flow and the anonymized flow. From it the adversary can exploit traffic analysis techniques to correlate the mutual information between the original flow and the anonymized flow to infer the communication relationship and the identities. Therefore, the key to mitigating the threats of traffic analysis attacks is to reduce the timing correlation between cells.

In this paper, we propose a multipath routing scheme for Tor hidden services (*m*TorHS) to defend against traffic analysis attacks. Our scheme routes data cells between the rendezvous point and the hidden server through multiple circuits, which exploits flow splitting and flow merging functionalities of multipath routing to remove identifiable patterns of the original flow. Through experiments on the Shadow simulator [1], we show that *m*TorHS is resistant to traffic analysis, even if the most robust watermarking-based attack is applied.

II. Previous Solutions

To mitigate this threat described in Section I, efforts can be made from two perspectives: (1) preventing an adversary from controlling two edges of a circuit to impede the occurrence of traffic analysis (2) reducing the success rate of traffic analysis even when two edges of a circuit are both compromised.

"Entry guards" [3] tried to solve this problem following the first direction. Each user constructs its guard set consisting of three more reliable routers by default,, which will expire in 30 to 60 days. With entry guards, whenever the hidden server builds a circuit to the rendezvous point in response to a client's request, it will picks an entry guard from the set for its first hop instead of choosing a random router in the network. The insistence on fixed first-hop routers prevents the adversary from inducing HS to choose a malicious entry node. As a result, the chance that an adversary controls both edges of a circuit is significantly reduced. However, it is unreliable that the security of hidden servers merely relies on the goodness of the entry guard set. Given enough time, a user will eventually selects a malicious entry node into his guard set. Johnson et al. showed that for an adversary with moderate bandwidth capacity, it only takes 50 to 60 days to include a malicious router to a user's guard set [4]. As noted by Elahi et al., the design of entry guard is still an unclear research problem [5] and subtle parameter selection is required to achieve expected protection. Therefore, it is critical to develop parallel protection mechanisms to enhance the resistance of Tor hidden services to traffic analysis attacks in cases where the guard set is compromised.

III. Multipath Tor Hidden Services (*m*TorHS)

In this paper, we make efforts following the second direction to reduce the success rate of traffic analysis when

978-1-4673-7876-5/15/$31.00 ©2015 IEEE

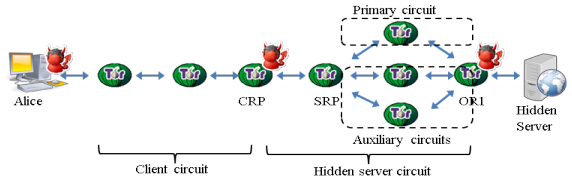

Fig. 1: *m*TorHS architecture

the attacker has successfully controlled both edges of a Tor circuit. We present a multipath routing scheme illustrated in Figure 1 for Tor hidden services. This scheme is based on the key insight that the traffic pattern observed or intentionally generated at the malicious entry guard (e.g., OR1) will be somewhat distorted by flow splitting and flow merging operations in multipath routing and by the multiple routes with different network dynamics.

The client's connection initialization remains the same as the current Tor hidden services: Alice first selects a client rendezvous point (CRP) and constructs a rendezvous circuit to it. Then, she builds an introduce circuit to one of HS's introduction points and sends an introduce message to request the hidden service and inform HS of the CRP's address.

After receiving the introduce request, HS decrypts it with its private key and extracts the address of CRP. Then, HS selects its own rendezvous point (SRP). The selection of SRP is very critical. From Figure 1, we can see that subflow merging occurs at SRP. Hence, even when multipath routing is adopted between SRP and HS, if the adversary controls SRP and OR1, she can observe traffic patterns from both ends of each subflow and thus perform traffic analysis successfully. When the adversary sends a large number of requests, If HS selects a new SRP for each received access request, it may eventually select one of the compromised router. Inspired by the entry guard idea, we propose "rendezvous guard" for SRP selection, which is a set of reliable routers selected by the hidden server. A hidden server initially selects three routers to compose its rendezvous guard set, each of which stays in the set for a random period between 30 and 60 days. Whenever HS builds a rendezvous circuit in response to the client request, it sticks to the same rendezvous guard set and randomly picks one router from it.

Once SRP is selected, HS builds an anonymous tunnel consisting of m circuits to it, following the same approach described in [6], where m is a server specific parameter. A primary circuit is first built to SRP to obtain the tunnel identifier (TID), and others *m-1* auxiliary circuits join the tunnel using the same TID. All m circuits go through the same entry guard OR1 and merge at SRP, which further relays the merged flow towards the client. HS then splits the original flow onto m subflows and attach each subflow to a circuit in the tunnel. HS is responsible for assigning data cells to subflows. To reduce the likelihood of inter-cell correlations, HS randomly assigns data cells to subflows with different capacities. As a result, a data cell from a fast circuit needs to wait at SRP for its earlier cells arriving from other slow subflows to be merged in an orderly manner. In this way, we maximize the network properties of different circuits to distort or destroy the possible traffic pattern inserted by the malicious guard, which greatly reduces the inter-cell correlation.

IV. Experiment Evaluation

To evaluate the enhanced anonymity, we test the performance of *m*TorHS against a well-known active traffic analysis attack, i.e., interval centroid-based watermarking (ICBW) [7]. The adversary embeds a watermark to the victim's flow at OR1 and observes it at CRP. We implement *m*TorHS on Tor v0.2.5.6-alpha and build a private Tor network in the Shadow simulator with 50 Tor routers, 1 hidden server, 20 general HTTP servers, 1 malicious client and 100 general web clients. Among the 50 routers, two are configured as malicious CRP and OR1. We also choose an extreme setting for comparison, where the adversary is the only client in the network.

We perform the ICBW attack on the original Tor and *m*TorHS, where m is set to 2, 4, 6 and 8. Table I shows the comparison in terms of Hamming distance between Tor and *m*TorHS with different settings. A larger Hamming distance indicates that the anonymity system can better transform the original flow and prevent the traffic analysis. No matter for general cases or extreme cases, *m*TorHS can better obscure the embedded watermark in the victim's flow.

	Tor	*m*TorHS			
		m=2	4	6	8
General case	6	9	9	10	12
Extreme case	3	8	9	9	11

TABLE I: Comparison of Hamming distance between Tor and *m*TorHS with different m where each flow is encoded using different watermarks.

V. Conclusion

Tor hidden service is a very important tool to provide receiver anonymity to server operators, but it is vulnerable to traffic analysis attacks especially when the entry guard protection is broken. In this paper, we propose a multipath routing based scheme that exploits flow mixing and flow merging to distort or destroy inserted traffic patterns in a victim's flow. We believe this is an effective complement to the existing protection mechanism.

Acknowledgments

This work was partially supported by the National Science Foundation under Award EPS0903806, KU General Research Fund under Award GRF2301075, and KU Research Investment Council Strategic Initiative Grant under Award INS0073037.

References

[1] R. Jansen and N. Hopper, "Shadow: Running Tor in a Box for Accurate and Efficient Experimentation," in *Proc. of NDSS Symposium*, 2012.

[2] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proc. of USENIX Security Symposium*, 2004.

[3] L. Øverlier and P. Syverson, "Locating hidden servers," in *Proceedings of the 2006 IEEE Symposium on Security and Privacy*, May 2006.

[4] A. Johnson, C. Wacek, R. Jansen, M. Sherr, and P. Syverson, "Users get routed: Traffic correlation on tor by realistic adversaries," in *Proc. of the 20th ACM conf. on Computer and Communications Security*, 2013.

[5] T. Elahi, K. Bauer, M. AlSabah, R. Dingledine, and I. Goldberg, "Changing of the guards: A framework for understanding and improving entry guard selection in tor," in *Proc. of the 2012 ACM WPES*, 2012.

[6] L. Yang and F. Li, "mTor: A Multipath Tor Routing Beyond Bandwidth Throttling," The University of Kansas, Tech. Rep., 2015.

[7] X. Wang, S. Chen, and S. Jajodia, "Network flow watermarking attack on low-latency anonymous communication systems," in *Security and Privacy, 2007. SP'07. IEEE Symposium on*. IEEE, 2007, pp. 116–130.