Exploiting Distance-Indexed IP Traceback Schemes

Lin Cai Jianping Pan Sherman Shen

Denial-of-Service (DoS) attacks and their distributed variants (DDoS) have become a serious threat to the healthy proliferation and growth of the Internet. Most Internet service providers, including many high-profile ones, have suffered severe DoS/DDoS attacks, and have sustained a considerable loss of service capabilities, customers, and revenues. Besides other causes, the lack of *source accountability* in the TCP/IP protocol stack is a major concern that sometimes even encourages these attacks. Attackers can easily forge their identities (usually the source IP address, protocol identifier, and port number of their outgoing packets) when they have no intention to obtain services from DoS/DDoS victims, but just want to prevent legitimate users from doing so. The so-called *source spoofing* does not affect the destination-oriented Internet routing fabric that successfully transports both attack and legitimate flows to victims.

Source traceability is a victim-oriented approach to achieve source accountability. With the assistance of anomaly and intrusion detection tools, victims or their agents first identify attack flows, and then initiate a request that traces back toward the real sources of these flows. Traceback can occur in higher layers (e.g., by correlating SMTP server signatures in email headers), but the traceability of IP packets is essential, due to the fact that many DoS attacks do not exchange application-layer data at all. However, IP-level traceback is much more challenging, since each IP packet is self-contained and can carry different source identities even from the same attack source. An *ideal* traceback scheme should correlate attack packets efficiently, identify or isolate attack sources effectively, and more importantly, allow an incremental deployment. In addition, such a scheme should be lightweight and only impose minimal changes to existing Internet infrastructures (especially in core routers).

Many IP traceback schemes [1] have been proposed in the last few years. In terms of how traceback characteristics are extracted and where the information is stored, most schemes follow one of the following two approaches: *router stamping* and *packet stamping*, which emulate the traceback techniques in postal and telephone systems, respectively. In router-stamping schemes (e.g., [2]), a router identity (or its fraction) is stored in packets when they travel through routers; victims collect these stamped packets, and after having enough stamps, recover a reverse path toward attack sources, which is identified by the stamps of traversal routers. In packet-stamping schemes (e.g., [3]), routers keep a copy (or a digest) of forwarded packets for a while; victims should initiate a traceback request within a certain time-period, which is facilitated by a traceback authority consulting routers that still have the matching packet stamps. In general, packet stamping incurs higher computation and storage overhead in routers; therefore, router stamping appears more attractive for IP traceback schemes.

In this poster, we focus on the router-stamping approach. We first discover an overlooked *buffer overflow* vulnerability that is intrinsic to many distance-indexed IP traceback schemes. When there is limited space in packets to store stamps, traversal routers only stamp packets with a given probability, in order to preserve the already-inscribed stamps made by upstream stamping routers. Also, these stamps are indexed by the incremented distance to the packet destination, in order to allow victims to *independently* recover a *reverse* path from the destination in a *hop-by-hop* manner. To promote an incremental and favorable deployment, these schemes have to follow open protocols and adopt well-known parameters, which are also available to attackers. Substantiated by extensive efficacy analysis and numerical results, we design *extension*, *split*, *branch*, and *synthesized* exploits that can actually take advantage of the buffer overflow vulnerability by creating different types of forged reverse paths in a very efficient manner when compared with the traceback effort attempted by victims. These forged paths are statistically indistinguishable from the genuine ones; in addition, with unconstrained exploits, even genuine paths appear to be forged from the viewpoint of victims. Consequently, we show that the design goal of these traceback schemes can be seriously compromised in practice.

Moreover, we discuss the distance-related vulnerability in a general context relevant to network protocols, and examine a few possible alternatives. We also point out that the overflow vulnerability in distance-indexed IP traceback schemes cannot be effectively eliminated by some *quick fixes* (e.g., dropping or flagging overflowed packets) without interfering with their stateless, low-overhead, and incrementally-deployable design, *unless* the distance increment procedure becomes overflow-resistant. We still agree that IP-level traceback schemes are essential and promising to circumvent ever-increasing DoS/DDoS attacks, if these schemes are properly designed, developed, and deployed; therefore, it is vitally important to identify and understand their intrinsic weaknesses, if any, at the earliest possible stage. This poster is an attempt to serve this purpose.

L. Cai and S. Shen are with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, Ontario N2L 3G1, Canada. E-Mail: {cai,xshen}@bbcr.uwaterloo.ca. J. Pan is now with NTT MCL, Palo Alto, CA, USA. E-Mail: panjianping@acm.org.

REFERENCES

[1] A. Belenky and N. Ansari. On IP traceback. *IEEE Communications Magazine*, 41(7):142–153, 2003.

[2] S. Savage, D. Wetherall, A. Karlin, and T. Anderson, "Practical network support for IP traceback," *Proc. of 16th ACM SIGCOMM (SIGCOMM'2000)*, pp. 295–306, 2000.

[3] A. Snoeren, C. Partridge, L. Sanchez, and C. Jones, "Hash-based IP traceback," *Proc. of 17th ACM SIGCOMM (SIGCOMM'01)*, pp. 3–14, 2001.