Kevin Bauer
CSCI 7000: Doctoral Level Independent Study
November 27, 2006

“Envisioned” Ph.D. Thesis

A Secure and Scalable Peer-to-Peer Privacy Enhancing System

Anonymous communication systems have been vitally important to protecting end-user’s privacy while using the Internet. These systems are essential to journalists reporting from countries where the freedom of press is not guaranteed, various law enforcement activities, and individuals concerned about personal privacy while online. There are currently several popular systems such as Tor (The Onion Router), Tarzan, and MorphMix that attempt to provide secure anonymous communication systems. These systems provide a form of a “mix-network” in which the source of a traffic stream is obfuscated by routing the packets through a random and arbitrary set of intermediate nodes. Ideally, to a passive observer capturing traffic at the destination server, the true source of the traffic is unknown.

Although privacy enhancing systems are popular among Internet users, their development still remains in its infancy. There remain several challenges that limit their ability to provide a high degree of anonymity. Specifically, these systems have difficulty scaling to large networks and often even their most fundamental security goals are not achieved. In order to motivate the need for improved anonymous communication systems, consider Tor, perhaps the most popular anonymous communication system available for anonymizing TCP traffic. The Tor system architecture is composed of “onion routers” that are responsible for transporting an end-user’s traffic through the Tor network and on to its final destination server. The information about the onion routers is maintained by a set of directory servers. Since it is necessary for each end-user’s Tor client to keep track of the list of all onion routers, the system cannot scale to a large number of onion routers. When an end-user wishes to begin an anonymous communication session, their Tor client contacts the directory server to retrieve the list of available onion routers. Once the list has been received, the client chooses a path of precisely three onion routers through which to forward its traffic to its desired destination server. In order to provide a low-latency service, the onion routers are chosen based upon a probabilistic weighting scheme by which those routers that claim to offer the highest performance (in terms of bandwidth) are chosen most frequently. Since the directory servers have no mechanism to verify bandwidth claims, all advertisements are assumed to be legitimate. Therefore, a malicious user can, with very little effort, send false resource advertisements to the directory servers and thereby appear on a disproportionately high number of paths through the Tor network. A group of colluding nodes can position themselves strategically on these paths and compromise the anonymity of a high percentage of users in the network.

The goal of this “envisioned” thesis is to provide a solution to both the scalability and security challenges associated with anonymous communication systems in large-scale distributed systems. A new privacy enhancing system is introduced that provides a high-throughput, low-latency service and does not rely upon centralized directory servers for distributing peer router information to clients. “Gossip” protocols are used to share information about peer routers in the system, allowing this information to propagate very quickly to all nodes in the network. Furthermore, this system utilizes third-party verification for all resource claims that are to be propagated through the network. This mitigates the ability of malicious nodes influencing the system’s path selection process.