*Phase 1: Shrew Attack*

**Summary of How Shrew Attacks Work**

The shrew attack is a low-rate DoS attack that was created to avoid some of the weaknesses of the brute-force DDoS methods. The shrew attack works by taking advantage of TCP's retransmission timeout (RTO) transmission; when the client sends requests to the server the attacker sends a burst of requests at the same time to a bottleneck router, causing the router to suspend data transmission. Packets are dropped during RTO (RFC 2988 recommends minRTO to be 1 second[1]). After minRTO the client retransmits the lost packets, considerably slowing TCP transmission. By using this method the attacker uses minimal resources (compared to DDoS) to reduce costs and effort as well as making it harder to detect as an anomaly. See Figure 1 for an illustration of the shrew attack model.

**Fulfillment of Assignment's Requirements**

The shrew attack can be considered a man-in-the-middle attack because the attacker has to monitor the client's transmissions to the server in order to coordinate his/her attack as short on/off bursts at the same time as the client's communications. This attack does not require the client's packets to be modified; it only drops/delays the packets at the bottleneck. However, traffic is not completely blocked since bottlenecks will allow requests through eventually, just at a slower rate with increased requests.

**Details of Shrew Attacks**

The trade-off shrew attacks make is choosing to minimize attack costs at the expense of the throughput degradation which could be achieved by a DDoS. Sometimes the bottleneck is large enough to accommodate any attempts by the attacker to create overflow, making the shrew attack ineffective. In this section we will examine the details of shrew attacks.

When a TCP client ACKs packets from the server, the server increments the congestion window (cwnd, which dictates the transmission rate of the server) and the server sends the next packet. A client can do an optimistic ACK, meaning they ACK packets before they are received and can therefore get the next packet faster than a normal client. In this attack, the attacker dictates all levels of traffic by sending batches of optimistic ACKs (see Figure 2 and 3 for flood cycles). By controlling the traffic flow, the attacker can control the server's cwnd and RTO.

During high traffic congestion the client's legitimate requests compete with the attacker's packets, filling the buffer. Some of the client's requests will be transmitted successfully, but once the buffer is filled all TCP packets are dropped for the duration of RTO. After the duration of RTO, the server exponentially increases its cwnd until the server can handle half the buffer capacity (also called the slow-start threshold). After that, the cwnd is linearly increased until it reaches the router buffer capacity.

[1] V. A. Kumar, P. S. Jayalekshmy, G. K. Patra and R. P. Thangavelu, "On remote exploitation of TCP sender for low-rate flooding denial-of-service attack"

**Tables and Figures**

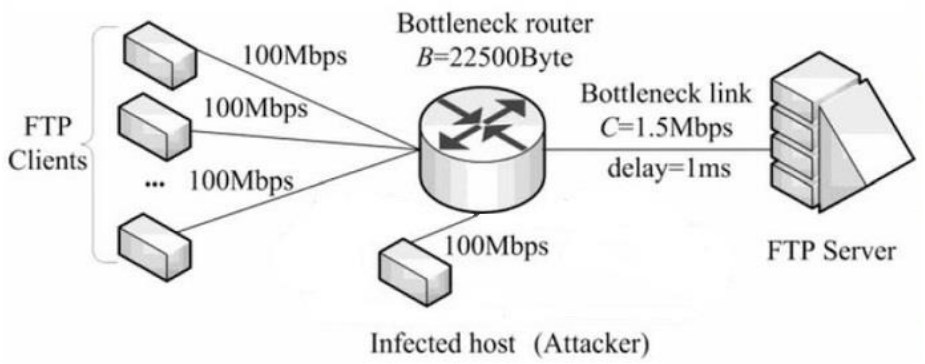


**Figure 1**: Shrew attack

$t_{short} = C/(R_{high} - R_{out})$
$N_p = R_{high} \times t_{short} = C/(1 - R_{out} / R_{high})$
**Figure 2**: Equation for determining the amount of packets needed to fill the buffer

"Each Shrew flood cycle consists of three different traffic rates of different duration: 1) a *highrate traffic* $(R_{high})$ for a *short-duration* $(t_{short})$ to fill the router buffer, 2) a *low-rate traffic* $(R_{low})$ for a relatively *longer duration* $(t_{long})$ to keep the buffer filled, and 3) a *null traffic* $(R_{null})$ with *duration* $(t_{null})$ adjusted in such a way that the *flood period* $(T = t_{short} + t_{long} + t_{null})$ is close to the minimum RTO of targeted TCP flows. Denote $C$ as the router buffer capacity in packets. If $R_{high}$ and $R_{out}$ are the packet arrival and departure rates at the router, then the time, $t_{short}$, *and* the no. of packets, $N_p$, required to fill the buffer."
V. A. Kumar, P. S. Jayalekshmy, G. K. Patra and R. P. Thangavelu, "On remote exploitation of TCP sender for low-rate flooding denial-of-service attack"

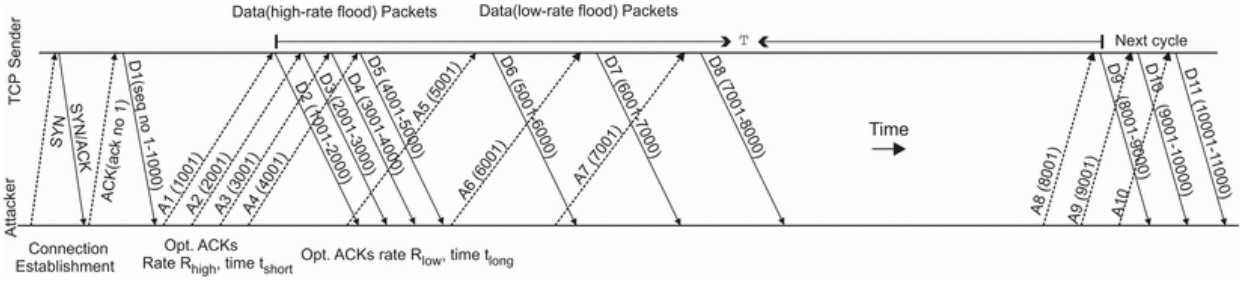


**Figure 3**: A shrew attacker sending optimistic ACKs to a TCP server
V. A. Kumar, P. S. Jayalekshmy, G. K. Patra and R. P. Thangavelu, "On remote exploitation of TCP sender for low-rate flooding denial-of-service attack"

**References**

J. Luo and X. Yang, "The NewShrew attack: A new type of low-rate TCP-Targeted DoS attack," *2014 IEEE International Conference on Communications (ICC)*, Sydney, NSW, 2014, pp. 713-718. doi: 10.1109/ICC.2014.6883403
URL: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6883403&isnumber=6883277

M. Yue, Z. Wu and M. Wang, "A New Exploration of FB-Shrew Attack," in *IEEE Communications Letters*, vol. 20, no. 10, pp. 1987-1990, Oct. 2016.
doi: 10.1109/LCOMM.2016.2596278
URL: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7524704&isnumber=7586084

V. A. Kumar, P. S. Jayalekshmy, G. K. Patra and R. P. Thangavelu, "On remote exploitation of TCP sender for low-rate flooding denial-of-service attack," in *IEEE Communications Letters*, vol. 13, no. 1, pp. 46-48, January 2009.
doi: 10.1109/LCOMM.2009.081555
URL: http://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=4753961&isnumber=4753944