Joseph Idziorek
Ph.D.
Computer Engineering
Major Professor: Doug Jacobson

Exploiting Cloud Utility Models for Profit and Ruin

A key characteristic that has led to the early adoption of public cloud computing is the utility pricing model that governs the cost of computing resources consumed. Similar to public utilities like gas and electricity, cloud consumers only pay for the resources they consume and only for the time they are utilized. As a result and pursuant to a Cloud Service Provider's Terms of Agreement, cloud consumers are responsible for all computational costs incurred within and in support of their rented compute environments whether these resources were consumed in good faith or not. While initial threat modeling and security research on the public cloud model has primarily focused on the confidentiality and integrity of data transferred, processed, and stored in the cloud, little attention has been paid to the external threat sources that have the capability to affect the financial viability of cloud-hosted services.

Bounded by a utility pricing model, Internet-facing web resources hosted in the cloud are vulnerable to Fraudulent Resource Consumption (FRC) attacks. Unlike an application-layer DDoS attack that consumes resources with the goal of disrupting short-term availability, a FRC attack is a considerably more subtle attack that instead targets the utility model over an extended time period. By fraudulently consuming web resources in sufficient volume (i.e. data transferred out of the cloud), an attacker is able to inflict significant fraudulent charges to the victim. This work introduces and thoroughly describes the FRC attack and discusses why current application-layer DDoS mitigation schemes are not applicable to a more subtle attack. The work goes on to propose three detection metrics that together form the criteria for detecting a FRC attack from that of normal web activity and an attribution methodology capable of accurately identifying FRC attack clients. Experimental results based on plausible and challenging attack scenarios show that an attacker, without knowledge of the training web log, has a difficult time mimicking the self-similar and consistent request semantics of normal web activity necessary to carryout a successful FRC attack.