S. Tarzia – EECS 443 Advanced Operating Systems – Paper Summary

Scalability, Fidelity, and Containment in the Potemkin Virtual Honeyfarm
Michael Vrable, Justin Ma, Jay Chen, David Moore, Erik Vandekieft, Alex C. Snoeren, Geoffrey M. Voelker, and Stefan Savage

There a tradeoff between accuracy (called fidelity or interactivity) and scalability in honeyfarms. Potemkin provides high fidelity by running a full VM for each target IP address (honeypot). This paper's main contribution is its achievement of scalability by flash cloning and delta virtualization. Flash cloning creates a new VM quickly by simply loading a memory image. All VMs on a given server will use this same initial memory image representing some fixed combination of OS and services (vulnerabilities). Fast VM creation allows idle IP addresses to be reclaimed (shut down) often without much of a performance penalty. Delta virtualization is a copy-on-write policy for VM state. VM memory is large but in practice most of this will be shared across the VMs on each server (since each started in the same state and is running the same services). Thus, with delta virtualization the memory requirement for each VM is limited to the memory that is modified; in their experiments this value is only a few MB).

The biggest problem with this paper is its dependence on Xen and thus limitation to Linux and BSD; in reality, Windows honeyfarms are much more useful. Also, the reliance on a single Gateway Router for all management tasks is unscalable. I also wonder if starting all of the VMs in the same state (with flash cloning) limits the variety of attacks that you can witness; specifically I am thinking of attacks that are only effective when the machine reaches some specific state. Attack attraction is always a sore spot for honeypots. Particularly, in this paper no detail is given on the type of traffic that the honeyfarm was exposed to (but I don't know how much you can really say about this).

I expect future work to address the 10MB Xen metadata size limit that was limiting their implementation to 116 VMs; they claimed that without this obstacle about 1500 VMs should be supportable on each server. What can we gain from massive-scale accurate honeyfarms? It would be nice to see some real results showing how Potemkin can contribute to network security. The authors also list DoS and honeypot detection vulnerabilities as open issues.