Virtual machine introspection in a hybrid honeypot architecture

Tamas K Lengyel & Justin Neumann

University of Connecticut

The role of the honeypot

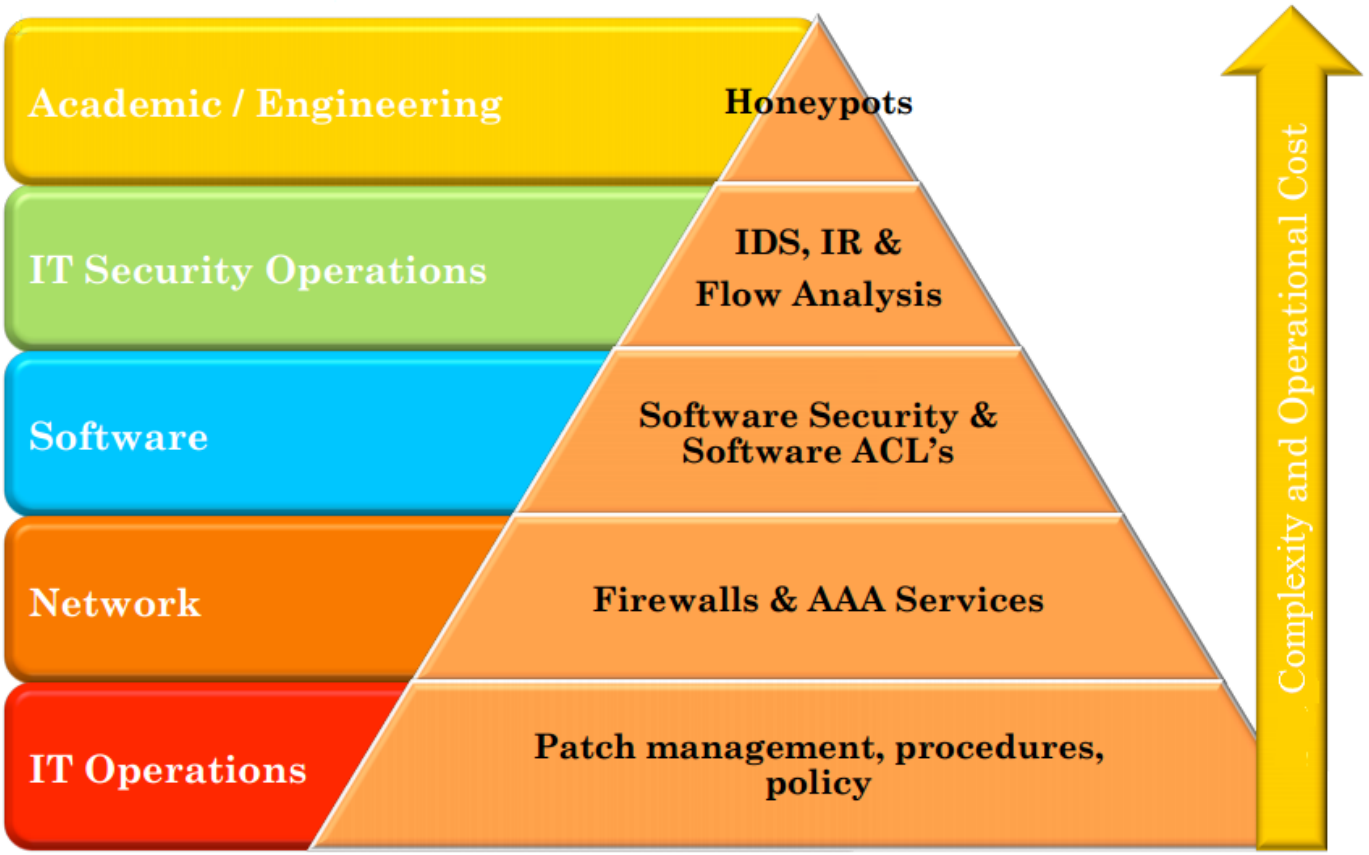

Bruce Potter, 2008, Defcon 16: Network flow analyses

The limitations

Low-interaction honeypots:

- "Artificial" attack surface
- Limited information about the attacks

High-interaction honeypots:

- Complexity
- Maintenance
- High risk

Hybrid honeypots

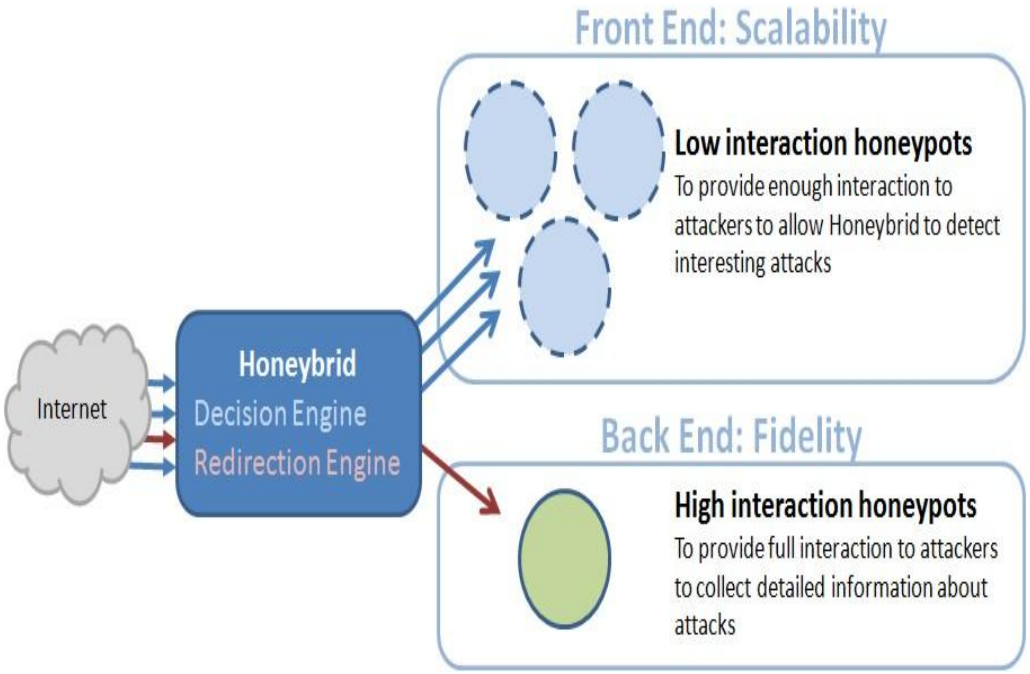

Robin Berthier, 2006: Advanced honeypot architecture for network threats quantification

Theory: Combining low and high interaction honeypots can provide the best of the two.

Original idea: switch an attack to a high-interaction honeypot based on predefined rules

Problem: What rules?

Further problems

Few choices for high-interaction honeypots

- Sebek
- Qebek
- Argos

Why?

"Regarding Reviewer #4's question as to whether we would consider releasing gateway and containment server code to the community, we indeed considered this. However, in our experience malware execution platforms differ substantially, and it would likely be hard to make our code work in a variety of environments. In addition, we lack the support to commit to the maintenance necessary for such a public release to be effective."

Kreibich et. al., SIGCOMM 2011: GQ: Practical Containment for Measuring Modern Malware Systems

Further problems

Virtualization based honeypots:

- Modified QEMU
- Malware can detect *monitoring* and alter its behaviour
- Most only work with Windows XP SP2

VMI-Honeymon http://vmi-honeymon.sf.net

- Built on open source tools
  - LibVMI
  - LibVirt
  - LibGuestFS
  - Volatility
  - Xen
- Full virtualization, no modification to Xen
- Works with ***all*** versions of Windows with no in-guest agent
- Read-only memory scanning and footprinting eliminates subversion attacks

System overview

- Honeybrid filters attackers who already dropped payload on Dionaea
- Only one attacker interacts with the HIH at a time
- An attack is transferred to the HIH when it is free (random samples)

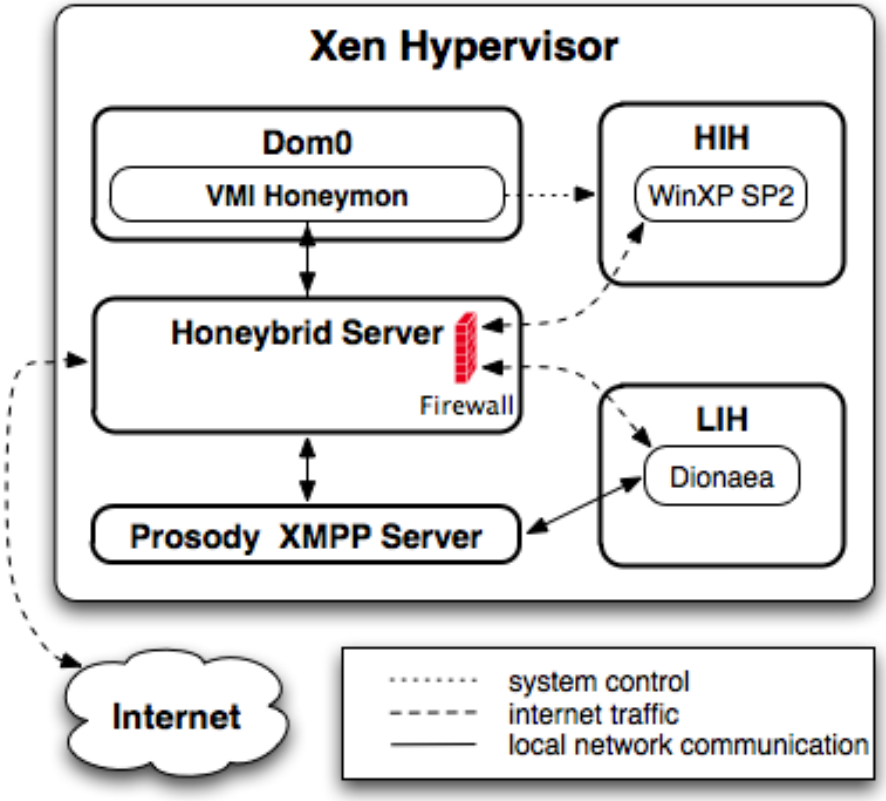

System overview

- Honeybrid detects outgoing connections from HIH, sends trigger to VMI-Honeymon
- On time-out Honeybrid sends trigger to VMI-Honeymon
- After attack session, HIH is reverted

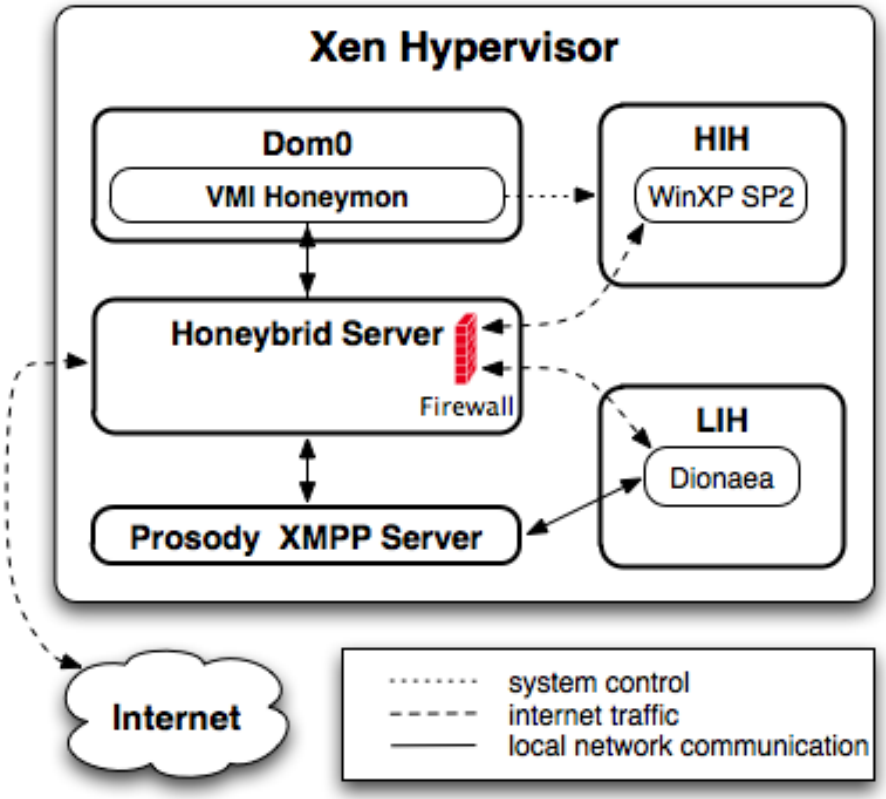

Results (in two weeks)

VMI-Honeymon: 886 binaries (6,335 TCP sessions)
Dionaea: 1,411 binaries (1,152,142 TCP sessions)

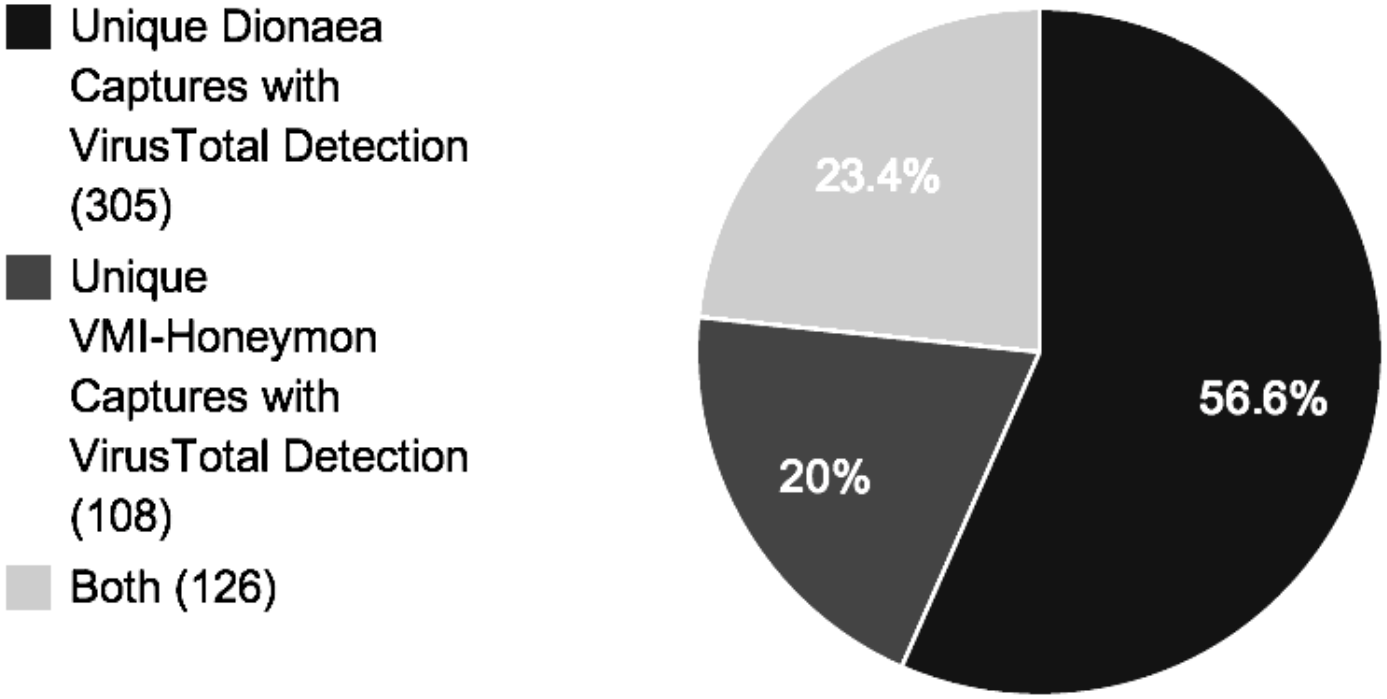

Performance

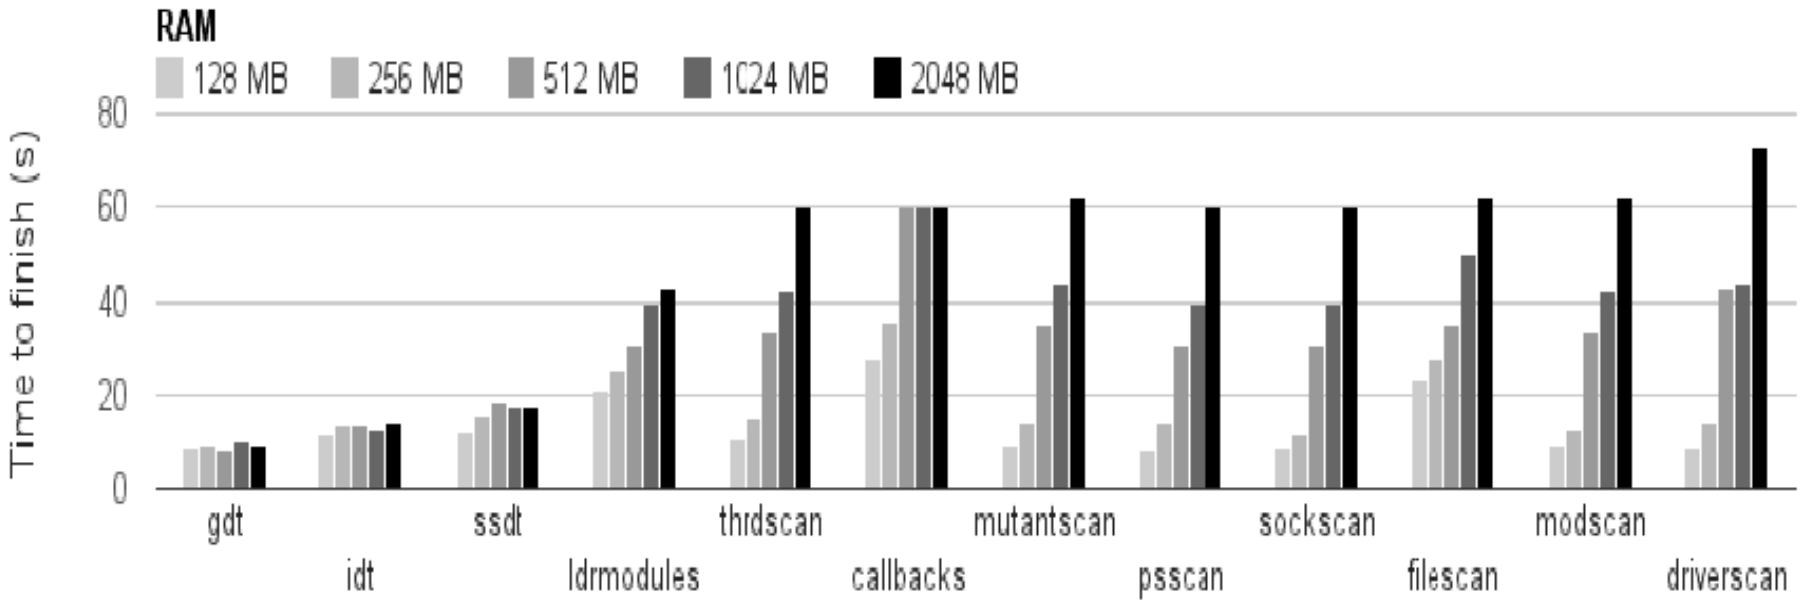

Future work

- Multiple concurrent HIHs
- Using Windows Vista, 7 and 8 as HIH
- Fast-clone/memory sharing of HIHs
- Automatic analyses of malware memory footprints to detect similarities

Thank you!