

Disk-Level Malware Detection



Nathanael Paul, Adrienne Felt, David Evans, Sudhanva Gurumurthi

University of Virginia Computer Science

<http://www.cs.virginia.edu/malware>

The processing capabilities of modern disk drives can be used to make malware detection more precise and efficient, and harder to circumvent.

Modern disk drive processors are now capable of general purpose computation, and we can harness this new power to implement malware detection directly on the disk drive.

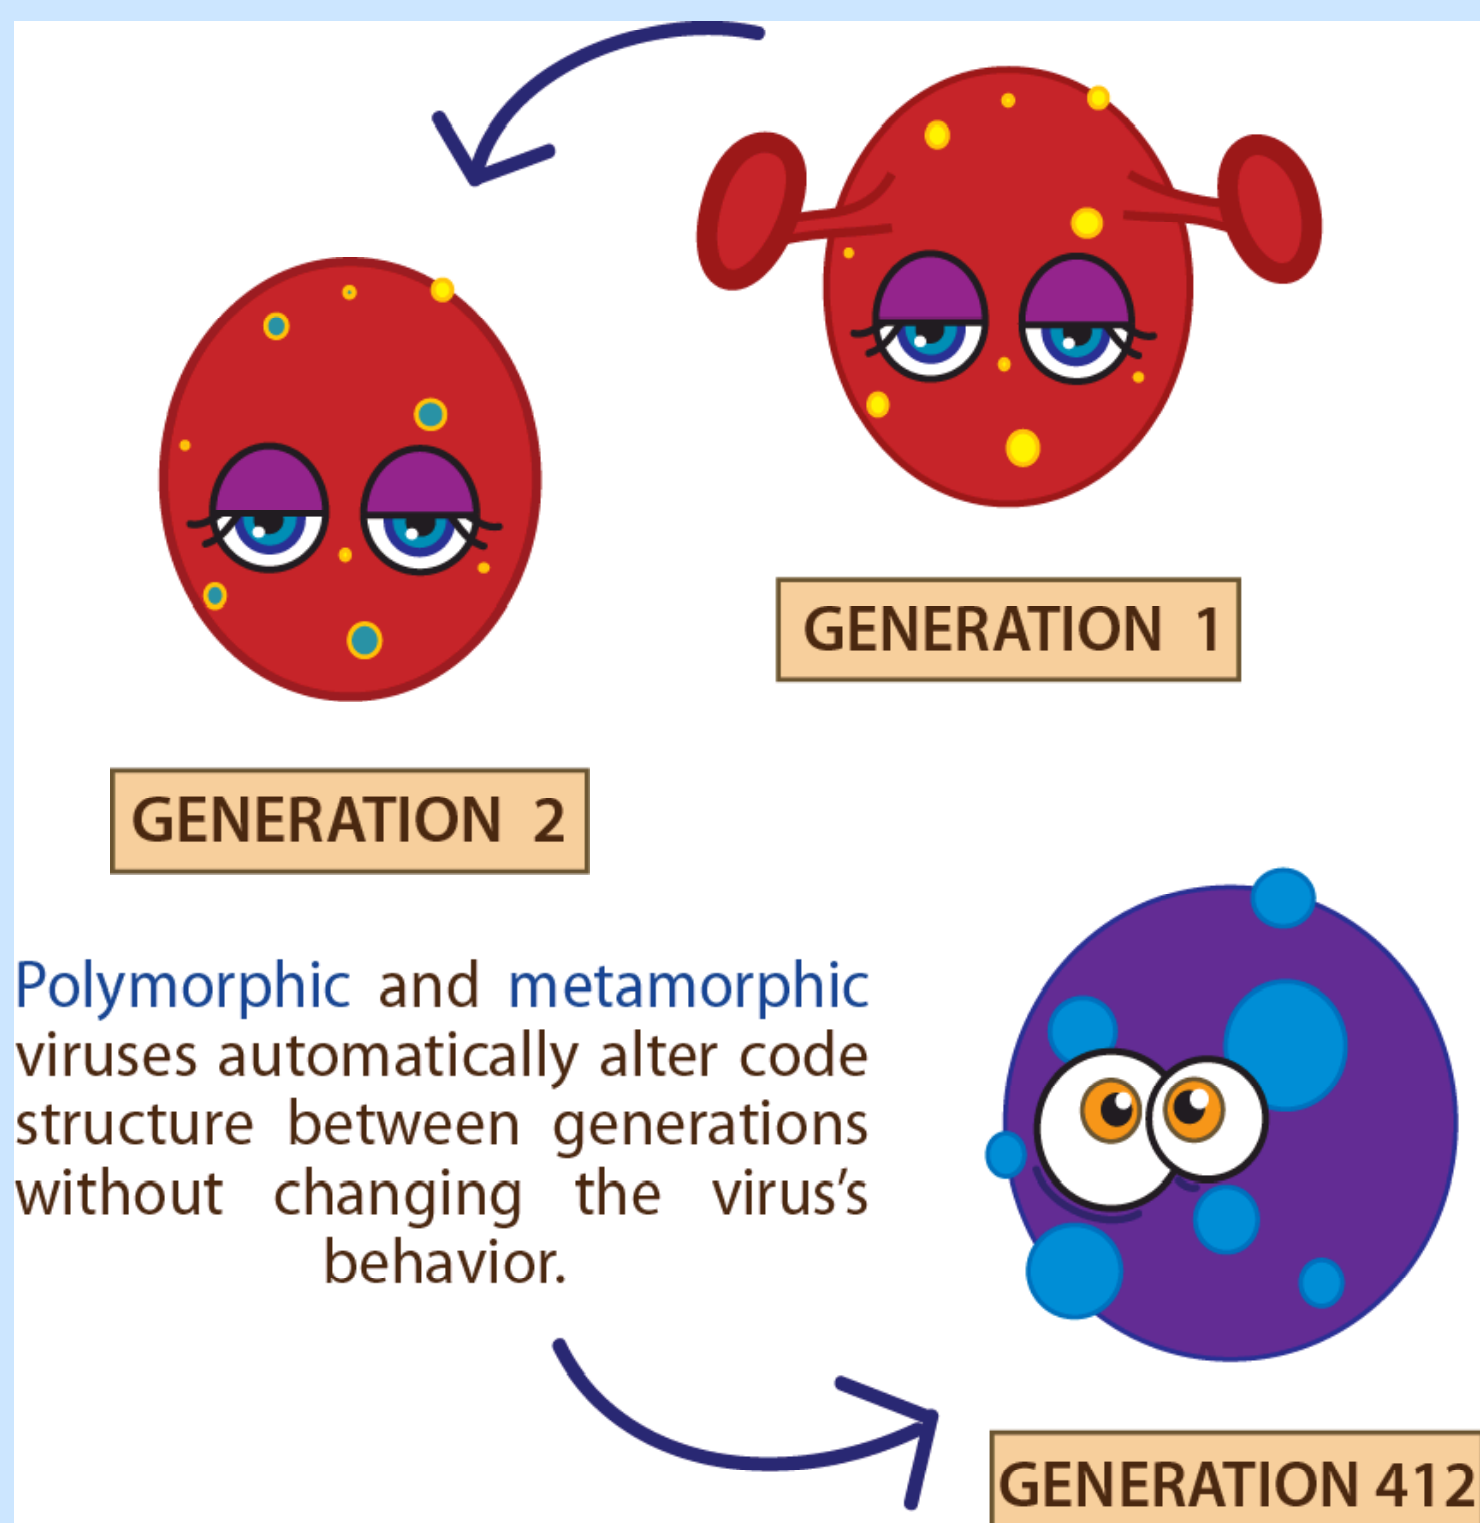
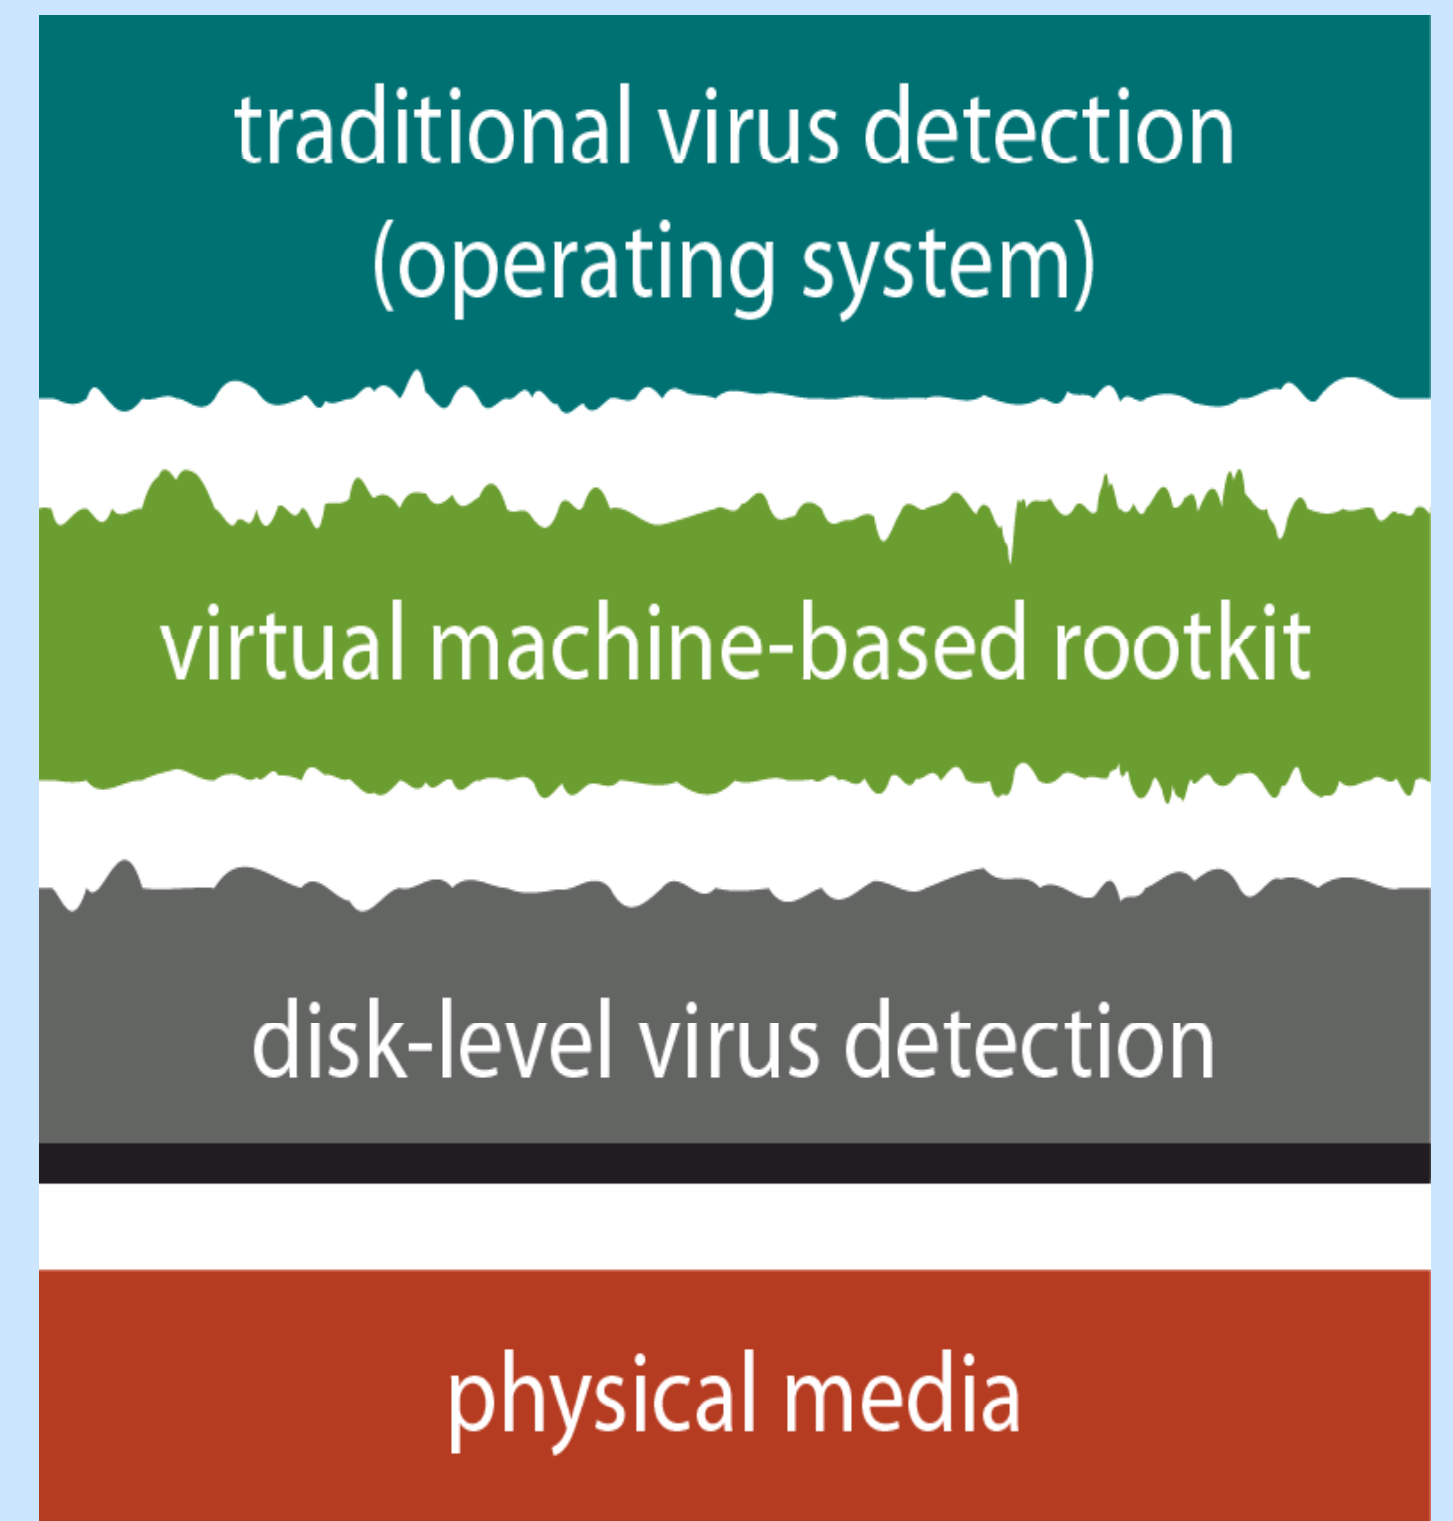
Disk-level malware detection offers these advantages:

Location: Below the host OS, sees all disk requests.

Isolation: Can operate while host may be compromised.

Difficult to evade: Hard to change disk requests without changing behavior.

Low overhead: Speed gap between the disk processor and mechanical data transfer system.



Limits of Current Approaches

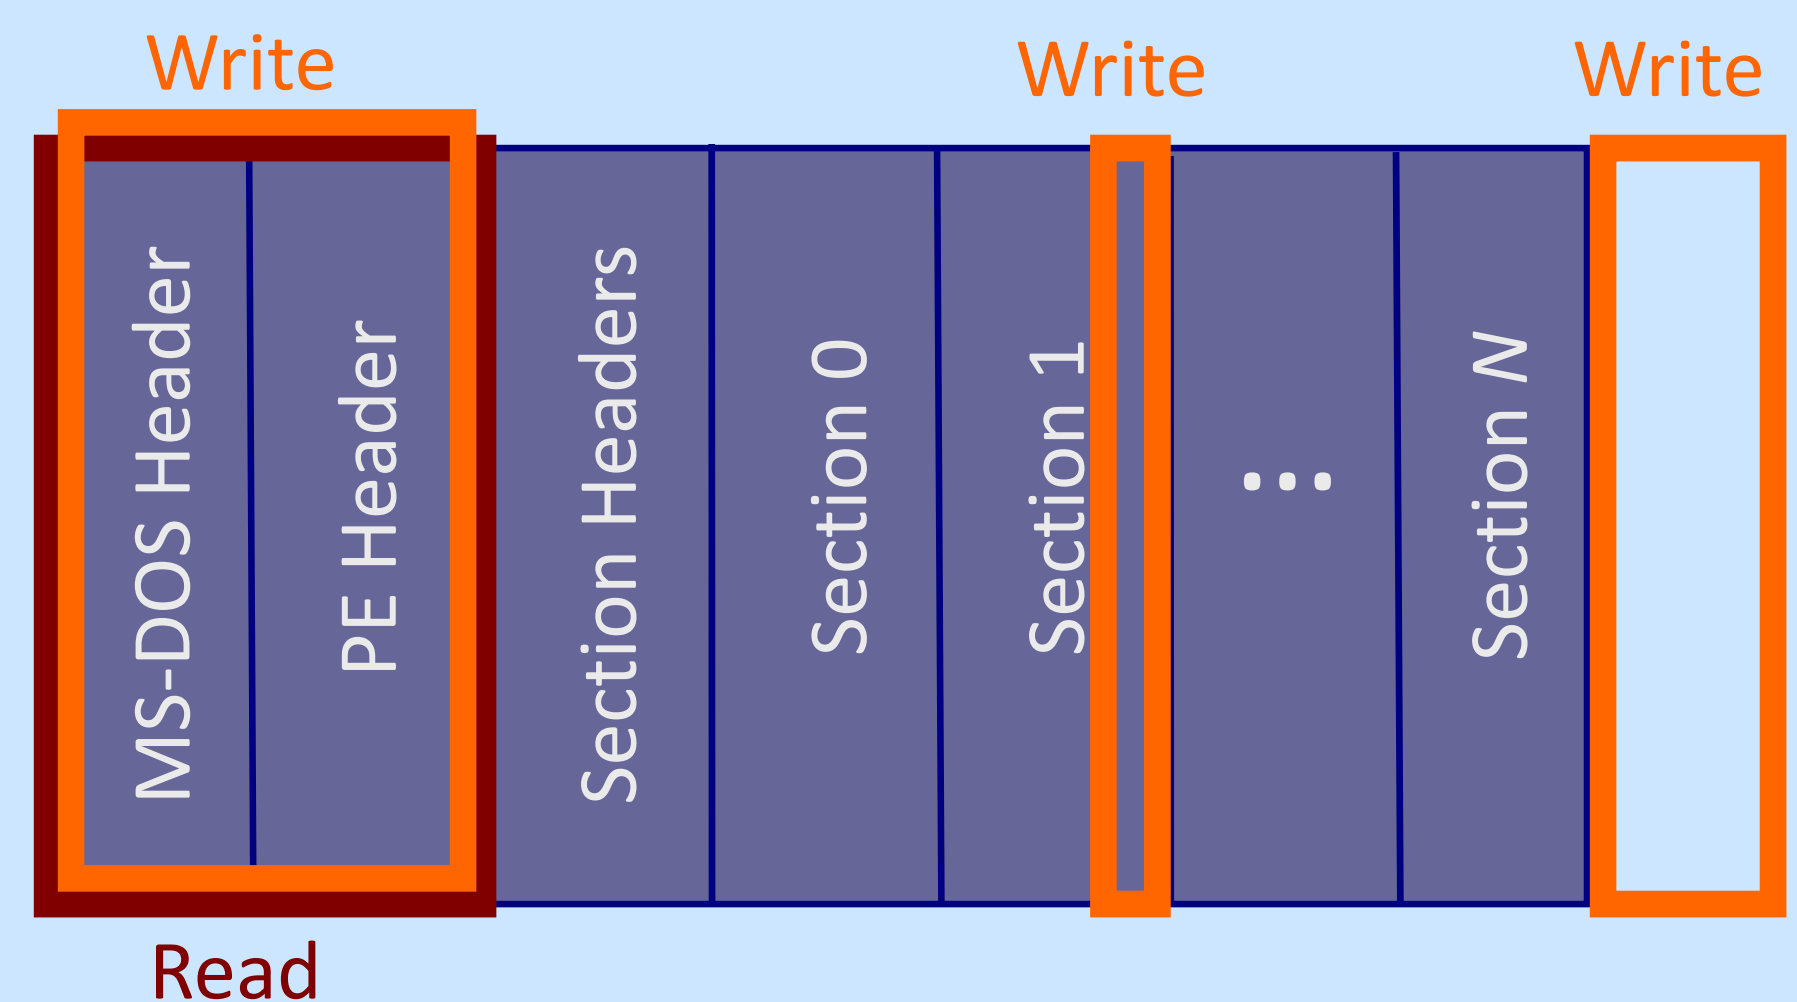
String scanning is the traditional - and primary - method of virus detection. Polymorphic and metamorphic viruses elude these detectors. It is easy to hand-craft variants that evade detection.

Emulation was designed in response to these complex viruses. Emulation is limited by its high computational cost and imprecision. Virus authors subsequently created anti-emulation techniques.

Rootkits are rapidly becoming serious threats. Virtual machine-based rootkits are nearly impossible to detect using host-level techniques because they run below detectors.

Generic Detection Rules

We have developed generic rules that detect viruses that infect Windows PE executable files. These rules are based on the fundamental disk-level activity typically used to infect an executable, which involves reading the file header and updating the header and other parts of the file.



Sample Detection Results

Virus	Generic Detection Rule			
	W	R_0W_0	RW_0W	R_0RW_0W
Adson.1559	5/5	4/5	4/5	4/5
Alma.2414	8/8	3/8	3/8	2/8
Belial.2609	12/12	12/12	12/12	12/12
Bika.1906	7/7	6/7	6/7	4/7
Champ.5714	7/7	0/7	0/7	0/7
Chiton.b	7/7	6/7	6/7	6/7
Chiton.r	40/40	39/40	39/40	39/40

We tested a random selection of malware from the VXHeavens (March 2007) repository using VMWare and our traced to log disk activity. If the malware added, deleted, or changed an executable file, we saved a log of the disk activity to scan. All traditional file-infectors are detected. For Alma.2414 and Champ.5714, the W rule matches a Windows anti-malware replacement of the protected executables. Chiton.b and Chiton.r each drop a malicious file that is detected by the W rule only.

False Positives

Various non-malicious activities exhibit disk-level behaviors similar to viruses. We evaluated the causes and frequency of these false positives by tracing disk activity for 8 users over several months.

Causes	W	R_0W_0	R_0W_0W	R_0RW_0W
Updates	73	0	0	0
Software Development	2	2	2	1
System Restores	33	33	13	2
Installations	10	0	0	0

All of the false positives are caused by a small number of particular activities. We have designed mechanisms to support these activities without generating false positives.