JOT
JOURNAL OF
OBJECT TECHNOLOGY



# Countering Anti-Debugging Techniques: Enhancing Transparency in Nested Virtualization using HyperDbg (Tool Demonstration)

**Anonymous Authors**
Anonymous Affiliation

**ABSTRACT** Modern malware increasingly employs sophisticated anti-debugging and anti-virtualization techniques to evade analysis, particularly targeting artifacts left by virtualization and nested virtualization environments such as VMware Workstation, Hyper-V, and KVM. HyperDbg, an open-source hypervisor-level debugger, introduces advanced mechanisms to mitigate both its own hypervisor footprints and those of the underlying nested virtualization stack. In this talk, we demonstrate the capabilities of adding a transparency layer on top of the HyperDbg debugger to detect, mitigate, and bypass common and advanced anti-debugging methods leveraged against such environments. Although achieving complete transparency remains infeasible, it significantly raises the bar for malware attempting to detect analysis environments, making evasion substantially more difficult. We further highlight the critical importance of these techniques in practical malware analysis workflows, particularly in scenarios involving snapshot restoration for analyzing and debugging internal malware behavior. By reducing observable artifacts, HyperDbg enhances the reliability of snapshot-based analysis and debugging, allowing researchers to stealthily investigate and understand the inner workings of evasive malware without premature detection or execution of anti-analysis payloads.

**KEYWORDS** Anti-Debugging, Anti-Virtualization, Nested Virtualization, Debugging Malware, Binary Analysis.

## Extended Abstract

Dynamic malware analysis and debugging are critical steps in understanding the behavior of sophisticated threats. Unlike static analysis, which can be hindered by heavy obfuscation, packing, or encryption, dynamic analysis enables researchers to observe real-time execution, track control flow, and extract hidden payloads. While static analysis tools like IDA Pro, Ghidra, Binary Ninja, and Radare are invaluable for reverse engineering, they become significantly less effective when dealing with packed or heavily obfuscated malware. In such cases, a debugger becomes the primary tool for malware analysts to unpack, deobfuscate, or automate the analysis of the malware's behavior.

Since regular user-mode or kernel-mode debuggers can leave detectable artifacts, hypervisor-based debuggers are often preferred for stealthier analysis. Hypervisor-level debuggers not only minimize their footprint but also provide system-level access, allowing observation of behaviors that occur even above the operating system kernel by employing hypervisor-level capabilities.

Hypervisor-based debuggers, among their benefits, come with certain artifacts that expose the presence of the hypervisor. This problem becomes even more serious in nested virtualization environments, where clues left by both the main hypervisor and the nested one can be used by malware to detect that it is being watched. Small differences in CPU behavior, hardware details, timing measurements, or system information can reveal that the malware is running inside a debugger or a virtual machine. Once detected, malware might shut itself down early or trigger fake behavior to mislead the analyst.

**JOT reference format:**
Anonymous Authors. *Countering Anti-Debugging Techniques: Enhancing Transparency in Nested Virtualization using HyperDbg (Tool Demonstration).* Journal of Object Technology. Vol. vv, No. nn, yyyy.  http://dx.doi.org/10.5381/jot.yyyy.vv.nn.aa

To address these issues, we present a transparency layer over the HyperDbg (Karvandi et al. 2022) debugger specifically designed to enhance transparency in nested virtualization environments. By not using OS-level debugging APIs, HyperDbg minimizes detectable artifacts at multiple layers, enabling analysts to conduct stealthy, fine-grained debugging sessions even against highly evasive malware. By mitigating both traditional anti-debugging techniques and emerging nested virtualization detection methods, HyperDbg significantly improves the reliability and effectiveness of dynamic malware analysis workflows.

**VM Detection Techniques and Mitigations:** Malware can detect virtualized environments using a broad range of heuristics. The transparency layer offers or facilitates mitigations against multiple categories of these techniques:

**CPU and Hypervisor Detection:** Many detection methods query CPU features to find virtualization artifacts. These include checking the hypervisor vendor ID (via `CPUID` instructions), looking for anomalies in CPU brand strings, hypervisor flags, and reserved CPUID leaves. Specific signatures associated with platforms like KVM or Intel's KGT branch can also be probed. Mitigations typically involve masking or spoofing CPUID (using nested VM-exits) to hide hypervisor presence.

**Timing Analysis:** Virtual machines often introduce timing inconsistencies due to emulation overheads or event handling delays. Malware may measure execution timings to infer virtualization. The timing resources include TSC (Time Stamp Counter) or system-wide resources like external devices, HPET (High Precision Event Timer), Performance Counters, APIC (Advanced Programmable Interrupt Controller) Timers, Windows APIs, or timing thread approaches (Schwarz et al. 2017).

Mitigations involve introducing random jitter or normalizing timing outputs to appear more like bare-metal systems by employing VM-exits for TSC (`RDTSC/RDTSCP`), PMC (`RDPMC`), Different MSR registers using MSR Bitmap, monitoring for system-calls querying timers, virtualizing APIC, or monitoring for PIO/MMIO ranges of external devices. Timing-based detection methods are less common in modern malware, because Windows 11 enforces Virtualization-Based Security (VBS) by default. As a result, malware must assume it is running inside a hypervisor to avoid triggering false positives.

**Windows-Specific Detection:** Malware may search for VM indicators specific to Windows environments, such as `MSSMBIOS` registry keys, loaded DLLs, known registry values, or system mutexes.

Mitigations involve monitoring system-calls that query the information related to VM and spoofing registry/file entries.

**CPU and Hardware Analysis:** By examining hardware characteristics, malware can identify virtualized environments. Checks include inspecting processor and core counts, temperature sensors, VHD boot status, thread numbers, driver names, and hypervisor memory pools.

Mitigations involve adjusting VM configurations to emulate realistic hardware profiles.

**Network Analysis:** VM network adapters often use MAC addresses tied to virtualization vendors. Malware can detect these known prefixes or look for VirtualBox network drivers. Mitigations include randomizing MAC addresses and renaming network drivers.

**Hardware Information:** VMs may have default or invalid chassis, device, and firmware information, such as missing thermal sensors, suspicious PCI bridge names, or generic BIOS details. Malware can also scan GPU capabilities and strings to detect virtualization. Mitigations include modifying DMI tables, spoofing firmware signatures, and emulating realistic hardware characteristics (e.g., intercepting I/O ports querying PCIe to spoof vendor ID).

**Instruction Set Analysis:** Certain CPU instructions, such as SIDT, SGDT, SLDT, or querying system registers can behave differently in virtual environments. Malware can use these differences to detect VMs. Some techniques rely on specific features, such as VMware I/O port backdoor. Mitigations involve emulating native CPU instruction behavior and hiding special I/O ports by intercepting VM-exits related to `IN` and `OUT` instructions.

**Filesystem and Storage Analysis:** Malware may search for files, drivers, or directories unique to VM platforms like QEMU, KVM, or VirtualBox. Moreover, VMs often use default or small-sized disks. Malware can check disk sizes, disk serial numbers, or memory allocation patterns.

Mitigations requires intercepting system-calls and spoofing filesystem traces.

**Process Analysis:** Running processes related to VM services, such as VMware Tools or QEMU guests, can reveal a virtualized environment. Mitigations involve masking or terminating these processes (e.g., through system-call monitoring).

**Specialized Techniques:** More advanced checks include analyzing CPU thread counts (e.g., detecting odd thread numbers), probing memory regions for VM signatures, or exploiting low-level system features like the OSXSAVE instruction. Mitigations require deep modifications to the hypervisor, and CPU emulation layers to mimic bare-metal behavior accurately.

Our work on top of HyperDbg demonstrates that substantial increases in transparency are achievable even in complex nested virtualization environments. Although perfect invisibility remains unattainable due to the fundamental limitations of software-based virtualization, HyperDbg significantly raises the bar for malware attempting to evade analysis. Its modular and open-source design enables continuous improvement and adaptation to emerging anti-virtualization techniques, making it a valuable tool for security researchers engaged in stealthy malware analysis, debugging, and reverse engineering.

## References


Karvandi, M. S., Gholamrezaei, M., Khalaj Monfared, S., Meghdadizanjani, S., Abbassi, B., Amini, A., … Schwarz, M. (2022). Hyperdbg: Reinventing hardware-assisted debugging. In *Proceedings of the 2022 acm sigsac conference on computer and communications security* (pp. 1709–1723).

Schwarz, M., Weiser, S., Gruss, D., Maurice, C., & Mangard, S. (2017). Malware guard extension: Using sgx to conceal cache attacks. In *Detection of intrusions and malware, and vulnerability assessment: 14th international conference, dimva 2017, bonn, germany, july 6-7, 2017, proceedings 14* (pp. 3–24).