CSAI Foundation | Cloud Security Alliance

DDRop: A Hardware Attack on Confidential Computing's Trust Root

Security Implications of the New Intel TDX and AMD SEV-SNP Memory Interposer Attack

2026-09-16

AI-assisted Rapid Research

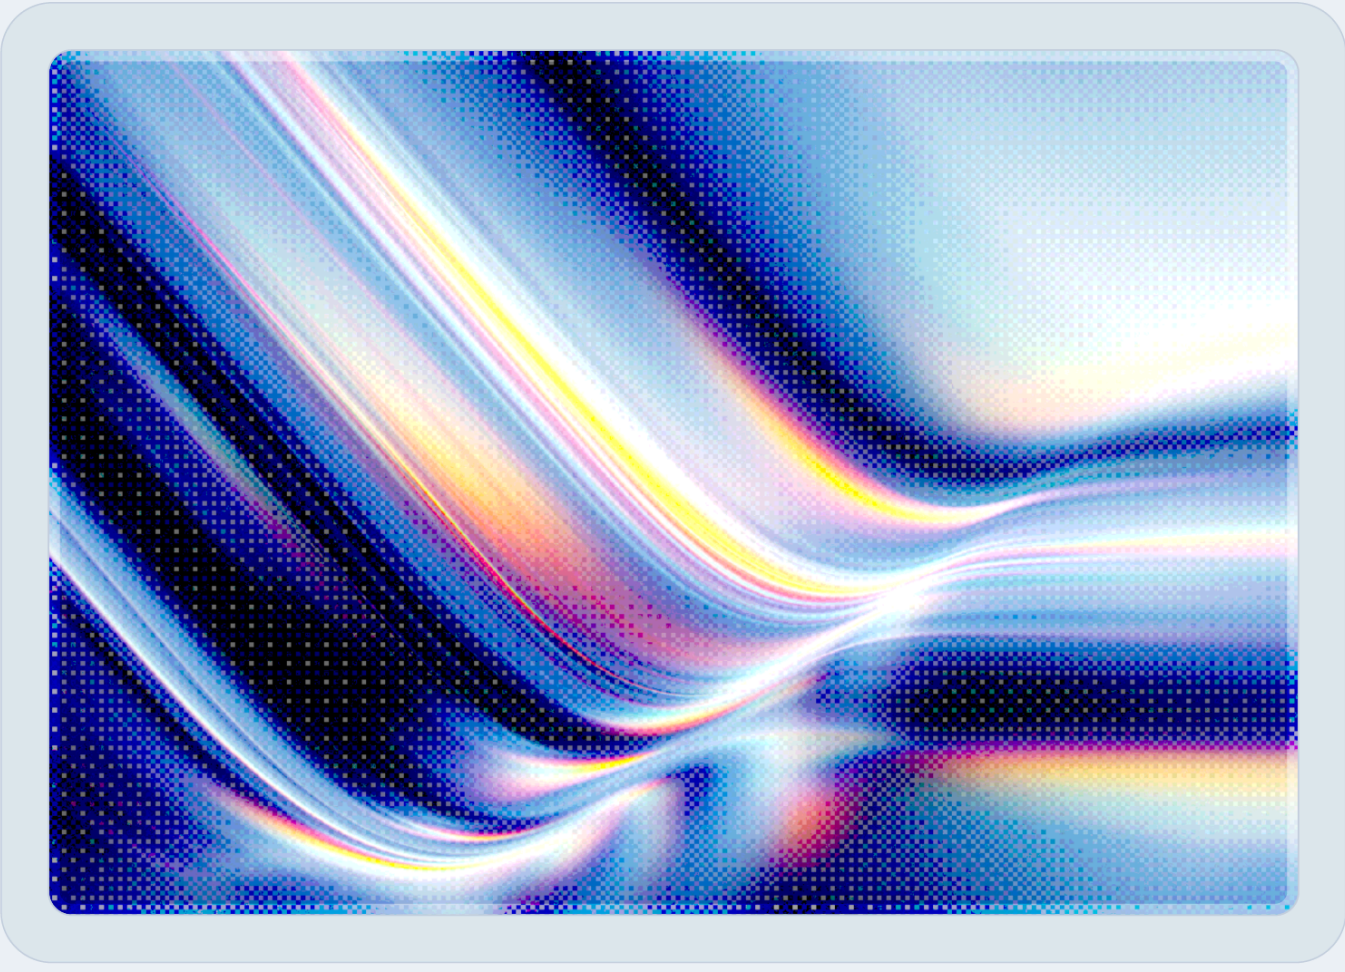

CSAI

CSA cloud security alliance®

Key Takeaways

A research team from KU Leuven, ETH Zurich, Durham University, and Google has publicly disclosed DDRop, a hardware attack that defeats the memory-integrity guarantees of Intel Trusted Domain Extensions (TDX), Intel Scalable SGX, and AMD Secure Encrypted Virtualization-Secure Nested Paging (SEV-SNP) [1][2]. The attack uses a purpose-built DDR5 memory interposer, costing roughly $159 in parts, that sits physically between the processor and a server's memory module and silently discards selected write commands so the processor keeps reading stale, attacker-chosen ciphertext [1][3][9]. On Intel TDX, the researchers demonstrated that dropped writes can suppress page-table initialization, expose protected memory contents, force a confidential virtual machine into debug mode, and forge remote-attestation reports carrying attacker-chosen measurements, all within about two minutes and without crashing the system [3]. Because the technique targets a physical DDR5 bus that memory-encryption schemes were never designed to cryptographically freshness-check, both Intel and AMD have classified it as outside their published threat models, declined to assign CVE identifiers, and offered no software patch [1][3]. The requirement for privileged software access plus brief physical proximity to the server narrows the realistic threat surface to insider risk, hardware supply chain compromise, and shared or co-located physical infrastructure, which, in CSA's assessment, directly challenges the assurance model that cloud providers and regulated customers have used to justify moving sensitive workloads into confidential virtual machines.

Background

Confidential computing extends cloud data protection beyond encryption at rest and in transit to cover data in use, isolating a workload's memory from the hypervisor, co-tenants, and cloud operator staff through hardware-enforced trusted execution environments (TEEs). Intel TDX, AMD SEV-SNP, and their predecessor technologies encrypt guest virtual machine memory with per-VM keys managed inside the processor package, and they pair that encryption with a remote-attestation mechanism that lets a customer cryptographically verify the integrity of the hardware and firmware stack before trusting it with sensitive data. Major cloud platforms, including AWS, Microsoft Azure, and Google Cloud, offer confidential VM instance types built on these technologies, and CSA has observed organizations in healthcare, financial services, and government beginning to treat them as a candidate control for data-in-use protection requirements under frameworks such as HIPAA, GDPR, and the EU's Digital

Operational Resilience Act. CSA's own Confidential Computing Working Group has documented how healthcare organizations increasingly rely on TEEs, hardware security modules, and remote attestation to process protected health information in the cloud while limiting exposure to insiders, co-resident workloads, and memory-scraping tools.

As the DDRop researchers frame it, prior academic research on confidential computing has largely focused on software-level and interrupt-based attacks, such as malicious hypervisor interrupts used to disrupt SEV-SNP guests, or on passive physical attacks that require slowing the memory bus to read data off it [1]. What distinguishes DDRop is that it operates as an active interposer on production-speed DDR5 memory, and that it breaks integrity rather than merely confidentiality: it does not just read data off the bus, it selectively prevents writes from reaching memory, and it does so on a fully patched, current-generation system rather than a legacy or degraded configuration [1]. The researchers behind DDRop presented their findings in a paper titled "DDRop: Active Memory Interposer Attacks on Confidential VMs by Dropping DDR5 Writes," authored by Jesse De Meulemeester, Stefan Gloor, Patrick Jattke, Daniel Moghimi, David Oswald, Martin Thompson, Kaveh Razavi, Ingrid Verbauwhede, and Jo Van Bulck, and coordinated disclosure with Intel and AMD concluded on September 14, 2026, ahead of a scheduled presentation at ACM CCS 2026 in November [1][3][9]. The team has also published the interposer's schematics, board files, and firmware on GitHub, which lowers the barrier for independent replication and, potentially, malicious reuse [10].

Security Analysis

DDRop exploits a gap between what memory encryption on TDX and SEV-SNP actually guarantees and the freshness protection that vendor marketing and compliance narratives have often implied. Both technologies encrypt memory contents so that an attacker who reads raw DRAM sees only ciphertext, but neither scheme, in its currently deployed form, cryptographically verifies that the ciphertext being read back is the most recently written version, a property researchers call freshness. DDRop's interposer, built around a low-cost microcontroller, sits on the memory bus between the CPU and a DDR5 registered DIMM and operates at the bus's native speed, differing from earlier passive-tap research, which could only read data by deliberately slowing the bus down [1][3]. When the interposer wants to drop a specific write, it deliberately induces a parity or command-bus error and then suppresses the alert signal the memory module would normally use to report that error, so the module silently discards the command while the processor and memory controller remain unaware anything went wrong [1][3]. The practical effect is that the protected workload continues operating on old, still-validly-decrypting data exactly when the attacker wants it to.

On Intel TDX, selectively dropping writes during page-table initialization let the researchers inject page-table entries of their choosing into a confidential VM's address space, which in turn let them access ciphertext that should have remained isolated, place the trust domain into debug mode, and forge attestation reports containing attacker-selected measurements, effectively defeating the mechanism a relying party would use to verify the VM's integrity before trusting it [1][3]. The debug-mode and attestation-forgery results are specific to TDX's current architecture. On AMD SEV-SNP, the team achieved a related but distinct result: by dropping writes during AMD's page-relocation feature, they were able to copy the contents of one victim memory page into another, an integrity bypass that undermines SEV-SNP's memory-integrity claims even though it did not extend to debug-mode activation or attestation forgery in the disclosed research [1]. Across both platforms, the attack requires an adversary with privileged software access to the target system and brief physical access to install the interposer, for example during a maintenance window or through a compromised or malicious insider with data-center access; it cannot be carried out remotely under the disclosed technique [1][3].

CSA reads the vendor responses as reflecting a disagreement about scope rather than severity: Intel's advisory, issued September 14, 2026, acknowledges the research but states that physical attacks against the DDR5 memory bus fall outside the threat model published for its confidential computing products, and the company has not assigned a CVE [4]. Intel's stated mitigation path is limited to standard physical-security and supply-chain guidance in the near term, alongside longer-term work evaluating architectural hardening options and a concept the company refers to as Platform Owner Endorsements intended to help verify physical possession of hardware for sensitive workloads [4]. AMD's bulletin similarly categorizes the technique as outside SEV-SNP's published threat model, which assumes the memory bus itself is not directly accessible to the attacker, and the company has not committed to a mitigation or a CVE at this time [5][2]. In CSA's reading, neither company disputes the researchers' technical findings; both are drawing a line around what confidential computing was designed to protect against, physical bus-level tampering versus logical or remote compromise. That distinction matters for how organizations scope their controls, but in CSA's view, it also means organizations that adopted confidential computing specifically to satisfy insider-threat or physical-security requirements should treat this as a gap between the assurance level often implied in vendor and provider marketing and the threat model the hardware actually implements. Notably, the researchers found that NVIDIA's confidential GPU implementations are not vulnerable because their memory sits inside the chip package rather than on an externally accessible bus, and that Intel's now-retired Client SGX was also immune because it used a hardware integrity tree capable of detecting stale data, a protection that TDX's current design does not carry forward [1].

Recommendations

Immediate Actions

Organizations running sensitive workloads on Intel TDX, Intel Scalable SGX, or AMD SEV-SNP confidential VMs, whether self-hosted or through a cloud provider, should inventory which systems rely on these technologies and confirm with their cloud provider or hardware vendor what physical access controls, chain-of-custody procedures, and tamper-evident measures apply to the underlying servers. Security teams should treat DDRop as a reminder to re-examine assumptions embedded in risk assessments, compliance narratives, or vendor contracts that describe confidential computing as fully mitigating insider or physical-access threats, since Intel and AMD have both stated that physical bus-level attacks fall outside their products' published threat models, which undercuts blanket claims that confidential computing fully mitigates insider or physical-access threats [4][5]. Teams operating on-premises or colocated confidential computing infrastructure should review who has physical access to server racks and memory modules, including third-party maintenance technicians, and tighten escort and logging requirements for any hands-on hardware work.

Short-Term Mitigations

No software patch closes the underlying gap, since the vulnerability stems from the absence of cryptographic freshness checking in the memory-encryption designs currently shipping, but several partial mitigations reduce exposure. Where available, Intel's optional cryptographic integrity mode should be enabled, as researchers noted it blocks some of the attack's vectors even though it does not eliminate the underlying weakness [1]. Organizations with the technical capacity to do so should evaluate boot-time and memory-training timing checks that can flag the latency signature an interposer introduces, implement runtime read-back verification for security-critical writes such as attestation state, and restrict or disable optional memory-management features, such as AMD's page-relocation function, where operational requirements permit [1][3]. Cloud customers should ask their providers directly whether these detection and hardening measures have been deployed on the physical hosts running their confidential VM instances, since customers themselves generally cannot verify or remediate host-level physical security.

Strategic Considerations

DDRop is best understood as a hardware-design limitation rather than an implementation bug, and researchers have been explicit that no simple patch exists because current memory-encryption schemes deliberately traded cryptographic freshness for performance and scalability [3]. That tradeoff will likely need to be revisited in future silicon generations, through mechanisms such as hardware freshness counters or integrity trees similar to those used in the retired Client SGX design, and organizations with long procurement horizons should factor this into evaluation criteria for next-generation confidential computing hardware. In the interim, security and compliance teams should recalibrate how much weight confidential computing attestation carries in zero-trust and insider-threat models, treating attestation as strong evidence against remote and logical compromise but not as a complete substitute for physical security controls. Given that DDRop's schematics and firmware are publicly available on GitHub, organizations should expect the barrier to acquiring the parts and design files to be low, even though building a working replica still requires DDR5 hardware engineering expertise; that combination argues for prioritizing physical-access hardening now rather than waiting for a hardware fix that may be one or more product generations away [10].

CSA Resource Alignment

CSA's Healthcare Confidential Computing and the Trusted Execution Environment white paper provides the most directly relevant prior analysis, since it catalogs how healthcare organizations depend on TEEs, hardware security modules, remote attestation, and root-of-trust concepts to protect data in use, and it explicitly discusses memory-exploitation threats and insider-access risk as the category of concern DDRop now demonstrates in concrete form [6]. Readers evaluating confidential computing for regulated data should revisit that paper's guidance on attestation verification and root-of-trust validation in light of DDRop's demonstrated ability to forge TDX attestation reports. The CSA Confidential Computing Working Group Charter, which established the working group's mandate to advance standards and best practices for TEE implementation, hardware roots of trust, and attestation services, frames the organizational context in which CSA is positioned to track and respond to findings like DDRop as they mature toward standardization discussions at venues such as ACM CCS [7]. In CSA's extrapolation, because DDRop concerns the integrity and trustworthiness of infrastructure that underlies both traditional and AI workloads, organizations assessing broader AI system risk should also map physical and infrastructure-layer assurance gaps of this kind against CSA's AI Controls Matrix (AICM) v1.1, which addresses infrastructure and system integrity controls relevant to workloads, including AI workloads, that rely on hardware-based isolation guarantees [8].

References

[1] The Hacker News. "New DDRop Attack Breaks Intel TDX and AMD SEV-SNP Confidential Computing." The Hacker News, September 2026.

[2] SC World. "DDRop attack bypasses Intel and AMD confidential computing defenses." SC World, September 2026.

[3] The Register. "New hardware device can RAM into encrypted memory, expose your data." The Register, September 14, 2026.

[4] Intel Corporation. "Intel Security Announcement: DDRop Physical Memory Interposer Research." Intel Security Center, September 2026.

[5] AMD. "AMD Security Bulletin AMD-SB-3048: Physical Memory Fault Injection Attacks on DDR5." AMD Product Security, September 2026.

[6] Cloud Security Alliance. "Healthcare Confidential Computing and the Trusted Execution Environment." Cloud Security Alliance, 2025.

[7] Cloud Security Alliance. "Confidential Computing Working Group Charter 2024." Cloud Security Alliance, September 2024.

[8] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." Cloud Security Alliance, 2026.

[9] DDRop Research Team. "DDRop Attack Project Page and Research Paper." KU Leuven, ETH Zurich, Durham University, and Google, September 2026.

[10] DDRop Research Team. "ddropattack/ddrop GitHub Repository." GitHub, September 2026.

