

CYBER SECURITY

FRONT /> CODE

SYSTEM SECURITY EVOLUTION & ADVANCED EXPLOIT

November 2025

GPU DRIVERS UNDER FIRE

**NEW YORK PASSES 72-
HOUR CYBERATTACK LAW**

ANALYZING HEISENBUGS

AI SUPPLY CHAIN SECURITY

**MEMORY-SAFE
PROGRAMMING**

**QUANTUM-RESISTANT
CRYPTOGRAPHY**

Cybersecurity 2025

November 2025/Volume 01

1 GPU Drivers Under Fire: Lessons from the Qualcomm and NVIDIA Flaw Disclosures

GPU drivers are back in the spotlight after recent security disclosures from Qualcomm and NVIDIA. Both companies were forced to fix flaws that allowed attackers to escalate privileges, run code, and bypass protections.

These cases highlight how vulnerable GPU software has become, especially when it operates close to the kernel.

Qualcomm's Zero-Day Issues in Adreno Drivers

In June 2025, Qualcomm patched three zero-day flaws in its Adreno GPU drivers. The bugs affected millions of Android devices running Snapdragon chips. Google's Android Threat Analysis Group discovered the flaws and reported them to Qualcomm.

Two of the bugs, CVE-2025-21479 and CVE-2025-2148, scored 8.6 out of 10 on the CVSS scale. These allowed memory corruption through specially crafted GPU commands.

The third bug, CVE-2025-27038, was rated 7.5. It involved a use-after-free condition in the rendering pipeline. All three issues let attackers move from user apps to kernel-level execution.

The flaws were already under attack before the fixes were made public. Google did not disclose who was behind the exploitation, but security analysts believe spyware vendors or state actors were likely involved.

Qualcomm delivered the patches to device manufacturers in May. Users still rely on OEMs to package and release updates, which delays protection.

This is not the first time Adreno drivers have caused problems. Past research from Google's Android Red Team showed that GPU code often lacks permission checks. That gives any app a path to poke at privileged driver functions, making kernel exploits more likely.

NVIDIA's Pile of Vulnerabilities

NVIDIA has published two major security updates this year. In April, it fixed CVE-2025-23244. This bug affected Linux systems and let attackers gain higher privileges from a non-root position. The vulnerability allowed remote code execution, denial of service, and data tampering.

The July update was broader. It fixed seven vulnerabilities across Windows and Linux. These included installer privilege flaws, memory errors, use-after-free issues, and information leaks.

CVE-2025-23276 and CVE-2025-23281 were among the most serious, both scoring above 7.5. Exploiting these flaws could result in local privilege escalation or data corruption.

NVIDIA issued patches for multiple driver branches. These included R535.247.01, R550.163.01, R570.133.07, and R575.51.02. Enterprise users had to apply the right update for their version. NVIDIA did not confirm active exploitation, but users were warned to update immediately.

Why GPU Drivers Keep Causing Trouble

These flaws all share one theme: GPU drivers carry too much power. They run close to the kernel. They manage memory. They accept input from apps without deep validation. They were designed for speed and throughput, not for safety.

As graphics hardware has evolved, the role of the GPU driver has grown. It now supports AI, gaming, encryption, and more. A GPU driver can expose a wide surface to attackers while operating with near-kernel privileges.

On Android, Google has flagged GPU drivers as a risk. Apps can access them freely, even when they contain bugs. An untrusted app can send malformed GPU commands and exploit memory handling issues to gain control. That puts billions of devices at risk.

Academic research has shown that GPUs can leak memory between users. One study found that uninitialized register access could let users spy on others' workloads. The problem extends beyond phones—it affects desktops, servers, and cloud environments.

Slow Patches and Poor Visibility

One big problem is the time it takes to deliver fixes. Qualcomm made the patches available in May, but most users did not get them until June or later. Android users depend on phone manufacturers to push updates. Many older phones may never get a fix.

NVIDIA controls its own update process, but still requires users to install the correct driver manually or through a managed system. Many users skip updates if performance is stable. That leaves them exposed.

Another issue is visibility. Most users don't know what a GPU driver patch really fixes. Manufacturers rarely explain that a graphics update closes a privilege escalation path. Unless a vulnerability makes the news, the threat may go unnoticed.

What Needs to Change

GPU drivers must be audited more thoroughly. Vendors should apply modern fuzzing and static analysis techniques to driver code. They should isolate GPU execution contexts and enforce strict boundaries around memory operations. Security researchers should keep probing GPU interfaces. Many of the flaws this year were found by independent teams, not by the vendors themselves.

Device makers must commit to faster patch delivery. If an OEM cannot patch GPU drivers quickly, it should not ship devices with exposed components. Google's push for modular driver updates through Project Treble and Android's Mainline modules is one possible fix, but adoption is still limited.

Users need to take updates seriously. On Windows and Linux, apply new NVIDIA drivers as soon as they're released. On Android, check for updates often and avoid devices that don't receive security patches.

To Conclude

GPU drivers have grown from performance tools into complex systems that run at the heart of computing. The recent flaws in Qualcomm and NVIDIA drivers show how dangerous poor GPU security can be. These aren't abstract bugs. They've been exploited in real attacks.

Vendors must treat GPU drivers like any other kernel component. That means full testing, rapid patches, and clear communication. It also means stripping out legacy assumptions that graphics code is low-risk. Security starts at the driver level. These disclosures make that clear.

2 New York Passes 72-Hour Cyberattack Reporting Law

Local governments in New York must now report cyberattacks within 72 hours. If they pay ransom, they must report that within 24 hours.

A detailed explanation must follow within 30 days. The new law also requires public employees to complete annual cybersecurity training.

Full Backing from Lawmakers

Senate Bill S.7672A/A.6769A passed both chambers without any opposition. The law builds on Governor Kathy Hochul's earlier cybersecurity plan. In 2022, she appointed the state's first cyber officer. Since then, the administration has pushed to tighten rules for public systems.

Clear Deadlines for Reporting

Under the new rule, any county, city, or agency hit by a cyberattack must alert the New York State Division of Homeland Security and Emergency Services within three days.

If the agency pays ransom, it must report that within one day. It then has 30 days to explain the decision. That report must include how much was paid and whether legal review was done.

Goal Is Faster Response and Less Damage

State officials say these deadlines help contain threats. Early reporting allows the state to track attacks, warn others, and send technical help if needed.

The law also closes a gap. Financial companies already follow similar rules under New York's Department of Financial Services. That regulation has required 72-hour incident reporting since 2017. New York's new law for public agencies brings the same urgency to schools, towns, and counties.

Governor and Cyber Officials Weigh In

Governor Hochul said the law protects basic services across the state. She called cybersecurity a core duty of public agencies.

Colin Ahern, New York's chief cyber officer, said attackers now target small towns and under-resourced systems. He said the new rule brings structure to a problem that has long been patchy across local governments.

Web Portal to Report Attacks

To help with compliance, the state launched a web portal. Local agencies can file reports and request help through a single platform. This lets state teams track incidents and coordinate faster.

Jackie Bray, head of Homeland Security and Emergency Services, said the system is built for speed. She said delays cost time and increase damage.

Public Training Now Required

The law also introduces yearly cybersecurity training for all public employees. The state will set minimum security rules that every public agency must follow.

The goal is to make basic protection routine. Local governments must now update their internal policies and schedule training sessions.

Ransom Reporting in Practice

If a city discovers a ransomware attack on August 10, it must report the breach by August 13. If it pays a ransom on August 11, it must file that report by August 12. By September 10, the city must submit a full written explanation of its decision.

These deadlines are firm. Reports must go through the state's new portal.

Some Pushback, but Support Holds

Some smaller towns have raised concerns. They say the timelines may be hard to meet with limited staff. State officials say the portal and training will help. They also say centralized reports will show broader patterns, which helps prevent future attacks.

Assembly member Billy Jones, one of the sponsors, said the law gives local leaders a rulebook. He said many towns didn't know what to do after being hit. Now, they do.

Law Is Now Active

The law took effect immediately. All public agencies in the state must follow it. Those that fail to meet the reporting deadlines may face penalties.

A First for Municipal Systems

New York is among the first states to require such reporting from local governments. Most state rules have focused on private companies and critical infrastructure. This law brings schools, town halls, and public agencies into the same system.

The focus is clear: act fast, report early, limit damage. Cyber threats grow every day. This law sets a simple rule, no more silence after an attack

3 Analyzing Heisenbugs with Hardware-Assisted Kernel Debugging

Heisenbugs are hard to detect bugs that can be omitted when observed in optimized and concurrent systems commonly used nowadays. Kernels are particularly difficult to debug with these bugs because the most traditional debugging approaches cannot be used due to timing, concurrency, and compiler optimizations.

Often surfacing only at scale in production, Heisenbugs require more than standard debugging. This article explores how combining structured techniques with hardware-assisted tools can help uncover and resolve these tricky system-level issues.

Understanding Heisenbugs


The heisenbug is usually a result of such complications as uninitialized memory, compiler optimization reordering, lack of determinism of memory layout, or race conditions.

Bugs can sometimes vanish during debug mode but remain in release, since debug builds, unlike release, zero out memory or serialize timing, as a means of saving programmer time to debug the program.

With kernel-mode, there is more to lose: tiny disturbances like enabling JTAG or tracepoints can change scheduling, serializers, or register spills, and jeopardize the bug without anyone realizing it.

Traditional Reproduction Techniques: Systematic Scheduling

Before hardware tools became mainstream, researchers used software-based approaches like CHES (by Microsoft Research) to systematically explore thread interleavings. CHES wraps .NET or Win32 programs, monitors their execution, and records the "happens-before" graph of thread interactions. This graph is then used to deterministically replay thread schedules and expose concurrency issues.



SOURCE

The architecture of CHES involves program instrumentation and a scheduler that records and replays thread behaviors. As shown, it integrates with the runtime through wrappers and feeds execution traces to a search module that systematically explores all possible thread schedules.

The CHES model enables reliable reproduction of hard-to-catch race conditions, even in programs with billions of potential interleavings.

Scaling Diversified Execution with Hardware Tracing

Beanstalk is a framework that improves Heisenbug detection by running the same kernel code across different configurations, varying instrumentation levels, and hardware setups.

This technique, called diversified execution, increases the chances of revealing concurrency-related bugs. Another prominent aspect of Beanstalk is a metric called the heisen factor, which measures the difficulty of reproducing a given bug to determine which bugs to fix first.

3 Analyzing Heisenbugs with Hardware-Assisted Kernel Debugging

To assist it, hardware-aided debugger tools such as WinDbg or HyperDbg are employed in a host-guest pairing, where the debugger is on the host, and accesses a virtualized kernel through named pipes or a virtual serial port. This method does not disturb kernel timing or behavior, making the trace capture accurate.

Integrated with hardware-level tracing, diversified execution offers a highly scalable way to detect and follow up on the hard-to-detect kernel bugs missed by other debugging tools.

Case Study and Empirical Data

A tool called H3, introduced at the USENIX ATC conference, shows how we can reproduce heisenbugs that happen in real-world kernel environments. It works by using special CPU features to capture and replay the exact sequence of events that caused the bug.

This lets developers recreate and study bugs that were previously impossible to reproduce, all on normal hardware. Even though this method is a bit older, it showed that it's possible to debug tricky kernel bugs without changing the system in ways that might hide them.

In another example, researchers studied silent data corruptions, errors caused by faulty hardware that behave like heisenbugs. Over 18 months, they looked at data from hundreds of thousands of computers and found hidden issues in CPUs.

They applied hardware-based debugging to trace backward the specific instructions that failed, and what caused the troubles, of all the components of the system, including hardware and up to the software.

Such examples point out that hardware-level debugging does not merely represent an option to detect software flaws, but it also has the potential to uncover hidden hardware flaws that may be circumvented by other conventional debugging tools.

Best Practices and Practical Tips

- Don't fill up the system by logging or debugging code: There is a fine line between sufficient debugging information and too much. Maintain lightweight tools for tracking.
- Rely on hardware features for clean tracking: Use hardware tracing and virtualization-based tools to see exactly what's happening inside the kernel, without changing how it behaves.
- Quantify how difficult the bug is: With a scale such as the heisen factor on the Beanstalk platform, you can determine how difficult the bug is to reproduce. This assists in prioritizing bug fixes.
- Automate the reproduction of the bug: Do repeatable tests that make the bug emerge, so that you can see whether the bug that existed before was actually fixed when you get the same bug conditions.

Conclusion

Heisenbugs are among the toughest kernel issues, often hiding in timing, optimization, and concurrency, and disappearing under normal debugging. But with hardware-assisted tracing, tools like HyperDbg and Beanstalk, and smarter debugging methods, these bugs can now be reliably found and fixed.

Kernel debugging is shifting from guesswork to a precise, repeatable process. Developers should explore tools like HyperDbg, CHESS, Snowboard, and Beanstalk to better understand and resolve these elusive system-level bugs.

4 AI Supply Chain Security — The Hidden Risk Behind Machine Learning Models

Artificial Intelligence (AI) has become central to digital transformation, but as reliance on pre-trained models and open datasets grows, so do the security risks hiding inside the AI supply chain. Recent research and real-world incidents show that attackers are exploiting these weak links — tampering with data, frameworks, or models to compromise AI integrity and trustworthiness.

A Growing Problem Hidden in Plain Sight

Modern AI systems are rarely built from scratch. Developers typically use pre-trained models from open repositories such as Hugging Face or TensorFlow Hub. These models are then fine-tuned using public or proprietary datasets before being deployed into production systems.

This interdependence of datasets, libraries, and cloud services has created an AI supply chain, one that mirrors — and multiplies — the software supply chain problems already seen with attacks like SolarWinds or Log4j.

In early 2025, multiple cybersecurity labs reported that malicious models disguised as open-source AI tools were circulating online. Once integrated, these models quietly transmitted data to remote servers. The infections were discovered only after months of undetected operation in enterprise test environments.

Real-World Attacks on AI Pipelines

In one widely discussed case, researchers found that a machine learning library in a public repository had been replaced with a poisoned version. The malicious file shared the same name and version number but contained hidden code to capture environment variables, including API tokens.

Another case involved data poisoning in facial recognition training sets. Attackers subtly altered a small percentage of labeled images, leading the resulting model to misidentify individuals of certain demographics. Because the manipulated samples were statistically insignificant, traditional validation didn't detect the corruption.

A 2024 report from MIT's AI Policy Lab warned that model supply chain risks could lead to "invisible influence" attacks, where poisoned models are used to skew financial, political, or hiring algorithms without immediate detection.

Why the AI Supply Chain Is Vulnerable

AI's architecture inherently encourages openness — open datasets, shared models, collaborative code. While this accelerates innovation, it reduces visibility into provenance and trust.

Unlike compiled software, AI models are binary blobs that embed millions or billions of parameters. There's no straightforward way to verify their internal integrity.

Moreover, existing security tools focus on application-level vulnerabilities, not on anomalies in model weights or dataset consistency. This makes the AI stack fertile ground for sophisticated attackers who prefer stealth over brute force.

Efforts Toward a Secure AI Pipeline

Security researchers and standard bodies are now pushing for AI provenance frameworks — systems that track every stage of model creation, from data ingestion to deployment.

NIST and ISO have begun exploring metadata standards for "Model Bills of Materials" (MBOMs), similar to Software Bills of Materials (SBOMs). Such documents would list model sources, training datasets, contributors, and dependencies.

Organizations like OpenAI and Google DeepMind have already introduced cryptographic model signing to verify authenticity. Some cloud platforms also offer "trusted training environments," which isolate model-building processes from external interference.

4 AI Supply Chain Security — The Hidden Risk Behind Machine Learning Models

Slow Adoption and the Road Ahead

Despite technical progress, adoption remains slow. Many smaller companies and researchers still rely on unverified public models to save time and resources. Awareness of model-level threats is limited compared to traditional software vulnerabilities.

As AI continues to integrate into critical infrastructure — from healthcare diagnostics to autonomous vehicles — the consequences of supply chain breaches could be catastrophic. Experts argue that without regulation and transparency, the next major cyberattack might not target software — but the data and models that shape it.

To Conclude

AI's promise depends on trust, and that trust begins at the source. The AI supply chain — once seen as a collaboration hub — has become a potential attack surface.

Protecting it will require not just new tools but a cultural shift in how AI systems are built, verified, and shared.

As the boundaries between software and intelligence blur, securing the supply chain means securing the future of AI itself.

5 Memory-Safe Programming — The Future of Secure Systems

Software vulnerabilities have long been the Achilles' heel of digital infrastructure. Despite decades of innovation, memory safety issues still account for the majority of high-severity security flaws. Now, a growing shift toward memory-safe programming languages like Rust and Go promises to reshape how we build secure systems — from the operating system kernel to web browsers.

The Persistent Memory Problem

For nearly half a century, software engineers have relied on C and C++ for system-level programming. These languages give developers powerful control over memory — but also make it easy to misuse it.

Bugs such as buffer overflows, dangling pointers, and use-after-free errors remain the root cause of critical exploits. Attackers can use these flaws to gain elevated privileges, corrupt data, or execute arbitrary code.

A 2024 Microsoft study revealed that over 70% of vulnerabilities in Windows originate from memory safety issues. Google echoed similar findings for Chrome and Android. Despite countless security patches, the pattern repeats because the underlying cause — unsafe memory access — persists in legacy codebases.

The Rise of Memory-Safe Languages

Memory-safe languages automatically prevent these dangerous operations. Rust, in particular, has gained traction for its strict compile-time guarantees that eliminate whole classes of memory errors.

By enforcing ownership and borrowing rules, Rust ensures that only one reference can modify a variable at a time. This prevents data races, double frees, and null pointer dereferences before the program even runs.

Tech giants are taking notice. Google is rewriting Android's media stack in Rust, leading to a significant drop in high-severity bugs. The Linux kernel added official Rust support in version 6.1, marking the first major shift in kernel development in decades.

Industry Adoption Accelerates

Microsoft has launched internal projects to rewrite Windows components in Rust, citing early reductions in crash telemetry. Amazon Web Services' lightweight virtualization engine, Firecracker, is entirely written in Rust — balancing performance with safety.

Mozilla's Servo engine pioneered safe concurrency, while Cloudflare's networking stack now leverages Rust to prevent memory leaks and overflows.

Even where full migration isn't possible, hybrid models are emerging. Rust modules are being embedded within C/C++ codebases using interface bindings. This allows gradual adoption without a complete rewrite, giving developers the benefits of safety with minimal disruption.

Why It Matters for Cybersecurity

The shift to memory-safe languages represents more than a trend — it's a fundamental change in software engineering philosophy. Instead of patching vulnerabilities after they're exploited, developers can now eliminate them at the source.

Modern compilers and tools, such as LLVM sanitizers and ARM's Memory Tagging Extension (MTE), complement this shift by detecting unsafe memory operations during testing.

However, these are stopgaps. True resilience comes from designing software where memory corruption is impossible by default.

Challenges and Transition Costs

The biggest barrier is legacy code. Rewriting decades of production systems is costly and time-consuming. Moreover, not all developers are trained in new languages like Rust. Yet the growing number of critical exploits has made modernization unavoidable.

Security experts argue that government-funded infrastructure and safety-critical industries — such as healthcare, defense, and aviation — should mandate memory-safe development for future software.

To Conclude

The move toward memory-safe programming is not just a modernization effort — it's a turning point for cybersecurity.

In the next decade, the phrase "written in Rust" could become a new trust seal for software, signaling safety by design.

After years of reactive patching, memory-safe programming offers something rare in security: prevention that truly lasts.

6 Quantum-Resistant Cryptography — Preparing for the Post-Quantum Era

The world's encryption systems, from online banking to national defense, rely on mathematical problems that classical computers can't solve efficiently. But quantum computing threatens to change that. Once considered theoretical, quantum machines are advancing fast enough to endanger today's cryptographic foundations — leading to what experts call the post-quantum security challenge.

Quantum Power Meets Classical Vulnerability

Traditional encryption — including RSA, Diffie-Hellman, and Elliptic Curve Cryptography (ECC) — depends on the difficulty of factoring large numbers or solving discrete logarithms. A sufficiently powerful quantum computer, using Shor's algorithm, could solve these in hours.

While such systems don't yet exist at full scale, intelligence agencies and corporations are already preparing. The fear is that attackers could harvest encrypted data today and decrypt it later once quantum computers mature.

The Global Push for Post-Quantum Cryptography

In response, the U.S. National Institute of Standards and Technology (NIST) has spent years evaluating new algorithms designed to resist quantum attacks. In 2024, it finalized the first generation of post-quantum cryptography (PQC) standards.

The selected algorithms include:

- CRYSTALS-Kyber for key encapsulation and exchange
- CRYSTALS-Dilithium for digital signatures
- Falcon for compact, high-performance signatures
- SPHINCS+, a hash-based backup algorithm emphasizing long-term resilience

These new cryptosystems rely on lattice-based and hash-based problems that are computationally infeasible even for quantum machines.

Industry Transition Begins

Major technology companies have already started adopting PQC in hybrid configurations. Google has implemented hybrid TLS in Chrome, combining classical and PQC key exchanges. Cloudflare and AWS are testing similar deployments.

Meanwhile, governments worldwide — including the U.S., Japan, and members of the EU — are mandating migration plans for all critical systems.

However, the transition is not straightforward. Post-quantum algorithms often have larger key sizes, require more processing power, and must be integrated without breaking legacy compatibility. Many IoT and embedded devices lack the resources to handle PQC's computational demands.

Migration Challenges

Organizations must first inventory their cryptographic assets — identifying what encryption they use and where. Next, they need crypto-agile systems capable of swapping algorithms without full redesigns.

Hardware vendors are also exploring new chips optimized for PQC, aiming to make future encryption faster and more energy-efficient.

One overlooked issue is data that's already been encrypted with classical methods. Even if PQC is adopted tomorrow, any sensitive data stored using traditional keys may remain vulnerable to "decrypt later" attacks in the future.

6 Quantum-Resistant Cryptography — Preparing for the Post-Quantum Era

Preparing for a Quantum Future

Experts advise organizations to begin transition planning immediately. The National Security Memorandum 10 (NSM-10) issued by the U.S. government requires federal agencies to begin PQC migration by 2027. Industry leaders suggest private enterprises follow suit to maintain interoperability and trust.

Cybersecurity researchers believe that the “quantum threat timeline” — once thought to be 30 years away — could shrink dramatically as companies like IBM, Google, and D-Wave achieve exponential growth in qubit performance.

To Conclude

Quantum computing represents both opportunity and existential risk. If the cryptographic world waits until quantum computers are ready, it will be too late to secure what's already encrypted.

The adoption of quantum-resistant algorithms must begin now, ensuring that today's secrets remain secure tomorrow. As one security researcher aptly put it: “Post-quantum cryptography isn't about the future of encryption — it's about preserving the past.”