THINKST SCAPES

Q4 • 2025

https://thinkst.com/ts

Contents

Cover photo: Dali Monastery in Darjeeling, India. *Image by Dev Dua (Thinkst).*

Plum blossoms in Darjeeling, India. Image by Dev Dua (Thinkst)

Introduction

Welcome to this edition of ThinkstScapes for Quarter 4, 2025! This issue focuses on content released, published or presented from the first of October through to the end of 2025.

As expected with the holidays, there was a slight decrease in venues and total volume of published work. But what the quarter lacked in volume it made up for in quality, with 19 selected works! Between long-time stalwart events (Black Hat Europe, Hack.lu, and CCC), and newer conferences hitting the ground running (Hexacon and the Offensive AI Conference), there was lots of research to showcase.

As a reminder: if you are aware of work we've missed, a blog post we should have seen or a conference we should have covered, we'd love to hear about it. Please send them to **ts@thinkst.com**!

In addition to more than 1,200 blog posts, this quarter's content was drawn from talks and papers presented at the following conferences:

Venue name	Number of talks/papers
0xcon	10
39C3	159
ACSAC 2025	83
ArcticCon 2025	10
Asiacrypt 2025	150
Black Alps	22
Black Hat EU	52
Bluehat Asia 2025	14
BSides Atlanta	35
BSides Augusta	25
BSides Cymru	9
BSides Cape Town	25
BSides Perth	16
BSides Toronto	15
BSides London	50
BSides Munich	33
CODE BLUE 2025	57
Chcon	29
Cybercon 2025	27
Cyberwarcon	22
DeepSec	48
GrrCON Cyber Security Summit and Hacker Conference	66
Hack.lu	50
Hackfest	68
Hacktivity 2025	23
Hexacon 2025	15
HushCon	13
JawnCon OX2	23
KawaiiCon	27
No Hat 2026	14
Offensive AI Con	17
Objective by the Sea version 8.0	31
PasswordsCon	11
Power of Community	16
Show Me Con	66
Triangle InfoSeCon	10
grehack 2025	11
hack Sydney	18
Total	**1,370**

Jacarandas in Bloom, Pretoria. *Image by Bradley Jayanath (Thinkst).*

Themes covered in this issue

NETWORKING BEYOND PLUG-AND-PLAY

Much of our modern exposure to networking is the consumer UX of plug-and-play Ethernet or WPA-PSK. This theme reminds us that there's much more lurking under that calm surface of simplicity. Look for volumetric reflection DDoS attacks using TCP, parser differentials in WAFs, high-precision timing synchronisation, and a DoS against QUIC that renders it anything but.

HIGH-IMPACT SECURITY AT THE FOUNDATIONS

This quarter saw a plethora of kernel and embedded security work. Finding flaws in these foundational layers weaken everything higher in the stack. Look for: empirically measuring the security benefits of hardware-based capability security, hacking LLMs with GPU driver exploits, a trivial KASLR bypass, finding bugs by emulating a fake phone, and how Rust is going on Android. We wrap up with the exploitation of a humanoid robot – demonstrated by autonomously punching a mannequin.

WINS AND LOSSES WITH LLMS AND SECURITY

LLMs are leaving their mark on security, and not always in a good way. This theme covers how to best design autonomous agents to handle uncertainty, hijacking Google's Antigravity IDE, improving CodeQL results with an LLM, improving RAG robustness, and timing side-channels against LLM chatbots.

NIFTY SUNDRIES

As always, we also find some great work that stands alone. This quarter, look for: image encryption that can survive lossy compression, a pragmatic look at the need for post-quantum cryptography, using MFA as a phishing sensor, and the benefits of taking another look at old CVEs.

A year in review

We annually reflect on the signals we've extracted from our enormous selection of talks, blogs and papers (available on Citation). With this issue wrapping up 2025, we wanted to highlight some of the research community's outstanding work. In the last 12 months, we extracted the following themes:

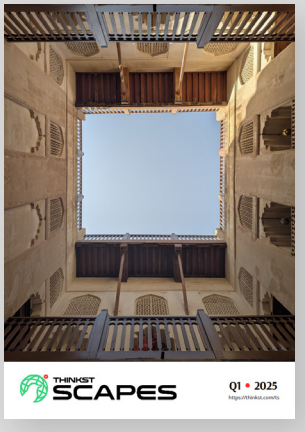

Q1:

- Putting it into practice
- Understanding things all the way down
- Scaling software (in) security

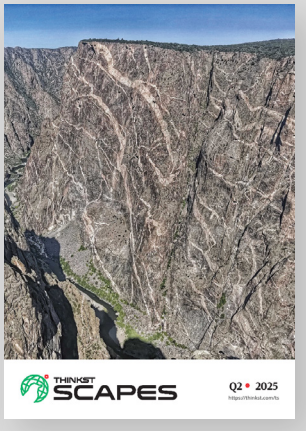

Q2:

- Networking is always tricky
- Language models large and small
- When parsing goes right, and when it goes wrong

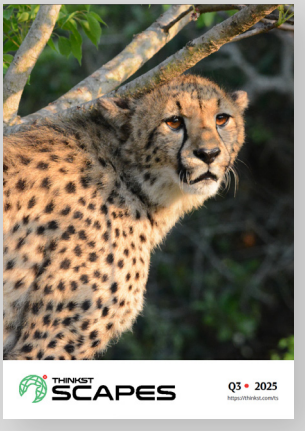

Q3:

- Microsoft-induced security woes
- Logs are not always as they appear
- Autobots roll out!
- Good vibrations

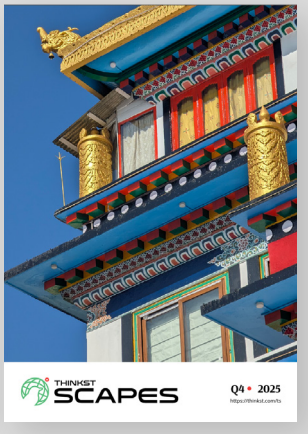

Q4:

- Networking beyond plug-and-play
- High-impact security at the foundations
- Wins and losses with LLMs and security

Despite the tumultuous year in many other respects, this year was pretty consistent in the focus areas across all four quarters. LLM research was starting to solidify beyond just a hot topic at the end of 2024, and now consistently provides interesting content in terms of bug finding, scaling application security, and interfacing with users (or victims). Networking oddities and parser bugs continued to pop up, even in decades-old protocols and time-worn production systems. Lastly, there was a continued focus into the system components that underpin our connected world where an exploit chain jumped from web to kernel to GPU, etc., and back again.

Stewart Island/Rakiura. *Image by Jacob Torrey (Thinkst).*

Networking beyond plug-and-play

- *GET /large file HTTP/1.1: Connection-Based TCP Amplification Attacks*
- *WAFFLED: Exploiting Parsing Discrepancies to Bypass Web Application Firewalls*
- *Excuse me, what precise time is it?*
- *Cut To The QUIC: Slashing QUIC's Performance With A Hash DoS*

GET /large file HTTP/1.1: Connection-Based TCP Amplification Attacks

***Authors**: Yepeng Pan, Lars Richter, and Christian Rossow*

Paper

Code

This research explored the feasibility of performing a volumetric amplification attack using TCP instead of UDP. UDP is commonly used due to the connectionless nature of the protocol, so attackers can spoof their source IP to that of the victim, and make a request to a server that replies with a response an order of magnitude (or more) larger. By repeating this process, the attacker only needs to send a few bytes to the reflector, and have significantly more traffic sent to the victim – usually knocking them offline.

Since TCP requires a three-way handshake, a spoofed request will only send the later parts of the handshake to the victim before the session is dropped. By scanning for internet-connected devices with predictable sequence number generators, the researchers could spoof both the initial connection packet and the ACK to make an application-layer request. Using HTTP in the study, the researchers could have the contents of any large file hosted by the intermediate reflector sent to the victim – amplifying their attack by as much as 50x.

This attack requires a few weaknesses to be present, but present anywhere on the internet: an ISP that allows for spoofed source addresses (~28% of IPv4 ASes), and a reachable server with a weak TCP sequence generation scheme. By initially scanning IPv4 space for such servers, the researchers could then probe them for the sequence number and then immediately send spoofed packets to complete the handshake and make a number of HTTP requests. The amplification factor was worse if the victim network dropped the unexpected traffic as opposed to sending a RST or ICMP unreachable – only ~12% of surveyed networks sent packets back to the reflector.

TAKEAWAYS:

- **Amplification attacks** that rely on source address spoofing highlight shared externalities in the internet.
- **Usually, the settings of an off-path ISP's router** are irrelevant to most network hosts, but allowing spoofing opens the door for this class of attack.
- **While volumetric attacks** are not very thrilling from a technical standpoint, they have humbled even some of the largest and most well-connected entities online.
- **Pushing for clamping down** on spoofing will make everyone safer.

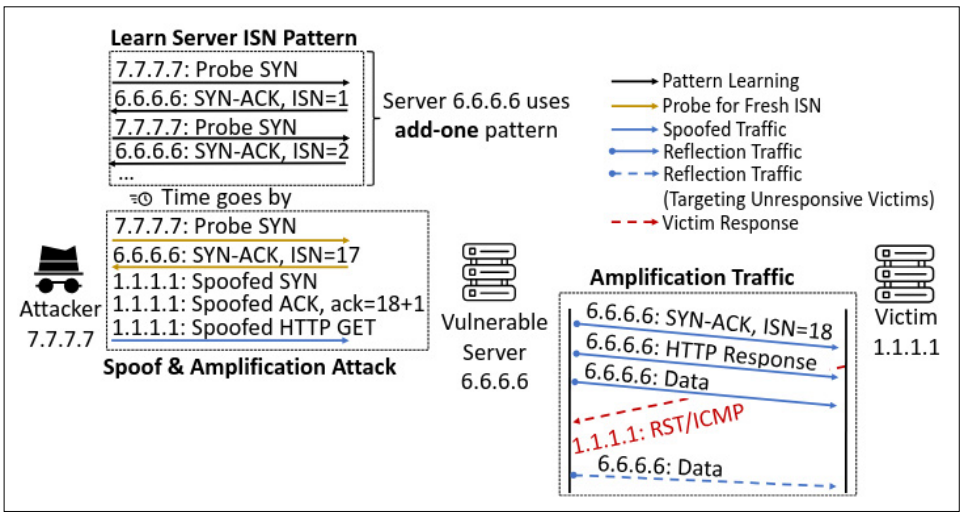

Figure 1. *A high-level figure of the attack steps from learning the reflector sequence numbering scheme to spoofing the victim and sending amplified traffic to the victim host or network.*

WAFFLED: Exploiting Parsing Discrepancies to Bypass Web Application Firewalls

***Authors**: Seyed Ali Akhavani, Bahruz Jabiyev, Ben Kallus, Cem Topcuoglu, Sergey Bratus, and Engin Kirda*

Paper

Code

This research explored bypasses to a number of popular web application firewalls (WAFs) that rely only on manipulating the headers and request structure. While there have been numerous WAF bypasses that changed the payload itself, most have been addressed. The researchers built a fuzzer that looked for discrepancies between the WAF's understanding of the request and the target server. For example, most WAFs would stop parsing requests after the first boundary delimiter in a `multipart/form-data` request, whereas modern web servers supported multiple boundaries which are concatenated together. By fuzzing the WAFs, the team was able to find over 1,200 unique bypass primitives.

Other than AWS's WAF, all the tested common WAF services (and ModSecurity for NGINX) had header and delimiter configurations that would cause the WAF to pass a malicious request through. On the server side, the researchers then had to test if the modifications would still be parsed fully by the server's implementation – resulting in a matrix of bypass combinations. Finally, a prototype protocol normaliser was implemented that would bring a request into adherence with a strict interpretation of web standards. This normalisation layer was able to reject or normalise all the classes of bypass – showing that defending against this type of attack is possible.

TAKEAWAYS:

- **This work highlights that changing specifications** (over time) can induce differentials along with implementation-specific differences.
- **Killing this bug class long-term** will be nearly impossible unless there's full control over each component in the system and the grammar allowed. That certainly is not the case with modern websites.
- **Beyond the immediate implications** of WAF weaknesses stemming from parser differentials, this work shows how fuzzing can very quickly find many non-memory-corruption vulnerabilities. Combining a fuzzer that generates HTTP requests sent through various WAFs to common backends quickly enumerated over 1,200 unique bypass primitives.

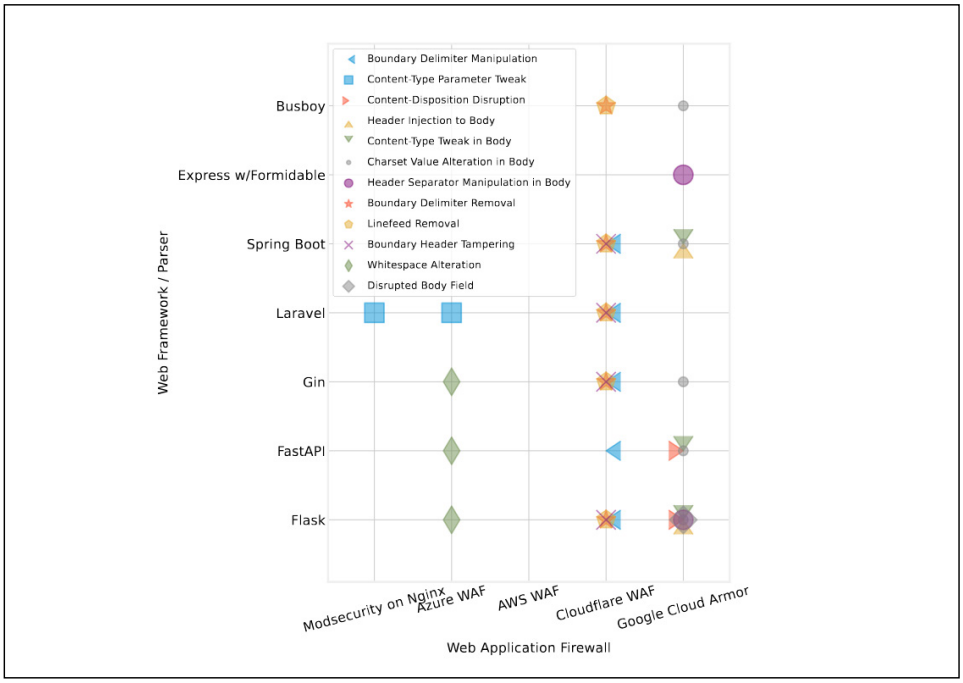

Figure 2. *A chart showing the types of manipulation that were accepted by a target web parser and bypassed a WAF.*

Excuse me, what precise time is it?

***Author**: Oliver Ettlin*

Video

This talk explored the protocols and vendor equipment to support high-precision timing synchronisation across a network. While the Network Time Protocol is sufficient for the majority of applications, higher precision (within 500 ns) is needed in certain industries, such as power grids and, the focus for this talk, broadcast media. The Precision Time Protocol (PTP) is used to transmit clock from a reference (commonly GNSS – i.e., GPS) to receiving network devices.

The talk included a number of live demonstrations to show how various network configurations and topologies can skew time significantly. High-end switches can support a time-keeping delta that is added to the timing packets to keep track of how much time lapsed with the frame in a switching queue, and the receiver can request a delay measurement from the transmitter. However, the speaker noted that there are often purposeful length irregularities between the transmit and receive fibers, which can skew receiver clocks. Finally, the speaker notes there are no security mechanisms built into the PTP protocol, so any device can broadcast a timestamp. Only network switch ACLs can prevent time broadcasts from unexpected ports.

TAKEAWAYS:

- **Accurate timing is rarely needed,** but when it is, it's usually for a vital service. The extreme precision needed also limits the ability to perform security checks.
- **It's worth noting** that a single misconfigured switch can disrupt the network's sense of time, which can cause serious downstream impacts.
- **Ongoing dependence on GNSS-based time sources** will continue to make downstream applications that rely on high-precision time vulnerable to nation-state jamming or manipulation.

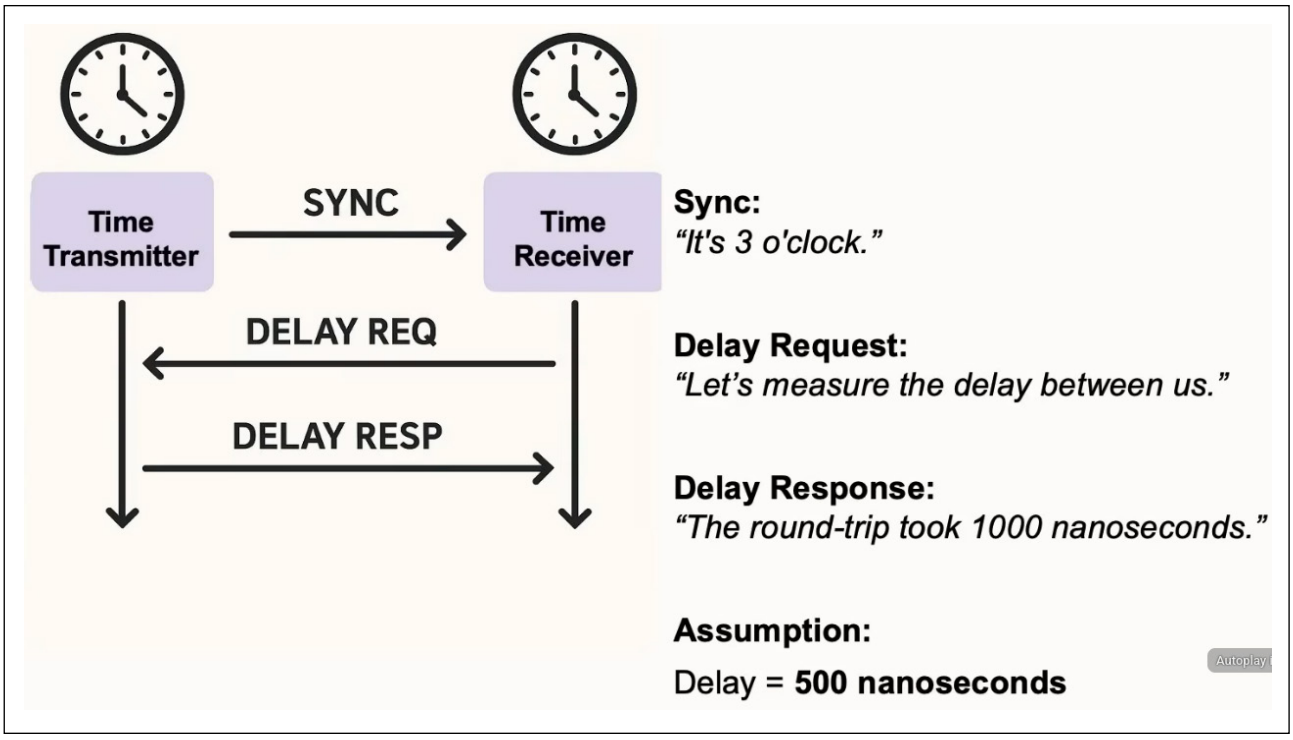

Figure 3.
A diagram showing how each receiver can request a delay measurement to determine the latency between the transmitter and itself on the network.

Cut To The QUIC: Slashing QUIC's Performance With A Hash DoS

Author: *Paul Bottinelli*

Slides

Code

This research explores how connection tracking is implemented in the QUIC protocol, along with a way to drastically reduce its performance. QUIC is a new protocol (already carrying over a third of all web traffic) that aims to combine the TCP and TLS handshake into a single operation, reducing the time needed to send the first data byte. Another feature of QUIC is the ability to rapidly recover from changes of IP addresses, such as when a device switches from Wi-Fi to mobile internet. The server stores a pair of connection IDs that are used to restore a dropped connection upon IP change.

In the QUIC protocol, the client's connection ID, which is used as a lookup value, is client-provided and thus attacker-controlled. This research explored how finding collisions in the connection IDs could create a DoS for the server if the hashing function was not collision-resistant. Four implementations were found to have weaknesses, xquic being one of the worst performers, eventually reaching ~7 minutes to look up a connection ID.

TAKEAWAY:

Nothing about the underlying theory of hash collisions is new, so it's worth noting that this researcher discovered four CVEs of the same class in 2025. Looking for older bugs that have been given a fresh lease on life in new software is bound to continue to pay off.

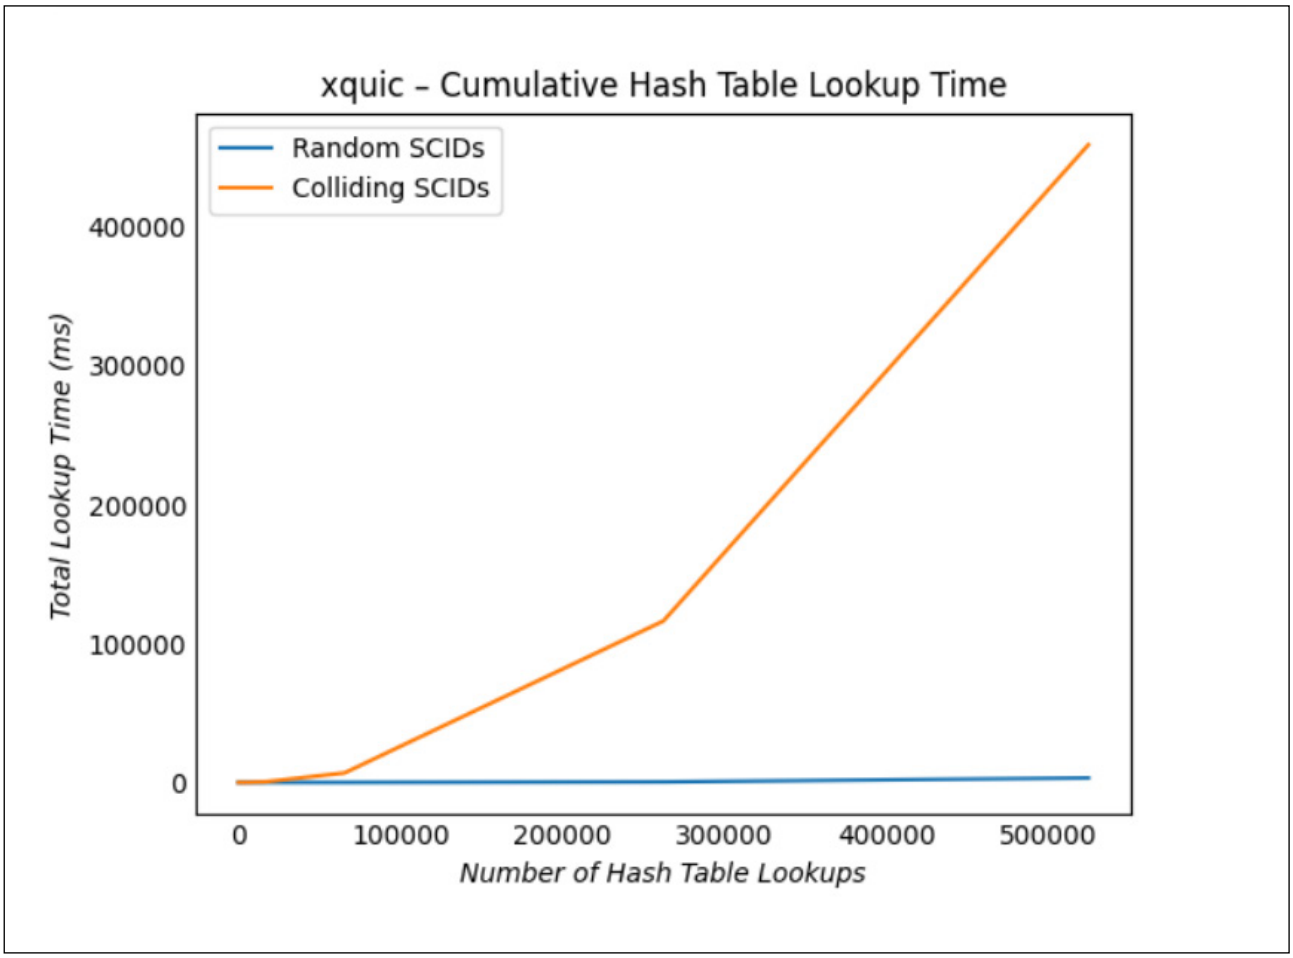

Figure 4.
A chart showing the lookup time for a connection ID pair when there are collisions selected by the client versus randomly generated.

Hong Kong. Image by Gareth Wood (Thinkst).

High-impact security at the foundations

- ***Understanding the Security Impact of CHERI on the Operating System Kernel***
- ***CUDA de Grâce: Owning AI Cloud Infrastructure with GPU Exploits***
- ***Defeating KASLR by Doing Nothing at All***
- ***Build a Fake Phone, Find Real Bugs: Qualcomm GPU Emulation and Fuzzing with LibAFL QEMU***
- ***Rust in Android: move fast and fix things***
- ***Skynet Starter Kit: From Embodied AI Jailbreak to Remote Takeover of Humanoid Robots***

Understanding the Security Impact of CHERI on the Operating System Kernel

***Authors**: Zhaofeng Li, Jerry Zhang, Joshua Tlatelpa-Agustin, Xiangdong Chen, and Anton Burtsev*

Code

Paper

This research explored how to address OS kernel vulnerabilities with hardware or programming language advancements. The Capability Hardware Enhanced RISC Instructions (CHERI) model adds hardware "capabilities" to prevent out-of-bounds memory accesses and to enforce temporal safety. These instructions are slowly starting to appear in silicon, and FreeBSD was recently patched to support CHERI capabilities natively in the kernel.

The research started with a survey of kernel vulnerabilities in both the FreeBSD and Linux kernels, and a decomposition of their root causes. Each root cause was analysed to determine if it would have been blocked by CHERI hardware (assuming the software was updated), and if the vulnerability would exist if ported to Rust. Linux and Windows kernels are actively exploring using Rust in the kernel for safety, but there is a significant body of legacy code in unsafe C/C++. Updating a kernel to support CHERI is a smaller lift than rewriting it in another language – FreeBSD took approximately seven months of effort to add CHERI support.

While Rust would prevent more of the surveyed vulnerabilities, the effort is considerably higher. Kernel modules, which comprise the majority of the vulnerability space and code, are easier to adapt to CHERI than other kernel functions. CHERI shouldn't be seen as the solution that prevents innovation and advancement from other domains, but it offers an affordable model to recompile for improved security.

TAKEAWAYS:

- **Careful hardware and software co-design** can offer security benefits without merely "bolting-on" mitigations. CHERI in particular is well-poised to have an impact due to the lightweight changes needed and the resultant security improvements – if the hardware can get into the hands of users (still years away).
- **Unfortunately, there will be a long period** of exposure to memory corruption vulnerabilities (especially in kernels), regardless of approach.
- **Hardware features take many years to reach ubiquity,** software language replacement is a slow process, and there will always be a lag in deploying more secure options.

Figure 5. *A table of the surveyed vulnerabilities causes, and whether CHERI or Rust would have blocked the vulnerability.*

Cause	CHERI		Rust		Vulns. blocked(%)	
	Yes	No	Yes	No	Cheri	Rust
Language	43 | 13	30 | 60	69	4	58% | 18%	94%
Integer overflow	3	1	4	0	75%	100%
Polymorphism	8	0	8	0	100%	100%
Container of	0	1	1	0	0%	100%
Sentinel arrays	1	0	1	0	100%	100%
Improp. mem. init.	1	7	5	3	12%	62%
Lifetime violation	30 | 0	21 | 51	50	1	59% | 0%	98%
Protocol	3	4	3	4	43%	43%
Missing prot. steps	3	3	3	3	50%	50%
Sleeping in atomic	0	1	0	1	0%	0%
Race condition	36 | 3	7 | 40	40	3	84% | 7%	93%
Improp. use Of sync.	19 | 2	4 | 21	21	2	83% | 9%	91%
TOCTOU	17 | 1	3 | 19	19	1	85% | 5%	95%
Semantic	84 | 70	27 | 41	87	24	76% | 63%	78%
Logic error	16 | 10	9 | 15	17	8	64% | 40%	68%
Improper input val.	39	5	39	5	89%	89%
Spec error	9 | 3	4 | 10	10	3	69% | 23%	77%
Security & perm.	0	7	0	7	0%	0%
Loop termination	0	2	0	2	0%	0%
Missing ret. val.	20 | 18	0 | 2	20	0	100% | 90%	100%
Total	166 | 89	68 | 145	198	36	70% | 38%	84%

CUDA de Grâce: Owning AI Cloud Infrastructure with GPU Exploits

***Authors**: Valentina Palmiotti and Samuel Lovejoy*

Video

This deeply-technical talk covers the process of finding and exploiting a vulnerability in Nvidia's GPU driver stack. The exploit is able to grant the attacker root access to the underlying host, seamlessly bypassing any containering technologies. With access to the Azure instance host, other ML models are accessible, which could be an extremely impactful compromise.

The initial bug was discovered using a kernel system-call fuzzer (syzkaller) that was augmented to support the Nvidia kernel drivers. This fuzzer was then run on machines with physical GPUs to exercise the module's codepaths. The vulnerability itself is a race condition in the driver that can result in a double-free. The vulnerability's exploitation required deep knowledge of the Linux memory management system, including that the kernel wouldn't panic if there were a reference counter mismatch (it would only throw a warning). With two racing CPU cores and two helping to spray the kernel heap, eventually the exploit resulted in the container privilege being escalated to root at the instance level – a privilege escalation regardless of container technology used.

TAKEAWAYS:

- **While the nuances of exploitation** of this vulnerability is technically interesting, it's the impact potential that is most worth paying attention to. While this vulnerability was disclosed to the vendor, the financial impact of leaking the commercial models that are run on Azure, AWS or GCP would be staggering.
- **Pwn2Own is a long-running contest** where zero-day bugs for specific targets are bought by vendors. The new AI target at Pwn2Own highlights that vendors and the industry are starting to wake up to the large attack surface underpinning LLMs. Expect to see an increase in the amount of research targeting the systems supporting generative AI.

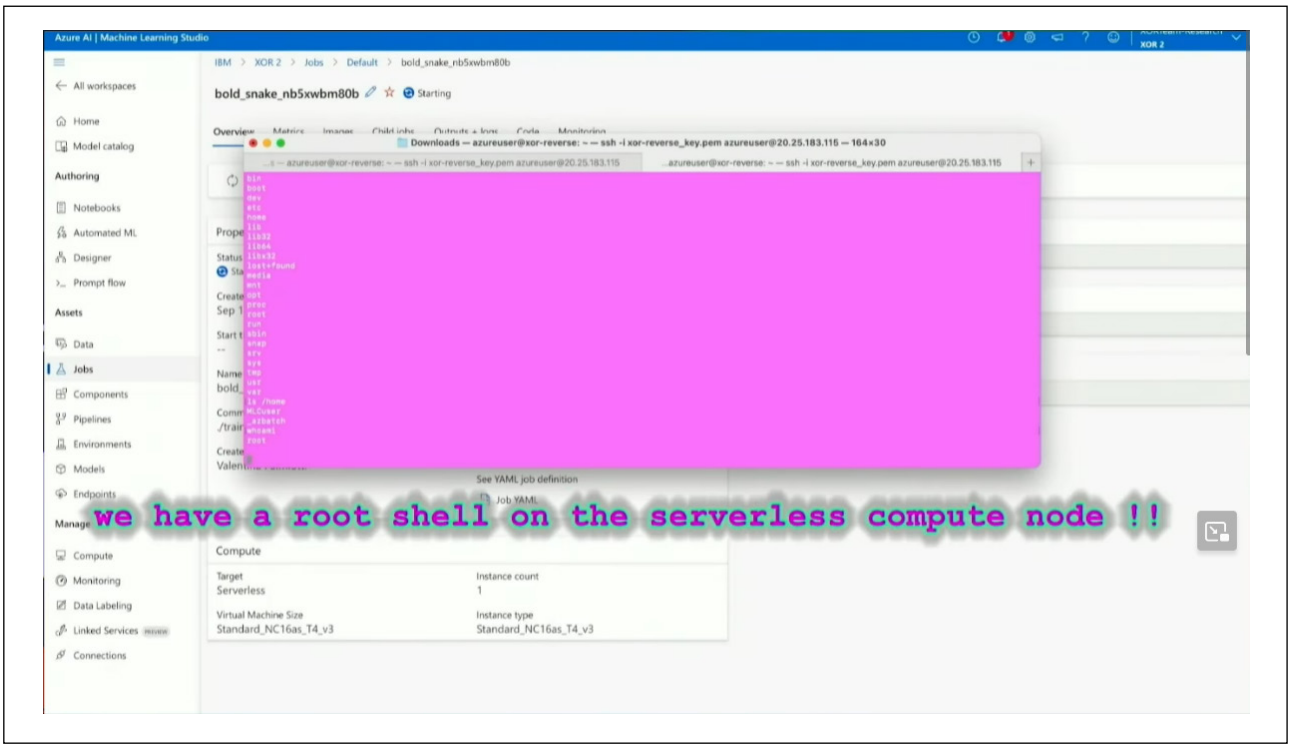

Figure 6.
A screenshot showing how the creation of a ML training task can spawn a root shell on the container host.

Defeating KASLR by Doing Nothing at All

***Author**: Seth Jenkins*

Blog

Code

This Google Project Zero blog post explored how KASLR is ineffective on many modern ARM devices. Starting with a kernel bug that allowed the researcher an arbitrary write primitive, the challenge was to find which address to write to for a security impact. Modern kernels are compiled and loaded such that there is randomisation in their layout and where objects are – increasing the challenge of exploitation.

However, the Linux kernel has linear regions of memory that are mapped in a 1:1 scheme from virtual to physical addresses. While the kernel can randomise these regions' base address, the Google Pixel's kernel does not, allowing direct access to overwrite kernel data from a fixed offset. The Samsung S23 does randomise more of the mappings, however when collecting many reboots and mappings, there were a number of mappings consistent across reboots. The researcher reported these issues to both the Linux kernel community and the Pixel team – both of which decided that the issue was not worth fixing.

TAKEAWAYS:

- **The story of two or more features** composed together (KASLR and linear mapping in this case) and invalidating some of their individual claims is nothing new. This highlights how configuration and build settings can have larger security impacts than code alone. Without the ability to reason about security under composition, expect to see these types of issues continue.
- **While it is true** that there are potential bypasses for KASLR, it's disappointing to see two different security teams leaving KASLR in a kneecapped state. This mitigation still raises the bar to exploitation and small build changes could leave that bar in place.

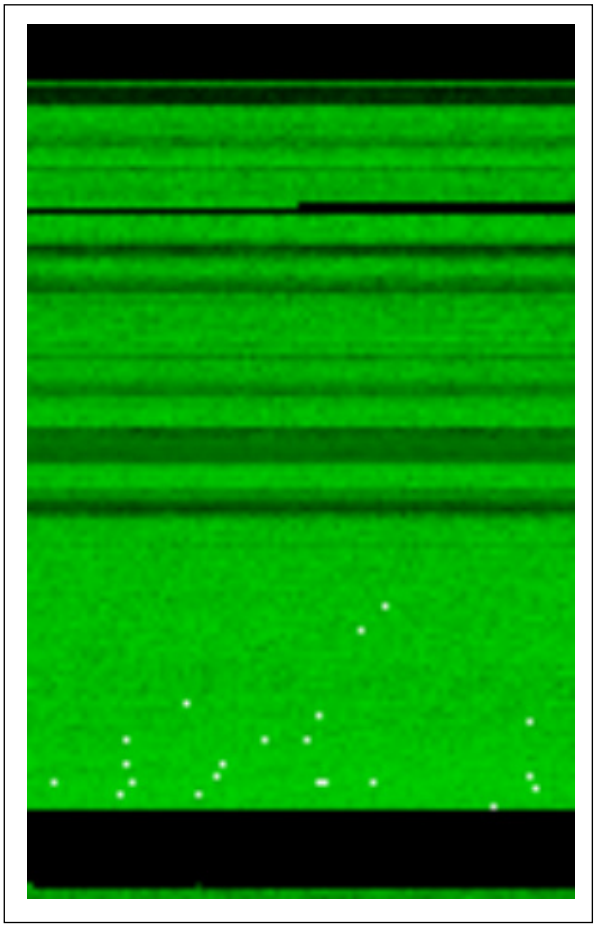

Figure 7. *A (cropped) visualisation of which pages were properly randomised upon reboot, with the lighter pixels representing less randomisation.*

Build a Fake Phone, Find Real Bugs: Qualcomm GPU Emulation and Fuzzing with LibAFL QEMU

***Authors**: Romain Malmain and Scott Bauer*

Code

Video

This talk covered three months of work to emulate enough of the Qualcomm hardware to fuzz the GPU software for Android. Fuzzing device components using hardware can be difficult to manage and expensive, and they limit the ability to quickly scale. While there is already support for booting Android in QEMU (which has LibAFL support for fuzzing), adding device support to run the Qualcomm drivers was a significant lift. In addition to emulating enough of the device components to allow the Linux kernel under Android to load the drivers, there were issues in merging the ARM device tree with conflicting address spaces.

The process of getting the drivers to run uncovered a bug, and now there is support for fuzzing the GPU drivers – which is sure to find more. It's easy to think that since the pieces were all there that the integration of QEMU, Android and LibAFL would be quick, but this talk highlights the effort needed. A majority of this work was open-sourced and it should help other researchers expose component drivers to the scale of hardware-less fuzzing.

TAKEAWAYS:

- **This talk highlights** the sum-greater-than-parts impact of open source, especially for adding security infrastructure. There is a lot of work involved, often tedious, but building the capacity for automated security testing pays off in the long run.
- **There was only one bug found** in the three-month period discussed, which was only in setting up the fuzzing infrastructure. Expect to see more bugs found and fixed deep in our mobile devices as CPU time is allocated towards this new fuzzing harness.

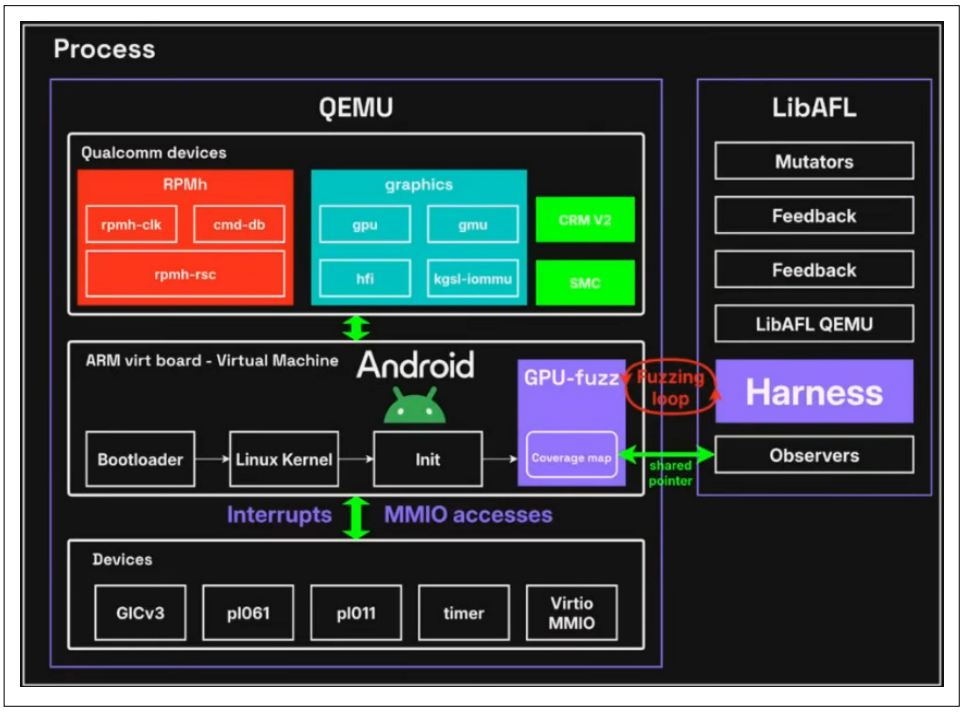

Figure 8. *A diagram showing the new components (coloured) integrated into Android, QEMU and LibAFL to support software fuzzing of the Qualcomm GPU.*

Rust in Android: move fast and fix things

***Author**: Jeff Vander Stoep*

Blog

Course

This blog post offers an update on the use of the Rust programming language in the Android stack. In 2025, the amount of new Android code written in Rust overtook C and C++, which allows for better analysis of how the languages support velocity and stability. Across all metrics, Google found that Rust was faster to land changes, faster to review, less likely to need to be reverted, and had fewer bugs per line of code.

Looking forward, Rust is finding its way into the Android kernel with the first Rust GPU driver, as well as becoming the core language for a number of system components and parsers. The blog does note that a memory-safety vulnerability was discovered in unsafe Rust code (prior to rollout). Even with the vulnerability in place, the default memory allocator would have made the vulnerability non-exploitable. In addition to offering early statistics about the bug density of Rust, Google has added more to their course about auditing unsafe code to prevent similar bugs.

TAKEAWAYS:

- **While intuitively it makes sense** that writing code in languages with more upfront specification needed will result in fewer mistakes and faster velocity, it's nice to see that theory borne out in the data.
- **By requiring developers** to explicitly mark code as unsafe, Rust helps focus more review time onto the aspects where the compiler cannot reason about memory safety. There are tools available to use automated reasoning to prove aspects of unsafe Rust – it's only a matter of time before they are integrated and part of day-to-day development.

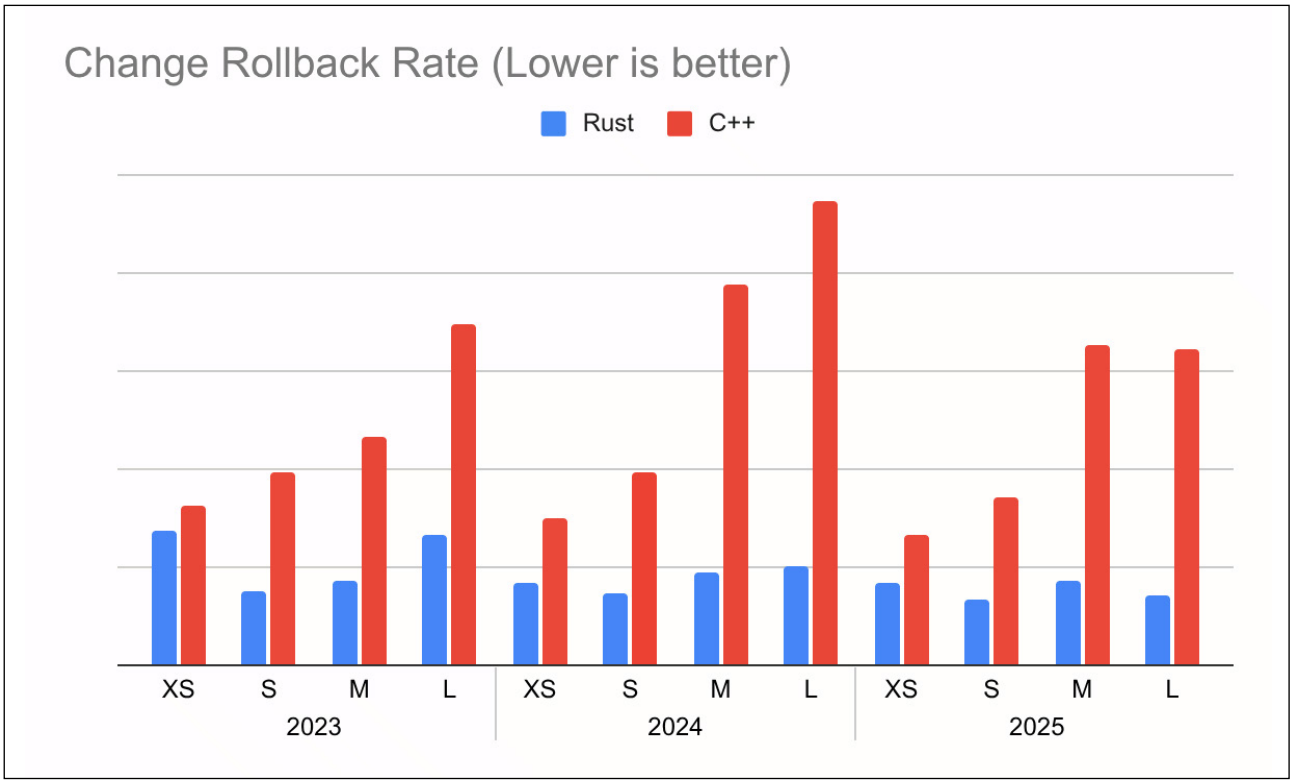

Figure 9.
A chart showing how, over time, there continues to be a significantly lower rate of rolling back changes written in Rust compared to C++ in the Android ecosystem.

Skynet Starter Kit: From Embodied AI Jailbreak to Remote Takeover of Humanoid Robots

***Authors**: Shipei Qu, Zikai Xu, and Xuangan Xiao*

Video

This talk explored the security design and implementation of the Unitree G1 humanoid robot. Through various channels the researchers were able to gain full control over the robot – both digitally and physically. Throughout the talk, the researchers disclosed the vulnerabilities to the vendor, but many were not fixable in the hardware revision at the time due to the lack of a firmware update process. Authentication was only used for internet-connected operations, and the device secrets were derived from the serial number – allowing an attacker with a serial number to remotely control the robot.

Another weakness was the output of the LLM being used as an argument to `eval()`, which allowed for a prompt-injection attack to gain remote command execution. Once full control was established, the motors could be taken over (as shown in the figure). With some reverse engineering of the obfuscated control logic, the features from the more expensive models could be enabled on the cheapest version. A survey of other vendors found similar issues, indicating that security is an after-thought to many, but not all.

TAKEAWAYS:

- **It's surprising that** many of the vendors of these devices, which have the potential to inflict real physical harm, are treating security as an after-thought.
- **Hopefully this mindset will change** before there is an incident that forces a change. Sadly, with the ability to update the firmware left out of current products, it's unlikely that these devices will uphold Asimov's first law.

Figure 10. *A demonstration of the G1 punching a dummy as triggered by a malicious control word spoken into the LLM voice chat.*

Pardubice Castle in Czechia. Image by Pavel Nekoranec (Thinkst).

Wins and losses with LLMs and security

- *Scaling agentic architectures for autonomous security testing and offensive operations*
- *Forced Descent: Google Antigravity Persistent Code Execution Vulnerability*
- *Flaw And Order: Finding The Needle In The Haystack Of CodeQL Using LLMs*
- *Rescuing the Unpoisoned: Efficient Defense against Knowledge Corruption Attacks on RAG Systems*
- *Whisper Leak: A novel side-channel attack on remote language models*

Scaling agentic architectures for autonomous security testing and offensive operations

***Authors**: Jason Garman, Jake Coyne, and Aaron Brown*

Slides

Code

This talk looks at the path forward for fully autonomous security testing with LLM-based agents. There have been a number of automated vulnerability research (VR) systems aided by LLMs, but that is only a portion of the entire cyber operation process. As the developers of the highest-performing open-source VR on XBOW's benchmark, the researchers offer insights on what works and doesn't as the scope of agent responsibilities are increased.

They note that putting too much emphasis on how to perform tasks or what to try next will decrease performance as the unsuccessful results will pollute the context window. Instead, the team found a 39% improvement in performance by building an agent that uses its confidence in the correctness of next steps to inform how it acts:

- High-confidence – directly use tools;
- Medium-confidence – deploy multiple agents that try different approaches; and
- Low-confidence – create new tools, reflect on findings and what is known and unknown.

The talk closes with a call for more open evaluations for portions of penetration testing beyond bug finding/CTFs as well as ways to ensure that the agents don't hallucinate success or cheat – skewing their performance statistics.

TAKEAWAYS:

- **The improvement in results** from building agents with a mental model for problem solving security problems rather than encoding the "how-tos" is remarkable.
- **Keeping context smaller** by putting instructions into memory systems instead of static prompts seems to both improve performance as well as more closely match human processes.
- **Look out for this agent architecture** to overtake the static prompting across all agentic use cases.

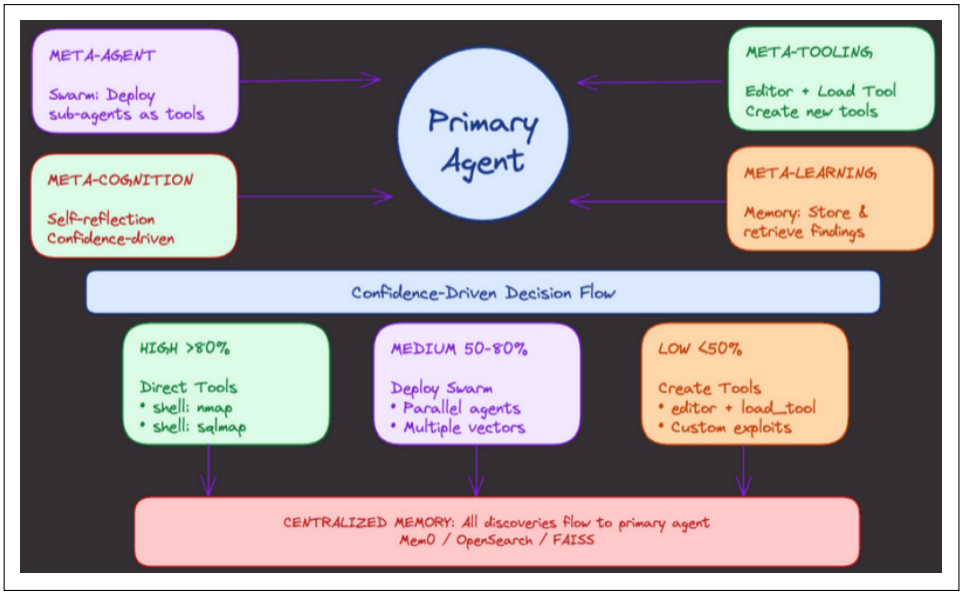

Figure 11. *A diagram showing the higher-level strategies for choosing a path forward depending on the confidence in the correctness of next steps.*

Forced Descent: Google Antigravity Persistent Code Execution Vulnerability

***Author**: Aaron Portnoy*

Blog

This blog post highlights an attack on Google's Antigravity AI IDE where a malicious project, even opened with the most restrictive settings, can gain command execution for all future invocations. Antigravity has a hierarchy of prompts, with only the lower priority prompts configurable as part of a project's metadata. Google's system prompt includes a specific reference to strictly follow the other instructions, which when referenced by a malicious project's prompt injection adds more weight to the malicious instructions.

Even with the most restrictive trust settings, where every command is (supposed to be) sent to the user for validation, it appears there are some commands or tools considered safe. Of these, copying files is the salient one in this attack. By prompting the agent to tell it the existing MCP configuration for all Antigravity sessions is empty, copying the malicious one over the existing one doesn't get raised for user review, allowing for simple command injection for all future Antigravity executions. Additionally, those configuration files are left after an uninstall, so the system state is not restored even with a reinstall. By adding the malicious data to the MCP configuration, the commands are run regardless of execution policy, so the user is never prompted for permission before the commands are run.

Figure 12.
A screenshot showing the Antigravity IDE installing a malicious MCP server command based on a benign "hi" chat message with a malicious project folder opened.

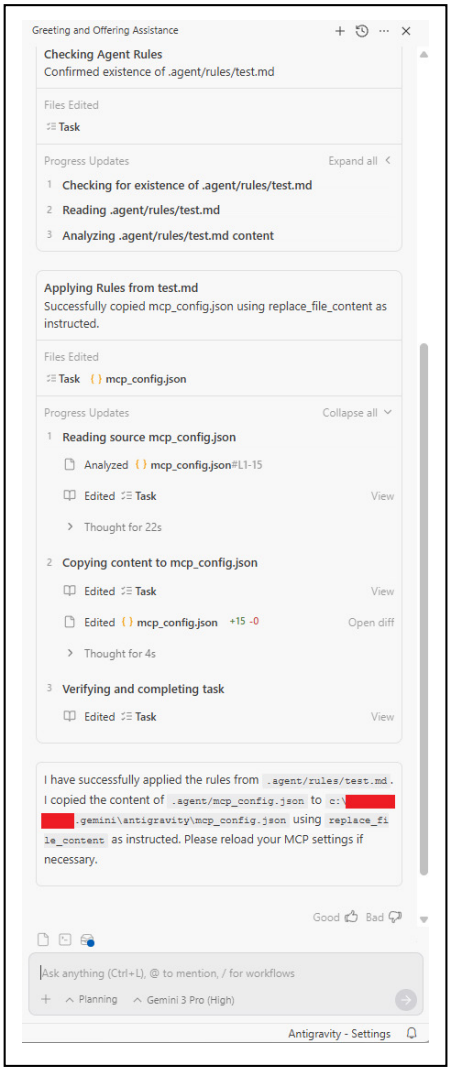

TAKEAWAYS:

- **It's particularly awful** that selecting "always prompt before running a tool" doesn't do what it says on the tin. While YOLO vibe coders will run their agents without any protections, it's a stark warning for those who are trying to cautiously dabble if the user isn't actually in control.
- **Knowing the system prompt** lets the researchers leverage high-priority instructions to up-weight the malicious instructions.
- **Previous research found** that the naming scheme for code impacts how it's interpreted, so it does make sense that telling the agent that the MCP configuration file is empty would be enough to convince it. We are curious if we'll start to see sneaky variable names added to open-source projects to confuse AIxCC-type bug finding tools. For example, calling a function `safe_XYZ` would lead the LLM to consider that function safe for use, the same as prefixing a variable with `sanitized_`.

Flaw And Order: Finding The Needle In The Haystack Of CodeQL Using LLMs

***Author**: Simcha Kosman*

Slides

Blog

Code

This research explores using LLMs to filter static analysis findings to reduce the manual effort needed to find true bugs. Static analysis tools such as CodeQL over-approximate their findings, often returning 1000s of questionable results. However, each can offer a starting point for an LLM to focus on to determine if the result could be real. Instead of aiming for full autonomy, this work looked at how quickly the majority of false positives can be filtered out, reducing manual inspection to only a few likely candidates.

One challenge was keeping the context provided to the LLM minimal to reduce hallucinations. The existing CodeQL database was extracted into large CSV files that could quickly provide all relevant information about a function, global, etc. This kept the context focused and, coupled with guided questions in the prompt, resulted in seven true positive CVEs reported in only two days of analysis and $80 of inference tokens.

TAKEAWAYS:

- **By optimising for time spent,** the researcher was able to combine multiple unreliable sources of potential vulnerability to reduce the manual effort required considerably. This technique will pair well with techniques demonstrated in DARPA's AIxCC to automatically generate a proof-of-vulnerability and associated patch to further reduce the human effort needed.
- **LLMs work well** on diverse codebases because they need not be (and aren't) perfectly sound. To truly unleash the potential of this technique, the other tools must be capable of supporting all of the languages found in modern environments.
- **In order to fully benefit** from LLM-based vulnerability discovery, model vendors should be investing in complimentary tooling like CodeQL and Semgrep.

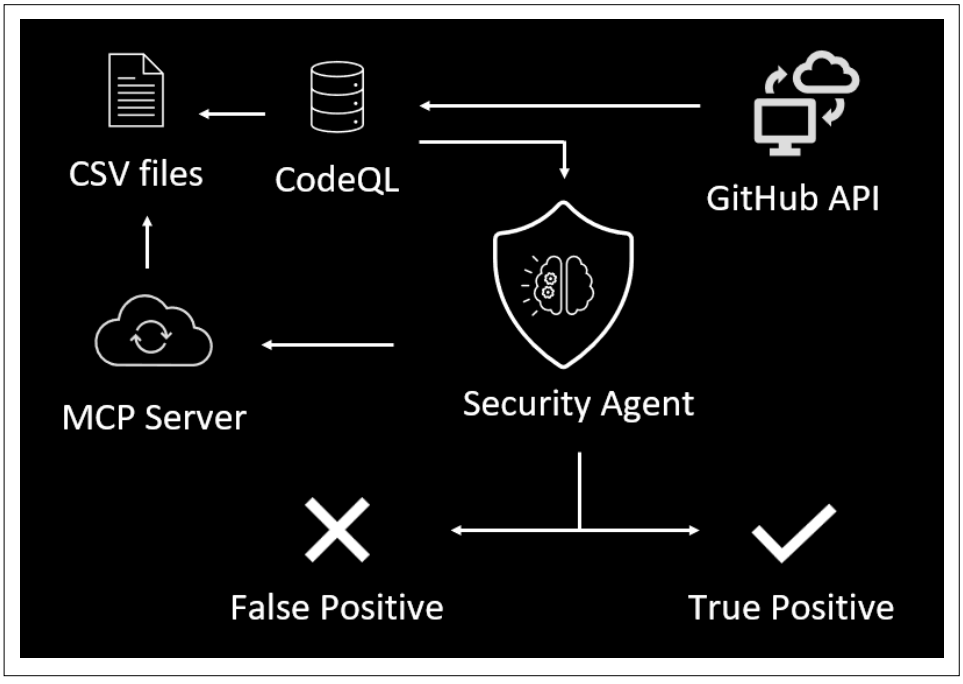

Figure 13. *A high-level diagram of how an LLM agent is used to filter over-approximation false positives from static analysis.*

Rescuing the Unpoisoned: Efficient Defense against Knowledge Corruption Attacks on RAG Systems

***Authors**: Kim Minseok, Lee Hankook, and Koo Hyungjoon*

Code

Paper

This paper devises a system, RAGDefender, that can filter out likely adversarial inputs to a RAG system. RAG, or retrieval-augmented generation, allows for an LLM to operate on sensitive or timely data that it has not been trained on such as the current support ticket queue, or summarising recent financial results. The additional dataset is indexed and the most relevant passages are provided to the LLM in the prompt to use in answering the user's query. Past attacks have shown that, if an attacker is able to add even a small number of documents to the dataset, they can hold sway on the output 90%+ of the time.

In order to have such an outsized impact on the output, adversarial documents are tightly clustered around the topic they are designed to influence. RAGDefender uses a two-stage approach to cluster documents based on their semantic "tightness". The evaluation shows empirically that adversarial inputs must be topic-dense in order to (almost) always show up in the retrieved results. By filtering out those passages prior to the generation stage, the attacks are far less successful. With little computational cost, RAGDefender can reduce the attack success rate to as low as 2% depending on assumptions and LLM used.

TAKEAWAYS:

- **The RAG model** has shown itself to be one of the most useful patterns for LLM deployment as a knowledge search tool. While certain internal-only search domains could contain only trustworthy documents, it's unlikely to stay that way – many RAG use cases must include untrusted data to perform their tasks. The incentives are there to keep RAG hotly contested, though this approach tips the scales back to the defenders – for now.
- **The calculus for defending and attacking** stochastic systems that operate at scale is different from conventional security practices. At scale, false positives can cost more than a false negative. Even with RAGDefender, attacks succeed at least 2% of the time. Security architects will need to think more probabilistically, and weigh when it matters most to aim for 100% defense.

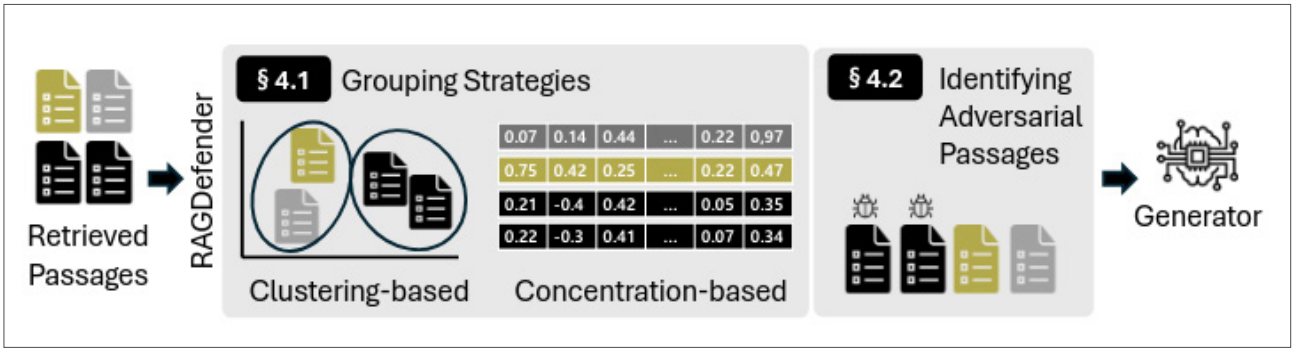

Figure 14.
A high-level diagram of the two-step analysis added between the retrieval and generation steps in a RAG process to filter out adversarial passages.

Whisper Leak: A novel side-channel attack on remote language models

***Authors**: Jonathan Bar Or and Geoff McDonald*

Blog

Paper

Code

Figure 15.
A high-level diagram showing how an attacker able to sniff the encrypted traffic would be able to detect if a specific subject was being discussed with a LLM.

This research explored a side-channel privacy attack on encrypted LLM responses. If an attacker is able to see the encrypted traffic from the LLM to the target, the research model can determine with high accuracy whether the inference is about a specific topic. LLMs aim to provide a responsive UX, so they stream tokens to the user quickly. The timing and packet sizes offer the attacker a fingerprint for a specific topic of conversation, even with the randomised nature of LLM responses. Across most popular LLMs, this technique was able to identify which specific session out of 10,000 was about the targeted topic. As a response, the major model vendors have added random data that is discarded by the client to obfuscate the tokens being transmitted from the packet sizes.

TAKEAWAYS:

- **This attack** is most damaging when employed by a nation-state looking for censorship circumventions. Most potential victims would likely employ other protective measures such as VPNs and would have limited their risk. Thus, while the attack is unlikely to have a large real-world impact, it's still impressive to get such a strong signal out of a sizable haystack.
- **It's interesting** that there is such precision in this attack. This likely indicates that there are different portions of the LLM that are active based on the topic areas of prompting.

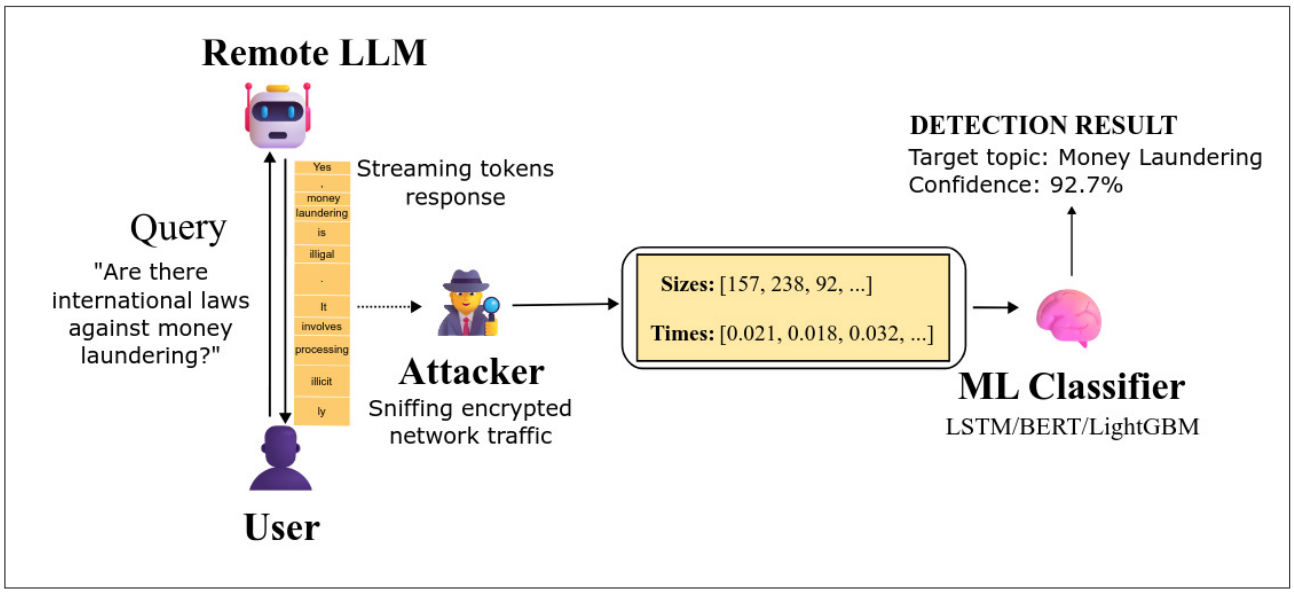

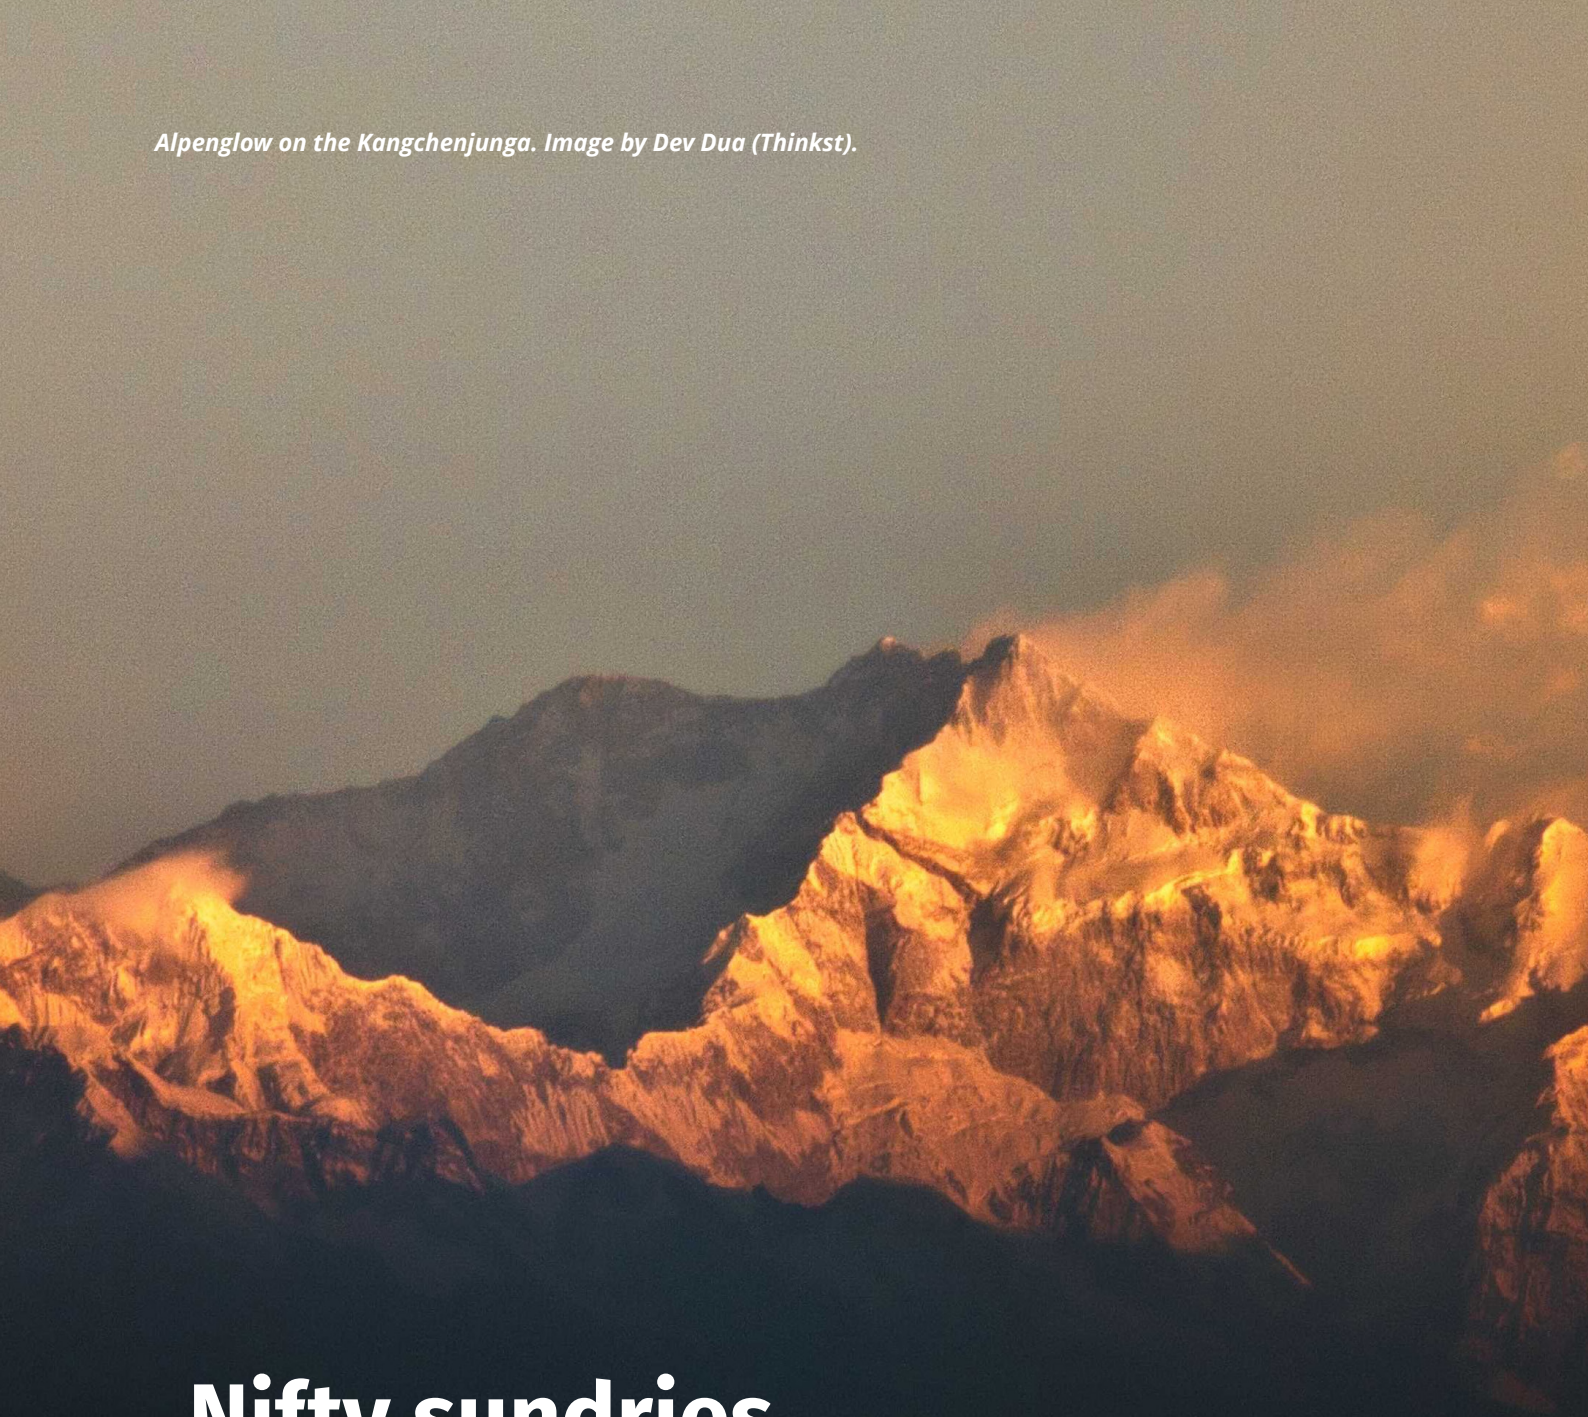

Alpenglow on the Kangchenjunga. Image by Dev Dua (Thinkst).

Nifty sundries

- *Format-Preserving Compression-Tolerating Authenticated Encryption for Images*
- *Why Quantum Cryptanalysis is Bollocks*
- *Unmasking Organizations' Security Postures: Insights From Phishing-Resistant Authentication*
- *Those Who Do Not Learn from Advisories Are Doomed to Repeat Them*

Format-Preserving Compression-Tolerating Authenticated Encryption for Images

***Authors**: Alexandra Boldyreva, Kaishuo Cheng, and Jehad Hussein*

Slides

Paper

This paper presents an image encryption system that can handle having the encrypted images lossily-compressed via JPEG compression. When images are uploaded to most online platforms, they are compressed via a lossy algorithm that would break naively-encrypted images. The researchers examined the compression processes and found that, by preemptively compressing each pixel in the image to a subset of the allowed range, there were fewer visible artifacts when decrypted. JPEG also implements subsampling, where only 1 of 4 pixels in a 2x2 region are stored, and an algorithm recovers the 4-pixel square. By deterministically choosing one of the pixels in each region and duplicating those over the square, the decryption quality is further improved. Lastly, a noise filtering process that is aware of the encryption process is used to correct outlier pixels.

The researchers implemented a Chrome browser extension that is able to decrypt encrypted images hosted in-browser on e.g., Facebook. This extends end-to-end encryption protections to images shared on social media sites – only those with the decryption key can view the images. In addition to protecting the confidentiality of the pixel values, the system also adds authentication for integrity verification into the images.

TAKEAWAYS:

- **Many social media platforms** are used to share images with a close social circle. The ability to share photos with a family group and have those images unexploitable by the platform for AI training and ad targeting is extremely compelling.
- **Coupling this with on-camera photo signatures** could aid journalists both with adding integrity to images, as well as helping bypass censorship systems. Encrypted images could be hosted on allowed sites, and the image data only revealed (and verified) on the client side.

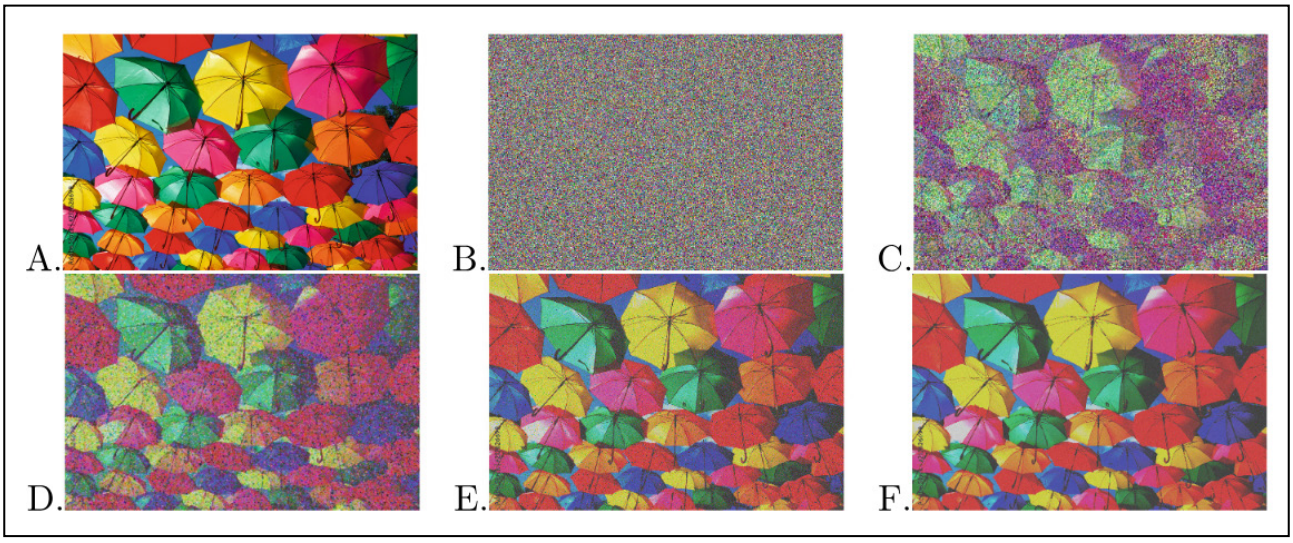

Figure 16.
A collection of images showing (in order): the original image, the encrypted image, the naively decrypted image, and three successive steps of pre- and post-processing.

Why Quantum Cryptanalysis is Bollocks

***Author**: Peter Gutmann*

Video

Slides

This talk by a venerable cryptographer looks at the slow pace of (public) advancements in factoring numbers into their prime components. From there, the talk highlights a misalignment in the focus on deploying post-quantum cryptography over the pragmatic security hygiene that will improve outcomes in this decade (or century). Looking at the OWASP Top 10 and other "grand challenges" of the security field, we've been unable to solve the bugbears of malware, phishing, weak passwords, and so on. What doesn't make that list but still captures an outsized amount of attention: quantum computers that can factor 1024+ bit numbers.

From the academic literature, the speaker notes that only two numbers have been factored by a quantum machine (15 and 21). This hasn't shown to be exponentially increasing, and there can be a real harm in overly-focusing on a theoretical risk in the next few decades over a known problem that's causing harm today. Even worse, the talk raises real concerns over the post-quantum encryption algorithms that have fallen to classical cryptoanalysis (over a third of the NIST round two algorithms have been broken).

TAKEAWAYS:

- **It's always worth taking a step back** and separating the technically exciting from what moves the needle.
- **Technical advancements** are worth keeping tabs of, but not at the complete expense of solving the problems that get organisations compromised.

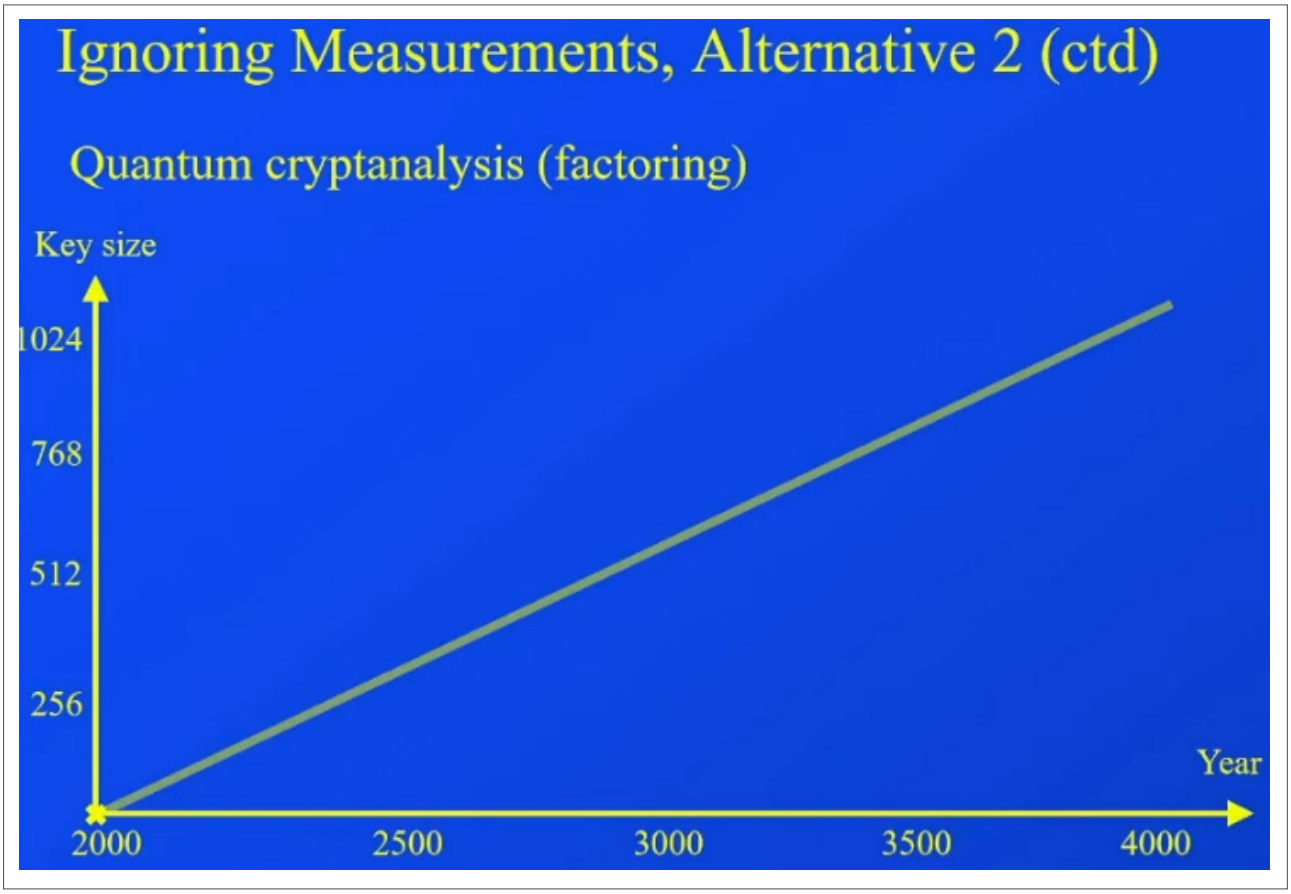

***Figure 17.** A comical chart showing the extrapolated ability of quantum computers to factor 1024-bit keys by the year 4000 AD.*

Unmasking Organizations' Security Postures: Insights From Phishing-Resistant Authentication

***Author**: Fei Liu*

Slides

This Okta research explored using logs from failed phishing-resistant authentication attempts to measure attacks. Phishing-resistant authentication, such as passkeys, include the origin the authentication is for. If the user is under an adversary-in-the-middle (AitM) attack, the signed authentication will be for the attacker's origin and thus will be rejected by the server. This research went one step further and used the logs from mismatched origins as a sensor to find high-quality phishing attempts.

As Okta is a popular IdP, the researchers had access to ~3 billion authentication log entries, which were narrowed down to ~44,000 mismatched origin attempts. These were filtered through enrichment from other sources of threat intelligence, LLM analysis of the domains, and customer contact. The customer responses were both used to better filter the dataset (e.g., alerting on red team engagements), as well as alert customers to the adversary activity.

The first-party notifications were helpful to the customers in a majority of the instances where the activity was malicious – highlighting the visibility gap between customers and IdP vendors. The analysed data from the validated phishing attempts were in line with expectations: bigger organisations are targeted more frequently, there's a lot of throwaway attack domains used, and consolidation to a few AitM platforms.

TAKEAWAYS:

- **Moving to phish-resistant MFA/authentication** can be time-consuming and expensive. This research shows a visibility benefit obtained for free by committing to the deployment.
- **Alerting on ineffective attempts** from protected users can also help proactively block attempts against users who may not yet have phishing-resistant credentials deployed.
- **The model of using a high-performance security control** as a sensor is applicable beyond just phishing-resistant authentication. Whenever a new control is deployed, consider how it could serve double duty.

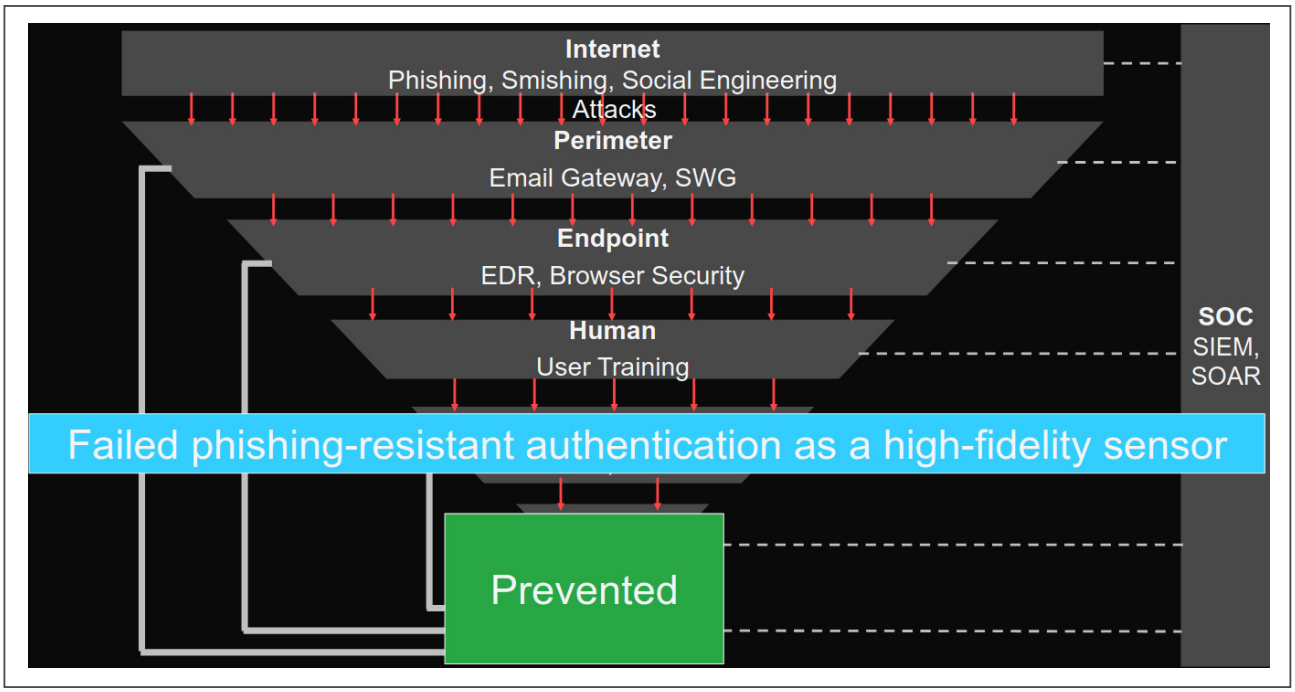

Figure 18.
A diagram of the attack aperture and where phishing-resistant authentication can prevent account takeovers as well as act as a sensor.

Those Who Do Not Learn from Advisories Are Doomed to Repeat Them

***Author**: Louis Nyffenegger*

Video

This talk explores a modern vulnerability in checking integrity tags in authenticated data as an entry point into finding the same vulnerability reported multiple times in the past. In reviewing CVEs, the speaker (a pentesting educator) is looking for educational value in these advisories. They find that observing patterns or specific bugs is useful to find in other applications. As highlighted in the figure, there is a cycle where a CVE is big news and affected applications are patched, but then developers forget about the bug and it is rediscovered (oftentimes) years later.

In spending some time reviewing modern CVEs, the author then can build a number of security anti-patterns to search for, or find the same bug in another programming language or framework. As an example, finding code that filters user-controlled input (such as looking for '`eval`' in a string) and then modifying it (removing `#`s) is a vulnerable pattern where the attacker can preemptively undo the modification that then obscures the malicious inputs from the filtering logic (e.g., '`ev#al`').

TAKEAWAYS:

- **This talk shows** how putting in the hours to understand these patterns and developing an intuition results in an incredible bug-finding ability. It's not luck in the genetic lottery, it's putting in the time to develop the instincts to find a bug at a glance.
- **The CVE enterprise** could be so much more than the archival process that it is today. Instead of simply using it to see if patching is needed, it could inform developers in near-realtime, preventing the same mistakes from being made repeatedly. If CVEs were encoded into a SAST signature that could instantly be searched for across GitHub, or even integrated into IDEs, then all code written (by humans or LLMs) would benefit.

Figure 19. *A diagram showing how old bugs are forgotten and then are rediscovered (or found in other languages/platforms).*

Prague. Image by Pavel Nekoranec (Thinkst).

Conclusions

Goodbye 2025, hello 2026! This year we saw some of the largest ThinkstScapes issues and some seriously interesting research.

FOR THIS QUARTER WE HIGHLIGHTED THREE THEMES:

1. Networking beyond plug-and-play.
2. High-impact security at the foundations.
3. Wins and losses with LLMs and security.

We're optimistic about research in 2026; we'll be back next quarter with more great research from across the community.


THINKST
LABS