ncc group

ADELARD part of ncc group

FOX IT part of ncc group

Annual Cyber Security Research Report 2025

World class cyber security research focused on real-world threats, vulnerabilities, and their impact on organisations and critical systems

Annual Cyber Security Research Report 2025

World class cyber security research focused on real-world threats, vulnerabilities, and their impact on organisations and critical systems

Introduction

The 2025 Research Report highlights a year of NCC researchers' contribution to expanding cyber security. Our cutting-edge cyber security research, from a global team of experts, continued to delve into AI, automotive security, hardware, home technology with experimentation, tooling, and engagement with industry and academia. NCC Group continues to look to expand experimentation and tooling, and to advance our cyber posture.

This year we published:

40 Articles

17 New or updated public tools

8 Public reports

8 Technical advisories

Contributions from Spain, the UK, the USA, the Netherlands, Singapore, the Philippines, and Australia

Over **1100** days of research

NCC Group research showcases broad and expanding cyber security research efforts across cryptography, hardware security, AI security, exploit development, and cyber-physical systems. Teams delivered dozens of public research outputs, technical advisories, blogs, tools, and assessments across emerging and established domains. A major theme was the rapid growth of AI-related risks from deepfake-enabled social engineering to insecure agentic-AI architectures alongside the need for formal security principles, cloud-AI hardening, and architectural controls. In parallel, established fields such as post-quantum cryptography, blockchain key management, and IoT/embedded security saw continued advancements, with researchers publishing new analyses, proofs-of-concept, and improvements to standards and tooling.

2025's key focus across practices was the evolving threat landscape shaped by increasingly complex systems, from authenticated automotive networks to quantum-enabled data centres. Transport security research emphasised system-level assurance and cross-modal methodologies, while exploit development efforts produced new techniques in automotive, USB, hypervisor, and embedded device exploitation.

Threat Intelligence documented real-world compromises-including ransomware intrusions, supply-chain attacks, SEO poisoning, deepfake vishing, and rapid social-engineering breaches-illustrating persistent failures in fundamentals such as credential hygiene and misconfiguration. These incident analyses reinforced the importance of defence-in-depth, stronger governance, and secure-by-design development across sectors.

Throughout the year, NCC Group researchers contributed numerous technical advisories spanning content management systems, QUIC implementations, IoT firmware, PDF libraries, and more, uncovering vulnerabilities ranging from XSS and memory corruption to cryptographic weaknesses. The organisation also continued extensive collaboration with industry and academia-supporting standards bodies such as C2PA, contributing to PQC adoption, and assessing critical technologies for Google, AWS, Meta, and VeChain. Public reports on confidential computing, secure AI services, and blockchain infrastructure demonstrated NCC Group's role in high-impact, security-critical reviews for global technology providers.

With continued investment in tooling, open-source contributions, and internal capability growth across practices, over 280 public GitHub repositories were updated or enhanced in 2025, enabling better analysis of kernel memory, blockchain protocols, MCP security, cloud posture, Bluetooth LE traffic, and adversary-simulation workflows.

The research programme reflects NCC Group's commitment to advancing cyber security through deep technical exploration, cross-disciplinary collaboration, and rigorous real-world assessments building a foundation for continued progress across AI security, hardware assurance, exploit development, cryptography, and emerging system architectures in 2026 and beyond.

We continue to move forward: aiming to make way for a safe and secure digital world for our global community.

The organisation also continued extensive collaboration with industry and academia – supporting standards bodies such as C2PA, contributing to PQC adoption, and assessing critical technologies for Google, AWS, Meta, and VeChain.

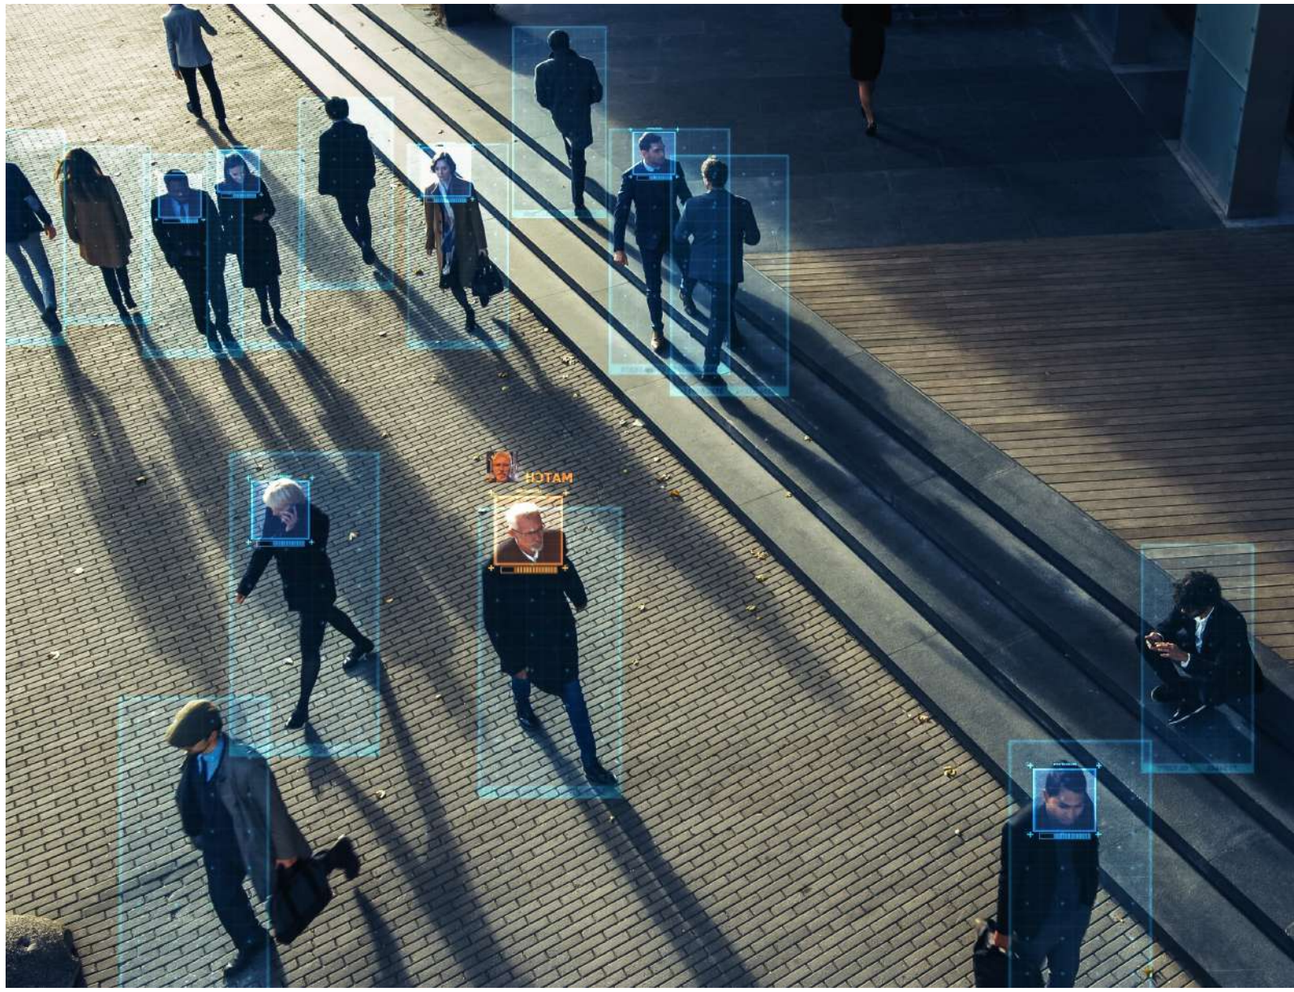

Overview of Our Research in 2025

Andy Davis
Global Research Director

2025 was a year defined by rapid change, deep technical progress, and a broadening of our research mission across every part of NCC Group. Across all practices, a clear theme emerged: security challenges are becoming more systemic, more interconnected, and increasingly shaped by fast-moving technologies - most notably, AI, quantum computing, cloud-scale cryptographic systems and cyber-physical infrastructure. Our research this year consistently demonstrated that defensive and offensive security now requires deeper interdisciplinary thinking and long-term, architectural approaches.

Rising Complexity and System-Level Security: Our Transport, Hardware & Embedded Systems, and Security Research Services teams all highlighted the same reality: today's risks stem less from individual components, and more from the interactions between them. Research into authenticated automotive networks, electrification ecosystems, and safety-critical systems showed that cryptographic maturity, supply-chain dependencies, and lifecycle governance now fundamentally shape cyber-physical risk. This was mirrored in our broader security assessments - many of the most significant compromises we analysed involved basic control failures (misconfigurations, weak credentials, and supply-chain blind spots) that became severe because of complex operational environments.

AI - From Novelty to Critical Infrastructure: 2025 was the year AI moved from experiment to production, and with it came a sharp rise in security-critical issues. Our AI research repeatedly showed that traditional guardrails are insufficient for agentic and multi-agent systems, and that secure deployment hinges on architectural controls, trust-boundary design, and data-code separation. We observed deepfake vishing becoming operationally viable, agentic AI gaining privileged access pathways, and organisations rapidly adopting standards such as MCP and ACP, without fully understanding the security implications. Our conclusion is clear: AI now represents one of the fastest-evolving and most consequential security domains, and NCC Group is well-positioned to lead the development of the methodologies, patterns, and frameworks needed to secure it.

Cryptography and the Post-Quantum Shift: Our Cryptography Services team continued to advance the state of applied cryptography, publishing research on post-quantum security, threshold signatures, secure messaging, and private processing systems for AI inference. Engagements such as our PQC reviews, blockchain key-management analyses and research into quantum-enabled data centres show a dramatic increase in industry demand for cryptographic assurance. The global transition to PQC is underway, and NCC Group's research is helping shape how organisations migrate securely while understanding the risks of "harvest now, decrypt later."

Exploit Development and Hardware Assurance: Our Exploit Development Group (EDG) delivered impact across automotive systems, embedded devices, hypervisors, and USB-based exploitation techniques, providing research that improves both vendor resilience and consultant capability. Hardware and Embedded Systems advanced tooling explored undefined behaviour in C, and deepened focus on secure firmware, OCP SAFE certification, and AI-assisted assessment methods. These areas of work underscore a broader trend: as systems become harder to exploit, research must become more specialised – requiring deep expertise, custom tooling, and long-term capability investment.

Threat Intelligence - Real-World Lessons from Active Adversaries: This year saw a rise in attacker sophistication but a continued prevalence of fundamental security failures. Our analyses of ransomware operations, supply-chain compromises, SEO poisoning, deepfake-assisted social engineering, and rapid remote-access intrusions all highlight a familiar message with new urgency: Misconfigurations and credential weakness remain dominant drivers of compromise, even as attackers leverage increasingly advanced techniques. Our incident investigations reinforce the need for defence-in-depth, high-quality detection engineering, secure-by-design principles, and a more resilient organisational security culture.

Safety and Assurance: Governance Meets Cyber-Physical Reality: Our Safety practice expanded guidance on Claims-Arguments-Evidence (CAE), assurance of AI/ML in regulated sectors, outcome-focused regulatory compliance, and the future of safety-critical digital systems. This research is increasingly aligned with cyber, reflecting a world where safety depends on cyber security and cyber security must be assured with the same rigour as safety. Across nuclear, automotive, medical, and industrial sectors, assurance is moving toward lifecycle-long, reasoning-based approaches - an evolution our research continues to help shape.

Collaboration, Tooling, and Open-Source Impact: We continued to deepen our collaboration with industry, academia, standards bodies, and global technology providers. Our contributions to C2PA, cryptographic standards, and high-impact reviews (Google, Meta, AWS) demonstrate NCC Group's role as an independent, trusted assessor of emerging technologies. More than 280 GitHub repositories were updated in 2025, expanding our capabilities across kernel analysis, protocol security, cloud posture assessment, MCP security, Bluetooth LE, blockchain protocols, exploit tooling, and red-team infrastructure. The scale and quality of our open-source contributions remain a defining strength of our research organisation.

Looking Ahead: The year ahead will challenge every organisation to secure systems that are more autonomous, more interconnected, and more adversarially-probed than ever before. The shift from human adversaries to AI-driven intruders, the persistence of unpatchable IoT and embedded devices, and accelerating regulatory expectations will demand new defensive models and deeper technical assurance. NCC Group's research community has shown that we are prepared, not just to keep pace, but to lead. Our global researchers continue to demonstrate the curiosity, rigour, and technical depth required to drive the field forward. In 2026, we will build on this foundation, expanding our capabilities across AI security, cryptographic assurance, secure system architectures, hardware security, and exploit research.

Together, we are shaping a safer digital world – through research that is bold, rigorous, collaborative, and deeply impactful.

Cryptography Services

Javed Samuel
Practice Director

NCC Group's Cryptography Services published cryptography research across a range of areas in 2025. Cryptography research is a key focus for the team as the field continues to evolve, with new developments, deployment challenges, and real-world applications emerging at pace.

- Published works by the Cryptography Services team include the following areas:
- Private Processing Systems for Artificial Intelligence Inference, where we aim for user data to still have the same privacy guarantees as they do with local-only computations.
- Post-Quantum Cryptography, including improved key generation, simpler and faster pairings, and auditing analysis.
- Threshold Signatures and Secure Messaging, such as verifiably encrypted threshold key derivation and Messaging Layer Security (MLS) in a Web3 environment.

As we enter the post-quantum era, Practice Director, Cryptography Services research goes beyond encryption and data protection to address emerging quantum threats, redefining digital trust, and enabling a resilient future. Through continuous research across this evolving landscape, Cryptography Services is committed to remaining at the forefront and delivering meaningful value to customers facing complex cryptographic challenges.

As data sharing increases - particularly with the greater use of Artificial Intelligence - cryptography research is advancing techniques that allow joint computation without revealing private inputs. This includes secure MPC for collaborative analytics, homomorphic encryption for computing on encrypted data, and zero-knowledge proofs (ZKPs) for privacy guarantees without revealing data. Our team continues to work actively on these areas as we endeavour to contribute to the advancement of the field.

Hardware and Embedded Systems

Catalin Visinescu
Technical Director

The Hardware and Embedded Systems team remained committed to delivering high-quality research throughout 2025, covering topics from IoT devices to cloud security. During the year, the team published numerous blog posts and tools, and took part in several conferences.

In 2026, we aim to continue our PowerG research. As additional targets, we plan to examine one or more reference bootloader implementations and analyse a selection of popular operating systems for vulnerabilities. Issues discovered in these areas often have significant impact due to their widespread use.

We also intend to present a comprehensive catalogue of undefined behaviour in C, along with practical guidance on how to avoid these pitfalls when building secure firmware.

Looking ahead over the next five years, AI-assisted security assessments are expected to play a growing role. Security and privacy challenges linked to AI adoption will become increasingly important.

Finally, OCP S.A.F.E. certification, promoted by major cloud providers to ensure secure, trustworthy hardware platforms across their infrastructure, is expected to become increasingly influential.

Security Research Services

Jon Renshaw
Director of Security Services

A number of themes have taken centre stage in our paid and collaborative research services in 2025, with some of the themes from previous years continuing. These include vulnerabilities in Internet of Things (IoT) devices, the security implications of the relentless development of quantum computing, building trust into digital media, and fostering relationships with some of the leading universities and collaborating on fundamental research.

Enterprise connected devices have been in the spotlight this year, with NCC Group research being published alongside a consultation by the UK Department for Science and Technology (DSIT) seeking views on the security of IoT devices. This consultation was targeted at business customers that may currently fall outside of the scope of the UK Product Security and Telecommunications Infrastructure (PSTI) Act, which came into force in 2024. Our research assessed the security of different classes of device a typical enterprise might connect to their network – such as meeting room panels, Voice over IP phones and IP CCTV cameras. The results were sobering, with many of the devices failing to meet multiple principles in the NCSC "Device security principles for manufacturers" and ETSI 303-645 "Cyber Security for Consumer Internet of Things: Baseline Requirements". In October, we presented the research at the ETSI Cyber Security Conference, alongside DSIT – with the timing proving serendipitous, as it coincided with ETSI launching their public repositories for the development of vertical standards for devices and software covered by the EU Cyber Resilience Act (CRA).

This research, our laboratory evaluations under schemes like Commercial Product Assurance (CPA), ioXt and Cyber Resilience Test Facilities, and our long history of independent research into device security, influenced the NCC Group response to the DSIT consultation, where we called for clearer incentives for product developers to invest in device security.

The security implications of Quantum Computing have been a hot topic within NCC Group research services during a year when many government cyber security bodies published their target timescales for migration to Post-Quantum Cryptography (PQC). For further information, you can get a thorough analysis from the 4th edition (October 2025) of our Global Cyber Policy Radar, which includes analysis from NCC Group's Government Affairs team, alongside interviews with our Cryptography Services lead, Javed Samuel, and Microsoft's Director for Cybersecurity Policy, Kevin Reifsteck.

Whilst the threat of a Cryptographically Relevant Quantum Computer (CRQC) is undoubtedly increasing as research into quantum computing continues at pace, it remains a relatively distant prospect by many estimations, with the primary hazard being the possibility of communications being intercepted now – under the assumption that they could be decrypted once a CRQC arrives (known as the "harvest now, decrypt later" threat). However, quantum computers are being developed, deployed and are available now through both private purchases and through shared Cloud infrastructure, with many organisations exploring how they could be used to solve problems too difficult or expensive to solve with classical computing. This year NCC Group concluded, alongside many collaborators from industry and academia, a project investigating the Quantum Data Centre of the Future (QDCF), where we shared our expertise on the threats to a quantum computing-enabled data centre, to ensure the developers and customers of a quantum computing service can adequately protect their sensitive intellectual property – critical in such a high tech and emerging field.

Exploit Development Group (EDG)

Alex Plaskett
Associate Director of Exploit Development

Over the past year, EDG primarily performed research into three main areas: Automotive, Embedded Device Security, and AI applied to vulnerability detection and exploitation.

In the Automotive space, the highlights of this were around presenting our prior Pwn2Own Automotive win against in-vehicle entertainment and electric vehicle chargers, as well as the media exposure around this. We continued working on automotive security throughout the year and developed novel USB exploitation techniques that we aim to present at an upcoming conference early in the new year. This research was inspired by real world exploitation of USB against mobile devices.

We also performed research on ECU security and published one of the outputs: an adb authentication bypass in Tesla's telematics control unit. We also have more ECU research to come – so, stay tuned!

In the Embedded Device space, edge device exploitation was found to be very topical in 2025. Many serious vulnerabilities were being found, as well as threat groups detected exploiting issues in the wild. To support the consultancy team, multiple internal exploits were created, which focused on edge networking devices (such as FortiGate and Citrix). This allowed consultants to really demonstrate the impact and allowed clients to quantify and understand the risks of compromised edge networking devices.

EDG also undertook multiple secondments throughout the year which led to an increase in the skills and knowledge of participating consultants. Brilliant research is often published after secondments in these areas have been completed, with consultants being able to apply this knowledge (such as VMWare Workstation host to guest exploitation!)

The third main area examined by EDG was the application of AI to both vulnerability detection and exploitation. In 2025, for EDG, this was primarily building our knowledge base of where AI could be most effectively applied in day-to-day activities and understanding the research published in this domain.

Looking forward

In the early half of 2026, we look forward to sharing more about some of the research outputs and findings from 2025. These findings have now reached a level of maturity where we can now share them with the security research community. The research areas examined in 2025 were all topical areas, and I don't anticipate any of these areas declining in 2026.

As a team, EDG looks forward to continuing to participate in hacking competitions (Pwn2Own and others) and further research into emerging technologies; this may include expanding into other events or challenges throughout the year. Whilst the time investment into these competitions is significant, the depth of research which can be performed, alongside the constantly evolving security postures of the targets, leads to great challenges and helps vendors in significantly enhancing their security posture.

Finally, I look forward to seeing the output from working in conjunction with the AI / ML Security practice on applying AI to common EDG tasks (which I will discuss more below).

Further forward

Year-on-year, we see more exploitation happening in the wild. While some of this could be explained by enhancements in detection technology, many publications have shown both increases in N-day and 0-day exploitation occurring.

For example, in 2024, Google (Mandiant) stated that exploits were the most frequently observed initial access vector, at 33%. However, what is particularly notable for those involved with effectively understanding the exploitation of vulnerabilities is that there is still a wide range of difficulty and maturity, based on target.

The more challenging targets require significant amounts of time to be invested, and often are a team effort. This is very different from a simple vulnerability, which could be weaponised in a matter of days or weeks. This is leading to the hyper-specialisation of researchers in these harder target areas, due to the massive amount of prior knowledge needed to be effective in the area or tooling development.

Overall, this is a huge credit to those vendors who are pushing the difficulty of exploitation to new heights and making their products more secure. However, this also really highlights that for a security company to really allow clients to both understand and defend against high-end threat actors, highly specialised technical research and teams like EDG is imperative.

It would be remiss of me not to mention AI and the emerging applications of it in both offensive and defensive security. Personally, I see great promise and already impressive outputs across the security research community. In 2025, EDG spent time building our understanding of these areas.

We designed a secondment (in conjunction with the NCC Group AI/ML practice) to use AI to investigate solving certain typical EDG tasks. This was to see which areas could immediately see benefits from the use of AI.

In 2026 and onwards, we hope to have tangible research outputs from this that we can share. At this stage, I will say that big swing projects like AIxCC and conferences like Offensive AI Con are bridging a cross-over between the ML practitioners and the traditional security research community, all leading to exciting outcomes.

One final observation is that the percentages of code being produced by AI are now well into double digits. Whilst these efficiency gains are commendable, this isn't without its security challenges.

Transport

Liz James
Managing Security Consultant

Across 2025, our research within the Transport Practice centred on understanding how system-level complexity shapes cyber-physical risk. Our analysis of legacy technologies across transport modes revealed that "legacy" is not simply a matter of technical age, but a compound effect of integration friction, inconsistent lifecycle management, and heterogeneous operational environments. These findings underscored a broader theme running through all our work this year: cyber security risk in transport increasingly emerges from interactions between technologies, not from isolated components. This insight informed our cross-modal methodologies and our continuing research into electrification, where reliable eHGV charging and vehicle infrastructure coordination depend on predictable, trustworthy behaviours across multiple independent actors.

A second major theme was the expanding role of structured assurance in transport cyber security, a key collaborative area of work with our Adelard colleagues. Across domains, from connected vehicle platforms to autonomy and UAS, we continued to explore approaches that move beyond checklist compliance toward Claims–Arguments–Evidence-driven reasoning, particularly exploring these workstreams as use cases for the ASCE software. Our research on Level 4 mobility highlighted how evidence fragmentation creates assurance debt, and our work on the Cyber Resilience Act demonstrated how transport organisations increasingly need defensible, regulation-agnostic assurance narratives. This direction was reinforced by our exploration of Euro 7's digital anti-tampering provisions, which link cyber security, software governance, and regulatory compliance into a single, continuous assurance problem. Together, these strands signal a clear industry shift toward lifecycle-long, reasoning-based assurance across transport technologies.

2025 was also the first time we, as a provider of independent testing and assurance services, had to fully confront the practical challenges introduced by the emergence of authenticated in-vehicle networks. While cryptographic standards like AUTOSAR SecOC and SAE J193991C have existed for some time, their increasing adoption in production-intent platforms meant that traditional penetration testing models were no longer sufficient. Systems that rely on authenticated messaging simply will not communicate without valid credentials, leading us to develop new tooling, new methods, and new ways of participating in a system's trust fabric without compromising its integrity. This led to early R&D into test domain PKIs, safe certificate handling mechanisms, and prototype J193991C state machine harnesses, all aimed at ensuring that independent assurance remains viable, as transport systems become more cryptographically mature.

Finally, 2025 saw significant internal capability development to support the growing breadth of transport cyber security challenges. We invested in maturing our testing frameworks, advancing methodologies for entropy assessment, authenticated protocol evaluation, and cyber-physical behavioural validation. We began the journey to strengthen our consultant capability through investment to build hands-on training environments, enabling teams to validate controls against documented intent, reason about system-level interactions, and communicate residual risk clearly. Our research into UAS assurance further broadened our transport remit, addressing safety, cyber security, data governance, and evidential integrity in a rapidly evolving operational space. These investments position the Transport Practice to meet the demands of 2026, where cryptographic networks, electrification ecosystems, cross-domain autonomy, and converging regulatory expectations will continue to drive the need for deep, system-wide assurance expertise.

Artificial Intelligence

Chris Anley
Chief Scientist, NCC Group

It has become increasingly obvious during 2025 that AI is changing business and society in many ways. Here are our top AI-related security takeaways.

Deepfake technology has now progressed to the point that deepfake video, images and voices are indistinguishable from the real thing. Deepfake voice clones, in particular, have proved to be extremely effective in our attack simulation work with customers. In the United States, the FBI has issued multiple warnings about the use of voice deepfakes in fraud. The UK Home Office has also issued warnings about Voice Cloning in its "Stop! Think Fraud" public information campaign. The crucial question for businesses is: are your staff empowered to resist a direct, voice-based instruction from a manager or executive? For example, to create an account, grant access to a system, or process a payment? Education, well-designed processes and solid governance are key to defending against these attacks.

Prompt injection remains a high-impact attack technique; this is amplified in an agentic system, where it may result in arbitrary code execution. Any content that appears in the context window of an LLM can cause a dangerous change in the output of the LLM; this has profound implications for AI security, as it can be extremely difficult for AI systems to safely process untrusted data.

Establishing a "secure" LLM, or a "secure" agentic framework has proven harder than expected. The "untrusted input leads to untrusted output" pattern means that guardrails are insufficient as a primary control in AI systems; architectural segmentation and data trust boundaries are needed to reduce risk in agentic systems.

Our research report, *The State of Supply Chain Security 2025*, highlighted AI as a top emerging risk.

- Supply chain risk now explicitly includes AI exposure. With increasingly widespread use of LLMs and even agentic systems in daily life, governance must address "shadow AI", where staff reach for the known and trusted AI systems they use personally, rather than the less familiar and (probably) less functional systems provided by the organisation they work for. We have seen incidents where fake AI applications have been used as the lure to persuade users to download trojans (an example is referenced in this report), and several similar supply-chain incidents have occurred over the past year in popular package managers.

- Mitigating supply-chain risk is a big topic; it typically involves a multidisciplinary team since there are legal risks, the issues touch on purchasing and contracts, as well as regulatory exposure and the more traditional cyber security issues of validating and scanning software and third-party software dependencies. The DNI and CISA in the US and NCSC in the UK publish excellent guidance on these issues.

The Future

The battle over the use of "AI" versus "ML" seems mostly resolved; ML lost. There is an old joke that ML is implemented in Python, and AI is implemented in PowerPoint; the joke is about 2 years old, which makes it roughly 20 in "AI years". Regardless, "AI" is the term used in the popular and scientific press, general academia and increasingly even among practitioners, so "AI" is the term that we must all now, somewhat grudgingly, embrace.

We are now all becoming used to the strengths and weaknesses of AI. Benchmarks are drawing the lines between various different commercial and open-source models. Agentic tooling and frameworks are emerging with a focus on different tasks – coding, writing, conversational user interfaces and generative systems for video, speech, music and graphic design. In the coming year, we expect to see these patterns become clearer as the industry continues to move toward greater commercialisation.

In security terms, AI has two main capabilities of interest: Automation and Imitation. In the latter part of 2025, we have seen agentic systems used to automate vulnerability research, infrastructure and web application scanning, and some limited initial reports of threat actor use of commercial AI systems. We expect this to be a significant growth area in the coming year, as these systems become more widely used and commercialised, and begin to approach the limits of the underlying technologies.

In terms of Imitation, voice and video deepfakes are now practically indistinguishable from the real thing. Having reached this pitch of near-perfection, there seems little room for improvement in the underlying technologies, but we expect to see much heavier use of these technologies in actual crime – from, voice deepfakes used in business and personal fraud, through to more widespread mis- and disinformation campaigns in the political sphere.

Regardless of which stage of grief you are currently at, in relation to AI, the technology is here to stay. Even with its considerable flaws, it offers clear business benefits in (marginally) improved efficiency, and (slightly) improved capability, and with these come some important shifts in the security landscape.

For better or worse, the future will be AI-assisted.

Safety

Peter Bishop
Chief Scientist, Adelard

Adelard, NCC's safety practice, specialises in the assurance of high-integrity safety-critical digital systems. Our research in the nuclear sector is highly regarded, and we have a long-standing programme of research projects sponsored by the nuclear industry, both nationally by the UK Control and Instrumentation Nuclear Industry Forum (CINIF), as well as internationally by the US Nuclear Regulatory Commission (NRC) and the UK's ONR We work across a range of other regulated sectors - including automotive and medical devices - and have developed guidance on security-informed safety and principles based assurance (PBA) for government agencies such as NPSA and NCSC.

The overarching goal of our research is to make it easier and more efficient for our clients to gain assurance that their systems satisfy stringent safety and security requirements, with a high degree of confidence. This is a multifaceted problem that needs to be addressed from a number of different perspectives. Complex systems are increasingly built using commercial-off-the shelf (COTS) components for reasons of cost and ease of availability, but such systems have an unknown pedigree and have not been subject to a rigorous development process or designed for safety-critical systems. Safety-critical systems are increasingly interconnected with other business systems, which makes them vulnerable to cyber-attacks. Systems are being built using novel methods and techniques such as Artificial Intelligence and Machine Learning (AI/ML), whose strengths and weaknesses are not well understood. Regulators are asking for more convincing assurance cases that focus on outcomes rather than prescriptive methods of compliance, which creates a skills gap for industry and a need for training and guidance in outcome- focused methods of assurance. As part of this trend, it is necessary to change the ways in which engineers think and create a more robust safety and security culture.

We work closely with both industry and the regulator to ensure that our research addresses current and future concerns about problematic aspects of assurance. We evaluate the approaches and methods we advocate with realistic case studies and industry trials, so that we can feed back lessons learned and improve our research outcomes. We also disseminate our ideas and learn about the latest academic research by attending conferences and workshops aimed at practitioners.

Claims-Arguments-Evidence

Claims-Arguments-Evidence (CAE) is a framework for reasoning about an assurance goal and for effectively communicating that reasoning. The use of such a framework helps to achieve a high level of completeness for a demonstration of safety or security and identify gaps and weaknesses. The technique is particularly effective for developing assurance cases that focus on "desired, measurable outcomes rather than prescriptive processes, techniques or procedures", in other words, "performance-based" or "goal-based" assurance cases. We have two long-standing research projects that are developing and evaluating guidance for using CAE to develop assurance cases for COTS components and safety cases for nuclear applications.

The Cogs (COTS Goal-based Safety assurance) project is a research programme funded by CINIF that focuses on the assurance of COTS components (such as smart devices or PC operating systems) for the purposes of regulatory compliance within nuclear safety and in particular in Instrumentation and Control (I&C) systems. Cogs aims to help nuclear safety engineers solve problems with their COTS assurance strategy by offering options for how to assess the impact of gaps in standard compliance and providing a route towards compensating for such gaps using the CAE methodology. This project has been running for almost 15 years and has developed templates and guidance documentation that the user community can directly use to justify COTS systems as part of the overall safety case. This year, the project is focusing on user training and support to help the community successfully adopt the approach.

The Declare (DEploying CLaims, ARguments, and Evidence) project is a research programme funded by CINIF aiming to support the deployment and uptake of the CAE approach for nuclear safety cases. The project has developed and trialled a suite of guidance documents on the application of CAE which have been made available to the user community (The Declare Guidance – CAE FRAMEWORK). This project has been running for 10+ years and is now focusing on engagement with nuclear licensees in the form of training and workshops to help them apply the CAE guidance.

In June 2025, we presented a paper about our experience with DECLARE at NPIC/HMIC'2025, the 14th Topical Meeting on Nuclear Plant Instrumentation, Control, and Human Machine Interface Technologies:

- Adopting CAE: A Deployment Approach for a Performance-Based Nuclear Industry

The CAE approach is not specific to nuclear systems and can be used to construct an assurance case for applications in other sectors that require formal regulatory approval. We already use CAE to provide assurance cases for medical devices and have recently started to explore the possibility of using CAE in the automotive sector to present the case for type approval by constructing an outcome-focused assurance case that demonstrates that the vehicle complies with all the relevant safety, security and environmental regulations.

We presented a poster and position paper about our ideas at SAFECOMP'2025 entitled:

- Using assurance cases to support type approval and regulatory compliance in the automotive sector

Assurance of novel tools and techniques

Novel technologies such as AI/ML need to be justified before they are deployed in safety-related systems in the nuclear sector. Similarly, the strengths and weaknesses of formal methods need to be properly understood before they can be used to make claims about the correctness of systems. The assurance of systems built using such novel tools and techniques is of concern to regulators and industry alike, and our research in this area is sponsored by both.

For example, the UK Office of Nuclear Regulation (ONR) are concerned about the assurance of systems that rely on AI/ML and commissioned us to provide guidance for assuring the safe deployment of AI/ML technologies for nuclear systems, which we expect to be published by ONR in 2026.

Similarly, the US Nuclear Regulatory Commission (NRC) wished to improve the efficiency, effectiveness, and flexibility of licensing reviews of systems developed using advanced "correct-by-construction" approaches, which are typically used for the highest-criticality Digital Instrumentation & Control (DI&C) systems (such as reactor protection systems for large light-water operating reactors), and commissioned us to provide a research study of such systems. The research study produced a set of basic assessment criteria for use by the regulator and the licensee who produces the safety assurance case (SAC). Guidance was developed on how to construct a SAC that demonstrates compliance with the criteria using safety case construction rules based on the CAE methodology. The SAC construction approach was illustrated by applying it to a "correct by construction" nuclear protection system example developed using formal methods.

The nuclear industry is also interested in methods that go beyond current verification and validation methods. Conventional verification approaches are primarily based on testing, which cannot guarantee a system is fault-free. Formal methods provide an alternative approach to demonstrating that an implementation meets its specification. This year, we undertook research for CINIF on the use of model checking to formally prove that a control system design satisfies specified safety requirements. If the proof fails, the model checker provides a "counterexample" showing how the safety requirement can be violated. The research study examined the current use of model checking in different countries and industry sectors, available model checking tools, and the feasibility of implementing the technique in the UK nuclear industry context.

Security-informed safety

Safety-critical systems are increasingly exposed to security threats that could compromise their safety, and this is recognised by safety standards such as IEC 61508 that require security threats to be assessed as part of hazard and risk management. Security-informed safety is about ensuring that security threats to cyber-physical systems do not pose an unacceptable safety risk in the physical world.

We have developed extensive guidance on security-informed safety for NPSA, including a code of practice for security-informed safety in the rail sector and the automotive sector. In particular, we were the technical authors of PAS 11281:2018, Connected Automotive Ecosystems, Impact of Security on Safety, Code of Practice, which is an international standard on road vehicles that gives recommendations for managing security risks that might lead to a compromise of safety in a connected automotive ecosystem. During 2025, we revised and updated PAS 11281:2018 to consider more recent standards and guidance on automotive cyber security, and we expect PAS11281:2026 to be published in March 2026.

We attended an IFIP WG 10.4 workshop on Cybersecurity of Transportation Systems and contributed two presentations to the session on Automotive Cybersecurity:

- Automotive Systems Engineering – Standards and Regulations
- What the history of functional safety can teach us about the future of cyber security in automotive

As part of a wider project to deploy an Advanced Modular Reactor (AMR) in the UK, we performed a security-informed hazard and operability study (HAZOP) of a WirelessHART demonstrator system, with the aim of identifying possible vulnerabilities and potential mitigations to inform the deployment of such technology in a nuclear I&C application. A paper about this work was presented at NPIC 2025:

- Digitisation of NPP I&C Systems via Wireless Technologies

Human factors in safety and security

Human factors are central to the safety and security of critical systems, as human performance can either contribute to system failures or act as a mitigation.

https://www.adelard.com/

Our research focuses on understanding how human behaviour, decision-making, and individual and organisational culture influence the reliability and overall performance of digital systems in complex applications. This research emphasises the need for integrating applied psychology, organisational theory, adaptive learning and ethical principles in the design, development and evaluation of computer-based systems.

During 2025, we presented our research at a RITICS workshop and to a meeting of the IEC 61508 association:

- Understanding the role of human factors in cyber security
- Safety culture in the nuclear sector – challenges and thinking about the future (focus on artificial intelligence)"

Future Research

Ristin Rivera
Research Manager

We're no longer dealing only with human attackers. AI-driven intruders are already probing systems on their own, hiding intent, and moving faster than any manual defence can respond. At the same time, personal privacy is getting harder to protect as smart glasses, covert sensors, and low-visibility cameras make it easy to record or profile people without consent.

Infrastructure risk is rising too. A huge portion of lower cost and white labelled devices can't be patched, replaced, or taken offline. These permanent vulnerabilities mean we can't rely on traditional cleanup-and-patch approaches; we must assume that some systems will stay exposed with no way to communicate this with the users.

On the human side, 2026 will bring real fatigue: notification fatigue, training fatigue, and testing fatigue. Social engineering will exploit that exhaustion, especially as employees become numb to the constant stream of "approve," "verify," and "confirm" prompts.

And finally, when every organisation deploys the same tools, exploits scale beautifully.

Technical Sections

AI Security (Architectures, Attacks, Operations and Tooling)

Security Tips For Your AI Cloud Infrastructure
Javier Garcia

21st February 2025

Guidance on securing AI services deployed in Cloud environments, noting that many AI-as-a-Service platforms are not secure by default. It emphasises applying core Cloud security principles, such as least-privilege identity and access management, network isolation to reduce unnecessary exposure, and hardening of supporting infrastructure like notebooks and training workloads.

It also highlights the importance of monitoring and logging AI workloads to detect misuse, model abuse, or cost-based attacks, as well as protecting sensitive training and inference data through encryption and privacy controls. Together, these recommendations reinforce that securing AI systems requires the same disciplined approach as traditional Cloud infrastructure, with additional attention to the unique risks introduced by AI workloads.

Analyzing Secure AI Principles
David Brauchler

10th March 2025

Secure AI design must account for the fundamental ways large language models (LLMs) differ from traditional deterministic application code. LLMs are essentially unstructured text-completion engines whose outputs are governed by the data supplied at inference time, meaning that threat actors who can influence input can profoundly manipulate model behaviour. Centralised designs that allow a single model unfettered access to all application functions significantly increase the risk of compromise, as LLMs may be coerced into privileged actions. Mitigating these risks requires architects to shift trust away from blanket access and instead limit model capabilities dynamically based on execution context, ensuring that models operate only on trusted data and have restricted access to sensitive functions.

Design principles such as data-code separation and least privilege are critical: Models should not be exposed simultaneously to untrusted input and high-trust functions, and permissions granted to models must be scoped tightly based on the nature of the data they receive. Additionally, avoiding excessive delegation of responsibilities to LLMs – especially for tasks better handled by deterministic code – and applying robust authorisation controls that mirror traditional session authentication help prevent escalation and misuse. These architectural controls form the foundation for reducing severe vulnerabilities in AI-integrated applications by restricting the influence of untrusted inputs over model outputs and function execution.

Analyzing Secure AI Architectures
David Brauchler

10th March 2025

Secure AI systems require architectural patterns that address risks not inherent in the models themselves but in how those models interact with application components and data. Traditional security assumptions fail in AI contexts because vulnerabilities arise from mixing trusted and untrusted data and exposing models to sensitive functions without segmentation. Strategies such as Gatekeeper, Orchestration Tree, and State Machine patterns isolate trusted operations from untrusted inputs by separating execution contexts, limiting model capabilities based on trust zones, and dynamically controlling access to sensitive actions. These designs prevent prompt injection and privilege escalation by ensuring that sensitive functionality is only invoked within tightly circumscribed, backend-managed contexts.

The recommended architectures enforce data-code separation and dynamic least privilege, so untrusted data does not directly influence high-trust decision points. For example, a gatekeeper pattern routes untrusted input through models without access to privileged functions, while orchestrator or state machine designs delegate tasks to specialised model instances with scoped capabilities. Careful management of data flows and trust boundaries minimises the impact of poisoned inputs and limits opportunities for exploitation. These patterns have been observed to mitigate whole classes of AI-related vulnerabilities when implemented correctly, emphasising that secure integration of AI technologies depends on thoughtful architectural design rather than surface-level guardrails alone.

When Guardrails Aren't Enough: Reinventing Agentic AI Security With Architectural Controls
David Brauchler

16th September 2025

Traditional AI guardrails, such as prompt filtering or output sanitisation, are insufficient to secure agentic AI systems that autonomously plan, execute, and adapt actions across tools and environments. The research argues that effective security must be built into the architecture of AI systems, applying controls that separate untrusted inputs from high-privilege operations, constrain where and how models can act, and enforce trust boundaries during execution. Patterns like Gatekeeper, Orchestration Trees, and State Machines provide structural ways to limit what autonomous agents can do, isolating critical functions and protecting sensitive operations behind controlled interfaces rather than relying on surface-level filters.

Applying these architectural controls ensures that even if an agent receives adversarial or malicious prompts, its ability to misuse privileges or access sensitive resources is restricted by design. This approach shifts the focus of security from reactive detection to proactive containment, treating AI decision processes as part of a broader secure system rather than as standalone components, and underscores the need for disciplined design practices when deploying agentic AI at scale.

Where You Inject Matters: The Role-Specific Impact of Prompt Injection Attacks on OpenAI models

Helia Estévez

7th May 2025

Prompt injection attacks against AI models are significantly influenced by the role context in which malicious input is placed within a conversation. In experiments using OpenAI's gpt-4o-mini, payloads injected into the system and assistant roles succeeded at much higher rates (86% and 92%) than identical payloads in the user role (52%), showing that attackers can more effectively manipulate model behaviour when they influence higher-privilege prompt contexts.

These findings emphasise that where external or untrusted data enters a prompt affects the likelihood of successful exfiltration or control of the model's responses. Development and deployment practices should therefore confine untrusted input to user-level roles and tightly control system and assistant context data to reduce risk. Additionally, more granular role structures may help clarify and defend against indirect prompt injection attacks.

Autonomous AI Agents: A hidden Risk in Insecure smolagents "CodeAgent" Usage

Ben Williams

28th July 2025

AI agents that run code autonomously and can allow large language models to generate and execute Python as part of their reasoning loop, enabling flexible multi-step logic and tool usage. While built-in sandboxes limit imports and execution scope, overriding these defaults with broad libraries can expose file system access and tool invocation paths that lead to remote code execution or unintended behaviour, if an attacker influences part of the agent's prompts. Proof-of-concept examples show how unsafe import allowances can let an agent read or write code on disk under certain conditions.

Secure deployment requires strict sandboxing, avoiding unnecessary import expansion, and running generated code in isolated environments – such as containers or remote executors. Monitoring execution and sanitising inputs helps reduce risk when agents generate and run code automatically, underscoring that autonomous AI systems with code execution capabilities must be carefully constrained to prevent serious security vulnerabilities.

Comparing AI Against Traditional Static Analysis Tools to Highlight Buffer Overflows

Mark Tedman

10th March 2025

The research evaluated the effectiveness of large language models (LLMs) versus traditional static analysis tools for identifying buffer overflow vulnerabilities in compiled binaries. The study compared results from a conventional static analyser (CWE_Checker), with LLM-driven analysis performed locally using Ollama, and introduced a custom Ghidra extension ("Revit") to integrate LLM assistance into reverse engineering workflows. While the static analyser reliably flagged the known overflow, both the tool and human analysts required deep C expertise to interpret findings and occasionally missed related issues such as improper string termination.

LLMs demonstrated strong capability in generating high-level insights and explaining code behaviour, but they also produced false positives and inconsistent vulnerability detection when analysing decompiled or incomplete code. The research concluded that LLMs are not yet reliable as standalone automated vulnerability detectors, but they can be a valuable supplement to traditional analysis-accelerating understanding and assisting analysts when used with careful prompt design and contextual validation.

Proxying PyRIT for fun and profit

Jose Selvi

15th January 2025

The research demonstrated techniques for improving visibility into PyRIT (Python Risk Identification Tool for generative AI), an open-source framework designed to automate adversarial testing of large language models (LLMs). The research details how proxying PyRIT's network traffic and instrumenting its internal request handling allows analysts to observe prompt injection attempts, multi-turn attack flows, and model evaluation logic in detail. Enhancements to PyRIT's proxy support, including a contributed code change, enabled systematic inspection of how adversarial prompts are generated, transmitted, and scored during testing.

The findings highlight the need for transparency and observability in automated AI security testing tools. Without insight into how these tools interact with target models, organisations risk misinterpreting results or overlooking failures in test coverage. By demonstrating how PyRIT can be proxied and analysed, the research supports a more reliable assessment of LLM security and strengthens defences against emerging threats, such as prompt injection, model manipulation, and misuse of generative AI systems.

HTTP to MCP Bridge

Jose Selvi

13th May 2025

HTTP to MCP Bridge is a tool that allows traditional HTTP-based security testing tools (like Burp Suite) to interact with remote MCP (Model Context Protocol) servers by translating standard HTTP requests into the asynchronous MCP transport mechanisms, such as HTTP+SSE or Streamable HTTP. The bridge acts as an intermediary HTTP server that establishes an MCP session with a target server and exposes a familiar HTTP interface for sending JSON-RPC messages, enabling easier enumeration and interaction with MCP services using existing web-security tooling.

This capability simplifies the security assessment of MCP implementations by removing the complexity of direct MCP transport handling, allowing analysts to use well-known HTTP workflows to test tool invocation and protocol behaviour. The project continues to evolve with added support for additional transport mechanisms, improved error handling, and enhanced testing features.

MCP Bridge Upgrade

Jose Selvi

17th November 2025

The HTTP-to-MCP Bridge tool, originally created to let standard web security tools interact with Model Context Protocol (MCP) services, has been significantly updated with enhancements including Streamable HTTP support, automatic detection of transport mechanisms (choosing between Streamable HTTP or HTTP+SSE), a revised / mcp/ endpoint, improved error handling, and a simple example MCP server for testing.

These improvements make it easier for security analysts to assess MCP implementations with familiar HTTP-based tooling by abstracting away the complexity of MCP's asynchronous transports while enabling broader protocol support and more robust testing workflows.

5 MCP Security Tips

Jose Selvi

16th May 2025

Securing Model Context Protocol (MCP) deployments requires specific controls to address risks inherent in model invocation and transport. Key recommendations include enforcing authentication and authorisation for all MCP endpoints to prevent unauthenticated use and using transport encryption (TLS) to protect context and response data in transit. Limiting exposed protocols and ports and placing MCP services behind access controls such as API gateways or zero-trust proxies, reduces unnecessary attack surface.

AI Webinars

We've had several webinars on AI topics this year, which you can view on demand, along with all our other webinars at:

View Webinars

AI Risk Management & HITRUST Certification: Ensuring Security & Compliance in AI-Powered Healthcare

15th May 2025

Watch Here

Securing AI: Cyber Resilience in an AI-Driven World

14th May 2025

Watch Here

Securing AI: Navigating the New Frontlines of Cyber Risk

12th June 2025

Watch Here

Securing AI: Building Confidence with Security, Safety, and Trust

10th July 2025

Watch Here

Cryptography & Blockchain

Announcing the Cryptopals Guided Tour Video 18: Implement CTR

Eli Sohl

10th March 2025

NCC Group's Cryptopals Guided Tour continues with Video 18, which focuses on implementing CTR (Counter) mode, a block-cipher mode that turns a block cipher like AES into a stream cipher. The guided video walks through the parameters and operational details of CTR - including handling nonces, keystream generation, and correct encryption/decryption - while emphasising that CTR is easy to misuse and that careful implementation is critical for secure cryptographic code.

VeChain JavaScript SDK Cryptography and Security Review

Paul Bottinelli, Gérald Doussot, Marie-Sarah Lacharité and Eli Sohl

16th April 2025

VeChain engaged NCC Group's Cryptography Services team to conduct a cryptography and security review of its JavaScript SDK. The SDK allows developers to interact with the VeChain blockchain and includes essential components such as cryptography and network-related functions.

State of the Art of Private Key Security in Blockchain Ops - 1. Concepts, Types of Wallets and Signing Strategies

Mario Rivas

2nd October 2025

Effective management of private keys is foundational to secure blockchain operations. Organisations must first choose whether to use custodial, non-custodial, or hybrid solutions, balancing control with operational complexity and risk. Private keys are used to sign blockchain transactions and authorise sensitive actions; how they are stored, accessed, and approved directly influences the security posture of a blockchain project. Three broad categories of wallets-hot, cold, and operational keys-reflect differing usage patterns and protective requirements, with cold storage prioritised for high-value holdings and hardened protections, and hot/operational keys designed for frequent interaction with automated systems. Multi-signature (multisig) and Multi-Party Computation (MPC) techniques provide greater assurance by distributing signing authority and reducing single points of failure. On-chain and off-chain mechanisms each bring trade-offs in enforcement, transparency, and flexibility.

State of the Art of Private Key Security in Blockchain Ops - 2. Common Custody Solutions Architectures

Mario Rivas

9th October 2025

Architecturally, custody solutions vary from centralised off-chain systems to MPC-based and on-chain multisig designs. A robust custody architecture applies Zero Trust principles to limit implicit trust relationships and continuously authenticate and authorise every action. Critical components include a dedicated signing module-segregating key usage from other services-and supporting modules such as transaction managers, storage layers, and node providers. Centralised signing modules present a high-impact risk if compromised; using secure hardware like Hardware Security Modules (HSMs), secure enclaves, or cloud Key Management Service platforms can isolate keys and reduce exposure. Poor key storage choices – such as in configuration files, databases, environment variables, or source code substantially increased risk.

State of the Art of Private Key Security in Blockchain Ops - 3. Private Key Storage and Signing Module

Mario Rivas

15th October 2025

The signing module itself must enforce proper authentication of signing requests, ensuring only approved transactions proceed to signing. Key storage should leverage hardware protections wherever possible and avoid mechanisms from which keys might be extracted. MPC and distributed key schemes further mitigate risk by requiring a threshold of parties to collaborate before a signature is produced, removing dependence on a single trusted signer. Software wallets and insecure mobile or browser storage should generally be avoided for operational private keys, though securely protected mobile enclaves may serve as part of distributed approval workflows.

State of the Art of Private Key Security in Blockchain Ops - 4. Approvals and Policies

Mario Rivas

23rd October 2025

Approvals and policy mechanisms ensure that sensitive blockchain actions occur only after appropriate review. Approval policies should default to deny-all and explicitly permit only business-required actions, with thresholds defined for different operations and roles. Processes should enforce policy changes themselves be subject to endorsement, and coordination systems must ensure unauthorised parties cannot propose transactions or policy modifications.

Combining transaction decoding and simulation aids human approvers in understanding what they are approving, reducing the chance of error. Whether enforced off-chain or on-chain via smart contracts, clear policies and secure approval workflows are essential to resilient key management and risk distribution, protecting against external threats, insider misuse, and operational mistakes.

Papers

Efficient Proofs of Possession for Legacy Signatures
Thomas Pornin, et al.

12th May 2025

Digital signatures underpin identity, authenticity, and trust in modern computer systems. Cryptography research has shown that it is possible to prove possession of a valid message and signature for some public key, without revealing the message or signature. These proofs of possession work only for specially designed signature schemes. Though these proofs of possession have many useful applications to improving security, privacy, and anonymity, they are not currently usable for widely deployed, legacy signature schemes such as RSA, ECDSA, and Ed25519. Unlocking practical proofs of possession for these legacy signature schemes requires closing a huge efficiency gap. This work brings proofs of possession for legacy signature schemes very close to practicality. Our design strategy is to encode the signature's verification algorithm as a rank-one constraint system (R1CS), then use a zkSNARK to prove knowledge of a solution. To do this efficiently we (1) design and analyse a new zkSNARK called Dorian that supports randomised computations, (2) introduce several new techniques for encoding hashes, elliptic curve operations, and modular arithmetic, (3) give a new approach that allows performing the most expensive parts of ECDSA and Ed25519 verifications outside R1CS, and (4) generate a novel elliptic curve that allows expressing Ed25519 curve operations very efficiently. Our techniques reduce R1CS sizes by up to 200 times and prover times by 3-22 times. We can generate a 240-byte proof of possession of an RSA signature over a message the size of a typical TLS certificate (two kilobytes) in only three seconds.

Falcon on ARM Cortex-M4: an Update
Thomas Pornin

29th January 2025

This note reports new implementation results for the Falcon signature algorithm on an ARM Cortex-M4 microcontroller. Compared with our previous implementation (in 2019), run time cost has been about halved.

Constant-Time Code: The Pessimist Case
Thomas Pornin

8th March 2025

This note discusses the problem of writing cryptographic implementations in software, free of timing-based side-channels, and many ways in which that endeavour can fail in practice. It is a pessimistic view: it highlights why such failures are expected to become more common, and how constant-time coding is, or will soon become, infeasible in all generality.

Improved (Again) Key Pair Generation for Falcon, BAT and Hawk
Thomas Pornin

11th July 2025

In this short note, we describe some further improvements to the key pair generation process for the Falcon and Hawk lattice-based signature schemes, and for the BAT key encapsulation scheme, in a fully constant-time way and without any use of floating-point operations. Our new code is slightly faster than our previous implementation, and, more importantly for small embedded systems, uses less RAM space.

Radical 2-Isogenies and Cryptographic Hash Functions in Dimensions 1, 2 and 3
Giacomo Pope, et al

5th May 2025

We provide explicit descriptions for radical 2-isogenies in dimensions one, two and three using theta coordinates. These formulas allow us to efficiently navigate the corresponding isogeny graphs. As an application of this, we implement different versions of the CGL hash function. Notably, the three-dimensional version is fastest, which demonstrates yet another potential of using higher dimensional isogeny graphs in cryptography.

Simpler and Faster Pairings from the Montgomery Ladder
Giacomo Pope, et al

7th July 2025

We show that Montgomery ladders compute pairings as a by-product and explain how a small adjustment to the ladder results in simple and efficient algorithms for the Weil and Tate pairing on elliptic curves using cubical arithmetic. We demonstrate the efficiency of the resulting cubical pairings in several applications from isogeny-based cryptography. Cubical pairings are simpler and more performant than pairings computed using Miller's algorithm: we get a speed-up of over 40% for use-cases in SQIsign, and a speed-up of about 7% for use-cases in CSIDH. While these results arise from a deep connection to biextensions and cubical arithmetic, in this article, we keep things as concrete (and digestible) as possible. We provide a concise and complete introduction to cubical arithmetic as an appendix.

Presentations

Post-Quantum Cryptography Auditing
International Cryptography Module Conference (ICMC), Javed Samuel

10th April 2025

With the recent standardisation of some post-quantum cryptography algorithms, our Cryptography Services team has had the opportunity to review various implementations. During this talk, Javed highlighted the categories of vulnerabilities reported in post-quantum cryptography implementations. These include:

- Out-of-Bound Memory Read
- Non-Constant Time
- Invalid Parameters
- Missing Validation
- Missing Memory Zeroization

The talk also provided development tips and strategies for implementing and reviewing post-quantum cryptography implementations.

Threat Intelligence & Social Engineering

[Weak Passwords Led to (SafePay) Ransomware…Yet Again]
DFIR team

10th March 2025

A ransomware attack attributed to the SafePay group began with a misconfigured Fortigate firewall that allowed threat actors to bypass VPN MFA and gain initial access, compounded by a weak domain administrator password that facilitated privilege escalation. Once inside, the attackers deployed a QDoor backdoor and a ScreenConnect service for persistence, moved laterally via RDP and SMB, and ultimately executed the SafePay ransomware to encrypt servers and hypervisors, appending the ".safepay" extension to impacted files.

The incident highlights how basic security failures, misconfigurations and weak credentials continue to enable severe compromises, allowing threat actors to harvest credentials, deploy ransomware with anti-recovery techniques, and disrupt environments. Rigorous firewall policy management, strong password enforcement, and multi-factor authentication are essential controls to reduce exposure to similar ransomware campaigns.

[In-Depth Technical Analysis of the Bybit Hack]
Mario Rivas

10th March 2025

The Bybit hack on 21 February 2025 involved attackers stealing over $1.4 billion from Bybit's cold wallet by compromising the Safe{Wallet} multisig interface and altering its JavaScript. Instead of stealing keys, the attackers tricked authorised signers into approving a malicious transaction that changed the wallet's contract implementation, giving the attackers control.

The incident highlights that supply-chain compromises and "blind signing" of opaque transactions can defeat multisig protections, and it underscores the need for stronger transaction verification, clearer signing interfaces, and tighter security controls around wallet tooling.

[Masquerade: You Downloaded ScreenConnect not Grok AI!]
Molly Dewis

18th July 2025

A malicious campaign delivered a trojanised ScreenConnect remote access tool disguised as a Grok AI download via a Facebook advert that led the victim to a fake site. After the user clicked the ad and downloaded what appeared to be an AI-branded file, a ScreenConnect installer was executed, giving the attacker persistent remote access and a channel to transfer additional malicious payloads.

Among these was AsyncRAT, which enabled credential theft via keylogging and browser data collection from the compromised host.

The incident highlights how social engineering and disguised legitimate tools can be used to circumvent trust and deploy remote access malware. It underscores the importance of user awareness about malicious lures, avoiding storage of plaintext passwords, and disabling browser features such as autofill that can facilitate credential harvesting.

[Rapid Breach: Social Engineering to Remote Access in 300 Seconds]
Rodrigo Munoz

5th August 2025

Threat actors used social engineering to swiftly compromise corporate systems by impersonating IT support and convincing users to grant remote access via Windows' QuickAssist tool. Within under five minutes, the attackers executed a sequence of PowerShell commands that downloaded and decrypted a payload hidden in an image file, deployed a remote management tool NetSupport Manager, a malicious DLL, and established persistence through registry and scheduled task mechanisms. During the session, they also employed a GUI prompt to capture user credentials.

The incident demonstrates how quickly a targeted social engineering attack can escalate to full remote access and malicious footholds when legitimate tools are abused. It underscores the need for robust user awareness training, strict controls around remote support tools, and rapid detection/response capabilities to prevent similar breaches.

[Voice Impersonation and DeepFake Vishing in Realtime]
Pablo Alobera, Pablo López, and Víctor Lasa

30th September 2025

Recent advancements in AI have made real-time voice cloning a practical tool for social engineering and vishing (voice phishing), enabling attackers to impersonate trusted speakers with high realism using only a few minutes of audio. NCC Group's research shows that AI-driven voice deepfakes can be generated with minimal latency and convincing quality even on modest hardware, and when combined with tactics like caller ID spoofing, such impersonations can deceive targets into divulging sensitive information or performing actions they believe come from colleagues or authority figures.

This development significantly expands the threat landscape by lowering the barrier to highly convincing scams that exploit human trust in familiar voices.

Organisations should treat voice communications as inherently untrusted, reinforce multi-factor verification methods, and train staff to recognise and respond to unusual requests – even when delivered in a familiar voice – because AI-enabled vishing techniques are increasingly accessible and effective.

[Fake CAPTCHA led to LUMMA]
Molly Dewis

25th November 2025

A threat actor used a fake CAPTCHA page to trick a victim into executing hidden PowerShell commands that downloaded and installed the LUMMA infostealer malware on a Windows host. After a user visited an initial compromised site, a sequence of malicious redirects led to the fake CAPTCHA, which employed social engineering (ClickFix technique) to get the user to open the Run dialog and execute obfuscated PowerShell that fetched and launched a malicious ZIP payload. The installed malware quickly connected to a command-and-control server and harvested sensitive browser data.

The incident demonstrates how deceptive web content can rapidly turn routine browsing into a compromise, with attackers bypassing traditional protections and leveraging legitimate system tools (PowerShell, mshta.exe, AutoIt) to evade detection. Rapid execution and credential theft underscore the need for user education, layered defence, and aggressive blocking of malicious redirects and social engineering lures.

[Black Hole of Trust: SEO Poisoning in Silver Fox's Space Odyssey]
Dillon Ashmore and Asher Glue

18th December 2025

An advanced threat actor tracked as Silver Fox ran a long-running SEO poisoning campaign that manipulated search engine results to promote malicious download sites hosting trojanised installers for popular applications (e.g., communication tools, VPNs). By optimising poisoned pages to rank highly for legitimate software queries, the campaign lured users into downloading backdoor-laden installers, exposing victims to remote access malware without exploiting software vulnerabilities.

Analysis of the actor's exposed telemetry, infrastructure, and tactics revealed false-flag elements and targeting primarily of Chinese-speaking users, with impact extending to Asia-Pacific, Europe, and North America. The campaign deployed variants of ValleyRAT and highlights how threat actors exploit user trust in search results rather than technical flaws, underscoring the need for vigilance against SEO-based distribution and harder controls on software acquisition.

Enterprise, Cloud & Containers

Defending Your Directory: An Expert Guide to Combating Kerberoasting in Active Directory
Rafael Alfaro March and Rodrigo Munoz

10th March 2025

Kerberoasting is an attack technique targeting the Kerberos authentication protocol used in Microsoft Active Directory, where an attacker with domain access requests service tickets tied to service account passwords and then performs offline cracking against the encrypted ticket hashes to recover clear-text credentials. Because service ticket requests are legitimate operations, attackers can exploit weak encryption, such as RC4 or poorly chosen service account passwords, to minimise detection and escalate access within the environment, potentially leading to lateral movement or broader compromise. The attack leverages tools like Impacket's GetUserSPNs and Rubeus and is mapped to MITRE ATT&CK T1558.003.

Effective mitigation focuses on hardening service account security and authentication mechanisms: enforcing long, complex, and regularly rotated passwords; minimising privileges and unnecessary Service Principal Names (SPNs) - by using modern Kerberos encryption such as AES256 instead of weaker ciphers and restricting delegation to the least necessary scope. Monitoring and hunting for anomalous Kerberos ticket activity using Windows Security logs (e.g., Event IDs 4768/4769) can help surface suspicious patterns, though careful interpretation is required since normal operations generate similar events. Combining prevention, detection, strong configuration policies, and ongoing audits strengthens resilience against Kerberoasting and supports a robust Active Directory security posture.

EAP-TLS: The most secure option?
Apostolos Gioulis

9th April 2025

EAP-TLS is presented as a strong wireless authentication method that uses mutual certificate-based authentication instead of passwords, eliminating common credential-theft vectors and making it one of the most secure options for WPA2-Enterprise networks. It relies on digital certificates for both clients and servers and establishes a TLS-based handshake to authenticate peers, reducing exposure to brute-force and phishing attacks.

Practical risks arise when username privacy is not protected during the initial identity exchange and when clients are misconfigured to accept untrusted server certificates, which can enable evil-twin style attacks if users approve bogus certificates. Ensuring strict server certificate validation and limiting clients to trusted Certificate Authorities mitigates these issues, reinforcing EAP-TLS's security benefits when correctly configured.

Your point of departure for forensic readiness
Michael Alexander Heenes

2nd October 2025

Forensic readiness is the capability of an organisation to collect, preserve, and analyse digital evidence in a way that supports incident response, investigations, and legal proceedings without disrupting normal operations. Achieving readiness begins with clearly defined objectives-understanding what evidence is needed, why it is required, and how it will be used, so that data collection and retention strategies align with investigative and compliance goals. The article emphasises that readiness should be designed proactively, rather than retrofitted after an incident occurs, to avoid costly gaps and reactive firefighting.

Key elements of forensic readiness include establishing policies for secure logging, audit trails, and chain of custody; training staff on forensic-safe procedures; and integrating forensic considerations into system architecture and incident workflows. This approach enables organisations to respond swiftly and effectively to security incidents and legal inquiries, while preserving evidentiary integrity and reducing risk exposure.

Auditing K3s Clusters
Andrew Wade

10th March 2025

K3s is a lightweight distribution of Kubernetes designed for resource-constrained environments. While K3s maintains core Kubernetes functionality, its simplified architecture-such as bundling control plane components into single processes and using a sqlite state store instead of etcd-alters how audit data and configuration should be collected and reviewed. Default locations for configuration manifests and the potential absence of those files post-deployment can complicate traditional audit workflows, requiring alternative methods like querying systemd journal logs to reconstruct runtime parameters.

Differences in certificate management and storage backend behaviour are highlighted that affect audit procedures. Because K3s certificates are shorter lived and automatically rotated and because its sqlite database lacks standard CLI tooling, auditors must adjust their techniques compared with standard Kubernetes clusters. The guidance underscores that while K3s offers deployment and operational advantages, security auditors must tailor their approach to effectively gather and analyse evidence, ensuring comprehensive review and forensic readiness.

Azure Fabric Backdoor With A Twist
Viktor Gazdag

21st October 2025

A proof-of-concept backdoor was demonstrated against Azure Fabric, Microsoft's cloud-based analytics platform, showing how built-in features like notebooks and the Activator event engine can be abused to run malicious code that creates persistent access. By embedding Python code in a notebook and triggering it through workspace events monitored by Activator, the backdoor uses the Azure Python SDK to provision new resources – such as a virtual machine with SSH access, a managed identity, and an elevated service principal-thereby extending control beyond Fabric into the broader Azure subscription.

OCP S.A.F.E. How-to
Aaron Kondziela, Diana Dragusin, and Evan Anderson

24th June 2025

The OCP Security Appraisal Framework and Enablement (S.A.F.E.) provides a standardised process for auditing firmware security in hardware devices, offering assurance about the security of code running on those devices rather than a pass/fail certification. It begins with creating or evaluating a threat model, followed by a structured security review broken into three increasing scopes – component vulnerabilities (Scope 1), subsystem interactions (Scope 2), and physical resilience testing (Scope 3) – to uncover..." and address security issues. After the review, the security reviewer delivers detailed findings and remediation guidance, and the vendor can request publication of a Short Form Report (SFR) to document the assessment.

The process includes retesting after fixes and can be repeated for new firmware versions to ensure ongoing security assurance. By providing a common evaluation methodology for firmware and hardware security, OCP S.A.F.E. enables vendors to demonstrate transparency and helps device consumers make informed decisions about product security.

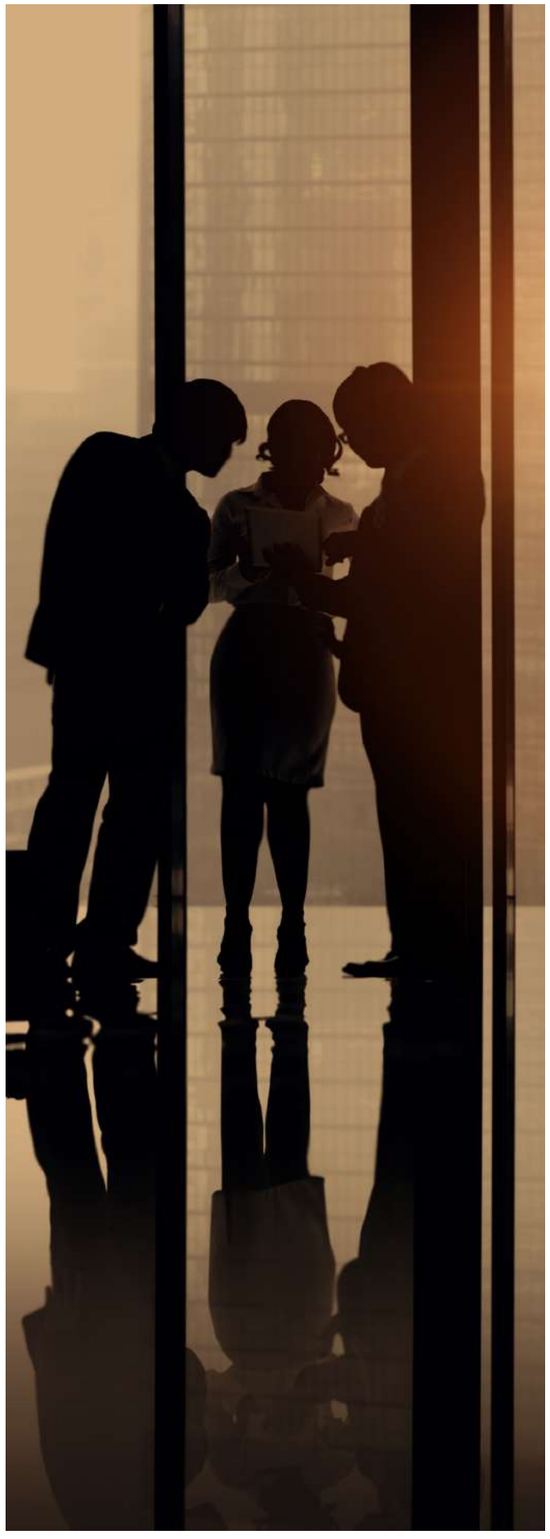

Automotive, Mobile & Embedded

44CON - Charging Ahead: Exploiting an EV Charger Controller at Pwn2Own Automotive 2024
McCaulay Hudson and Alex Plaskett

10th March 2025

McCaulay and Alex discussed their methodology, attack surface analysis, and demonstrated tooling they created to accelerate vulnerability discovery in firmware, as well as how this tooling was applied to a specific EV charger controller. EV chargers overall have a wide range of features and extensive connectivity, leading to significant attack surfaces. They described their journey from having zero knowledge of the specific target (Phoenix Contact CHARX SEC-3100) to remotely compromising it, performing privilege escalation and additional attacks.

They explored the intricacies of this "build your own" charging component and how this enabled the deployment of EV charging infrastructure, and then discussed weaknesses related to uploading arbitrary file contents, state switching, and injection techniques, and how these were combined to build an exploit chain that was eligible for the Pwn2Own competition. The audience gained an understanding of how multiple seemingly low-risk vulnerabilities can be chained together to amplify their impact, ultimately leading to code execution on the charger controller and demonstrating real-world consequences.

The talk concluded with an overview of EV charger post-exploitation and outlined several threat scenarios and potential impacts that could occur if an attacker were to compromise these devices and maintain persistence. Multiple demonstrations were presented, including tooling and exploits used to obtain a shell on the device. To illustrate full device control, the presenters also demonstrated a light show running on the EV charger. Finally, they shared their thoughts on building a robust security architecture for EV charging deployments.

Insomnihack - Pioneering Zero Days at Pwn2Own Automotive 2024
McCaulay Hudson and Alex Plaskett

13th March 2025

This presentation detailed the process of attacking an in-vehicle infotainment (IVI) system from start to finish as part of the Pwn2Own Automotive 2024 competition. Beginning with a hardware attack, the presenters discussed the sequence of challenges they had to overcome to gain initial access and subsequently establish a debugging environment that provided visibility into the device internals. This included overcoming hidden developer menus (including a Japanese language barrier), in-depth eMMC in-circuit programming, and exploring multiple hardware attack approaches.

Using this level of access, the presenters identified multiple software vulnerabilities that were used during the competition to gain remote code execution and compromise the IVI system. After examining Pioneer's novel operating system modifications and custom applications, they identified an arbitrary file write and directory traversal vulnerability that was severely constrained. They then abused a PKCS#11 configuration mechanism and chained it with a denial-of-service condition to transform the issue into arbitrary code execution, allowing them to satisfy the constraints of the competition.

During the talk, the presenters disclosed these vulnerabilities in detail and explained how they were chained together to compromise the device. In conclusion, they showcased a real-world attack scenario by demonstrating spyware exfiltrating data from the infotainment system to track an individual's location, contacts, and call history.

Samsung Galaxy S24 Pwn2Own Ireland 2024
Ken Gannon

13th June 2025

The white paper discussed the details of the vulnerabilities which were used at Pwn2Own Ireland 2024 to successfully compromise the Samsung Galaxy S24 and were presented at OffensiveCon 2025 in the talk Chainspotting 2: The Unofficial Sequel to the 2018 Talk "Chainspotting". The full exploit chain consisted of five different issues across several different applications, resulting in the ability to install arbitrary APKs.

A Look at RTEMS Security
Catalin Visinescu, Mark Paruzel, and Jonathan Lindsay

18th August 2025

A time-boxed security assessment was conducted of the Real-Time Executive for Multiprocessor Systems (RTEMS) project. RTEMs is a real-time operating system geared towards small and medium-sized embedded systems and has been deployed in a wide array of environments ranging from satellites and rovers to nuclear reactors and physics experiments. It is a multi-threaded OS that runs within a single address space with no separation of kernel and user space.

Streamlining Global Automotive Cybersecurity Governance to Accelerate Innovation, Assurance, and Compliance
Josh Kolleda

30th April 2025

Global automotive organisations face challenges in managing cyber security across multiple brands and regions due to differing regulations, siloed teams, and duplicated assurance efforts.

Effective governance requires centralised leadership and accountability, aligning policies and processes across the enterprise to balance innovation with robust security engineering and verification, reduce rework, and increase efficiency. Early integration of cyber security into development – and reuse of assurance artefacts like threat analyses – can shorten delivery times, lower costs, and reduce post-launch issues.

Transforming governance also involves strategic organisational design, stakeholder engagement, and change management to align resources and culture across brands. Establishing a cross-brand cyber security task force, reengineered funding streams, and consistent metrics supports continuous improvement and compliance with global standards while fostering shared innovation and risk management.

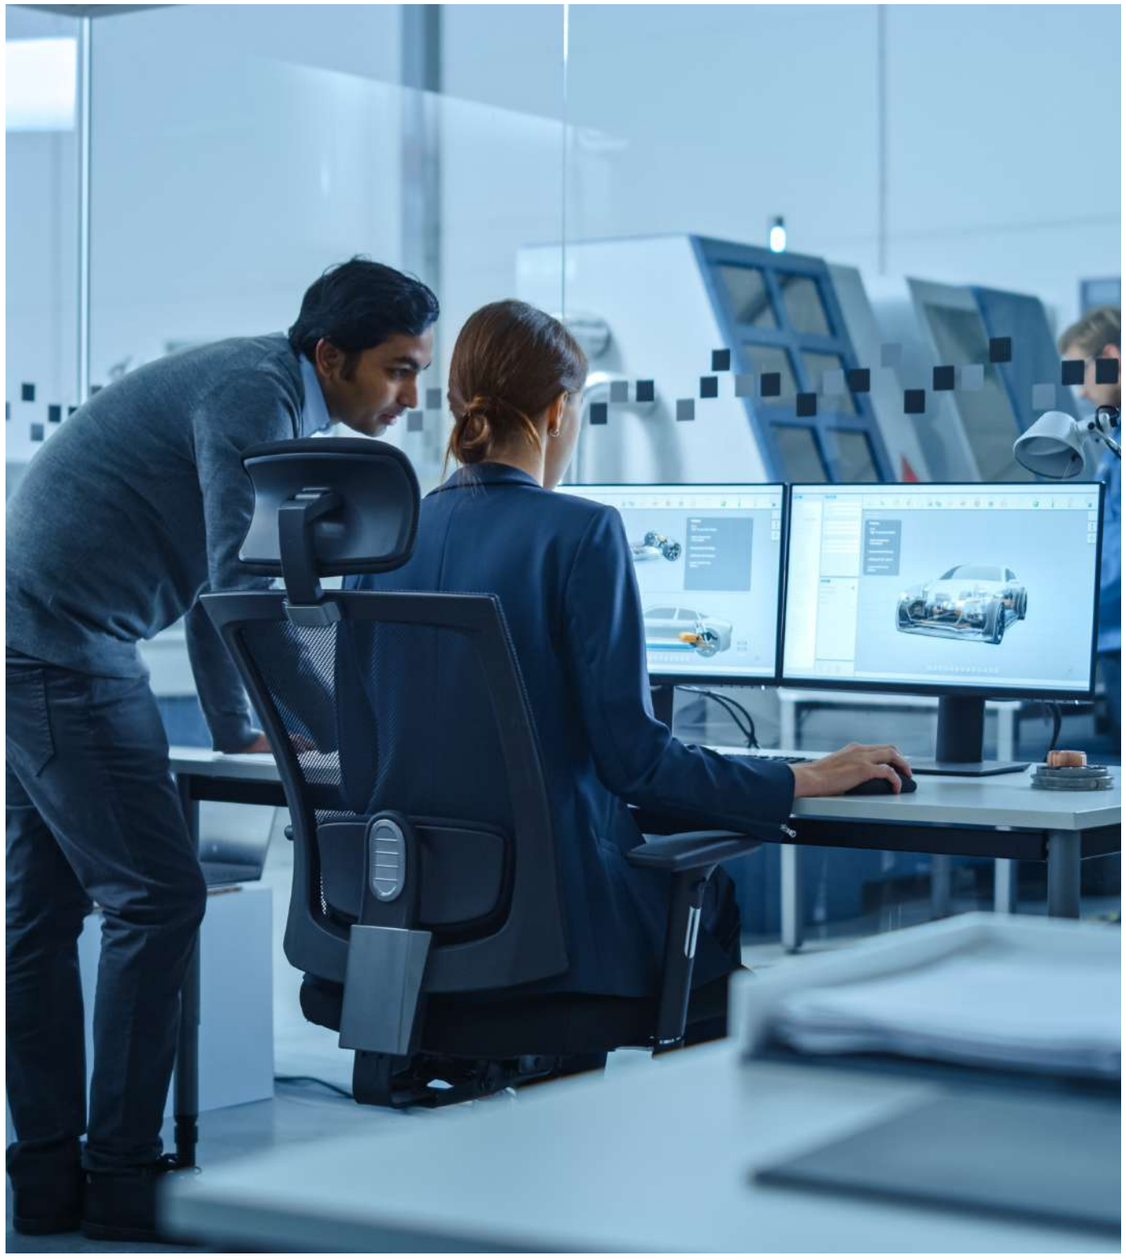

Exploit Development & Low-level Systems

VMware Workstation Guest-to-Host Escape Exploit Development
Alex Zaviyalov

30th September 2025

The research details the development of a guest-to-host escape exploit against VMware Workstation, demonstrating how a malicious actor operating inside a virtual machine (VM) can break out and execute code on the host system. The exploit chain combines an information disclosure with a stack-based buffer overflow in the virtualisation layer's device emulation, enabling memory corruption in the host process and controlled execution of a payload, such as a reverse shell.

This work highlights how vulnerabilities in hypervisor-level components can undermine the fundamental isolation between guest and host, elevating the risk of host compromise from a seemingly contained VM breach. The findings reinforce the importance of applying vendor patches promptly and hardening virtual devices to mitigate escape vectors in virtualised environments.

A Rendezvous with System Management Interrupts
Carles Pey

10th March 2025

System Management Interrupts (SMIs)-which switch CPUs into the highly privileged System Management Mode (SMM)-and found significant timing variability when measuring SMI execution on x86 UEFI platforms. The variability persisted even with identical code and parameters, and was traced largely to UEFI's SMI rendezvous mechanism, which synchronises all processor cores before servicing an SMI. Reducing the number of active cores reduced the timing differences for firmware security testing because it shows that SMI timing measurements can be influenced by normal platform behaviour rather than vulnerabilities. Accurate interpretation of SMM performance requires accounting for core synchronisation effects to avoid false positives when auditing firmware or investigating potential SMM anomalies.

Adventures in EM Side-channel Attacks
Aaron Kondziela

27th October 2025

The Hardware and Embedded Systems practice of NCC Group reproduced a sophisticated electromagnetic (EM) side-channel attack originally documented by NinjaLab against Infineon secure elements used in devices like YubiKey 5 Series tokens.

The original research showed that EM leakage could be exploited to extract secret ECDSA keys, potentially allowing hardware token cloning and complete compromise of the device's root of trust; patched firmware (5.7+) mitigates the issue. NCC Group successfully replicated this attack in a lab environment over two weeks using custom-fabricated equipment, demonstrating the practical challenges and techniques involved.

The work highlights that even well-designed cryptographic hardware can emit side-channel signals that leak sensitive information, underlining the importance of considering physical emanations – such as EM emissions – in security evaluations of hardware tokens and secure elements.

Governance, Policy & Strategy

Goal-Based Regulation
Liz James

6th October 2025

Goal-based regulation represents a shift from traditional checklist-style rules toward outcome-focused assurance, where regulators set high-level safety or security goals and expect organisations to demonstrate how those goals are met with reasoning and evidence. Rather than prescribing specific steps or technologies, this approach emphasises structured justification-organisations must explain their decisions, show evidence of effective risk management, and make their assurance cases defensible and traceable across audits. This reflects a broader trend in sectors such as maritime, automotive, aviation, rail, energy, and cyber-physical systems, where rigid checklists can stifle innovation and fail to capture the complexity of real-world systems.

Under goal-based regimes, compliance is not measured by ticking boxes but by evaluating the quality of evidence and the maturity of assurance practices. For example, in frameworks like UNECE R155 and principles-based assurance, regulators look beyond individual tests to understand how an organisation has framed, scoped, responded to, and learned from assessments as part of its overall management system. This method encourages continuous improvement, supports innovation, and helps regulation remain relevant in rapidly evolving technical landscapes by prioritising outcomes over specific mechanisms.

Bridging the Valley of Death
Liz James

4th November 2025

Promising innovations often fail between proof of concept (PoC) and minimum viable product (MVP), not just because of engineering challenges but due to gaps in assurance - confidence that a product is secure, reliable, and responsibly managed. Stakeholders – including customers, regulators, and investors – require defensible evidence of security and quality, yet many teams rely on ad-hoc activities ("just do a pen test") rather than structured confidence-building artefacts.

The article advocates using assurance cases - explicit frameworks linking claims, supporting arguments, and evidence - to make security and reliability visible and defensible throughout development. Treating outputs like vulnerability findings as assurance fragments that feed into a broader case helps organisations move innovations across the "Valley of Death" with trusted, auditable assurance rather than tacit or fragmented risk assessments.

Legacy Technology in Transport: More Than "Old Tech"
Liz James

9th October 2025

In transport and other long-lived, safety-critical sectors, "legacy" technology means far more than simply outdated hardware or software; it reflects systems that are deeply embedded into infrastructure, certification regimes, and operational processes, often enduring for decades due to safety requirements, regulatory constraints, and complex dependencies. Unlike in traditional IT – where legacy typically denotes "unsupported and risky" and drives straightforward replacement – transport systems may remain essential long after they become difficult to maintain, integrate, or upgrade, making clear definitions of legacy critical to effective risk assessment.

The ambiguity of the term can obscure discussions about investment and risk, because age, vendor support, maintainability, and safety certification can all influence whether a system should be modernised or retained. Transport organisations must therefore develop tailored frameworks for evaluating when technology has reached fragility, considering not only technical support but also embedded dependencies, ecosystem decay, and lifecycle constraints that influence safety and continuity of operations.

An Engineer's View: Operational Technology
Liz James and Jeffrey Hall

26th March 2025

Operational Technology (OT) systems control physical processes and must be managed differently from IT due to safety and engineering constraints. The article proposes evaluating OT along a Cyber-Physical Axis and Impact of Failure, helping organisations classify systems and prioritise security investments based on how integrated and critical they are. It emphasises that OT failures often cause partial degradation rather than complete outages, meaning risk assessments must account for nuanced operational impacts. This structured view supports better governance, resilience, and alignment between engineering and cyber security.

Pentesting V. Red Teaming V. Bug Bounty
Minali Arora

28th July 2025

The author explains the key differences between penetration testing, red teaming, and bug bounty programs, highlighting that each method serves a distinct purpose in an organisation's security strategy.

Penetration testing focuses on identifying and validating vulnerabilities within a defined scope, typically using white-box or grey-box approaches to provide actionable, proof-of-concept-backed findings. Red teaming simulates realistic adversary behaviour to assess an organisation's ability to detect, respond to, and recover from attacks, taking a broader and more holistic view of overall security posture. Bug bounty programs, by contrast, leverage the global security researcher community through continuous, incentivised testing to uncover issues that may fall outside traditional assessments. The post emphasises that these methods complement one another, and mature organisations often benefit from using all three in combination to build layered, resilient defences.

Quantum Data Centre of the Future
Jon Renshaw

26th June 2025

The article outlines NCC Group's role in the Quantum Data Centre of the Future (QDCF) project, an Innovate UK funded collaboration running from 2022 to 2025, that explores how quantum technologies can be securely integrated into classical data-centre environments. The initiative delivered a blueprint for hybrid quantum/classical architectures, developed quantum-ready computing modules, and designed quantum-secure communication mechanisms, including Quantum Key Distribution (QKD) and Post-Quantum Cryptography. NCC Group worked with partners such as ORCA Computing and BT to build threat models for quantum processors, quantum communications, and entropy-as-a-service components, highlighting the security implications of deploying quantum systems at scale. The project culminated in a demonstration at the UK National Quantum Computing Centre, showcasing use cases ranging from optimisation and simulation to entanglement-based security, and the blog summarises the associated risks and recommended security controls for organisations developing quantum-enabled data-centre capabilities.

Unmasking Techno Sophists
Andrew Havard

7th October 2025

Common signs of LLM-generated content to help readers distinguish AI-authored text from human writing, drawing a parallel between modern automated writing and ancient sophistry – persuasive, yet potentially shallow rhetoric lacking deep reasoning. It outlines observable patterns in punctuation, phrasing, and templated language that tend to appear in output from large language models, illustrating how such content can be recognised and critically evaluated rather than accepted at face value.

The Symbols of Operation
Chris Anley

14th October 2025

The article revisits an 1843 observation by Countess Ada Lovelace about the inherent confusion between code ("symbols of operation") and data, showing how this

The article revisits an 1843 observation by Countess Ada Lovelace about the inherent confusion between code ("symbols of operation") and data, showing how this blurred distinction has been a persistent root cause of major security issues in computing. Historical examples such as self-modifying code, buffer overflows, and return-oriented programming (ROP) illustrate how treating data as executable instructions enables classic exploitation patterns, while modern Agentic AI vulnerabilities similarly arise when systems misinterpret data as operational directives.

By framing security challenges through the lens of this fundamental "code vs. data" ambiguity, the discussion highlights the deep connection between theoretical computation concepts and practical vulnerabilities. Understanding and respecting the boundary between executable logic and untrusted input remains critical to preventing arbitrary code execution and similar classes of attacks across both traditional software and emerging AI-driven systems.

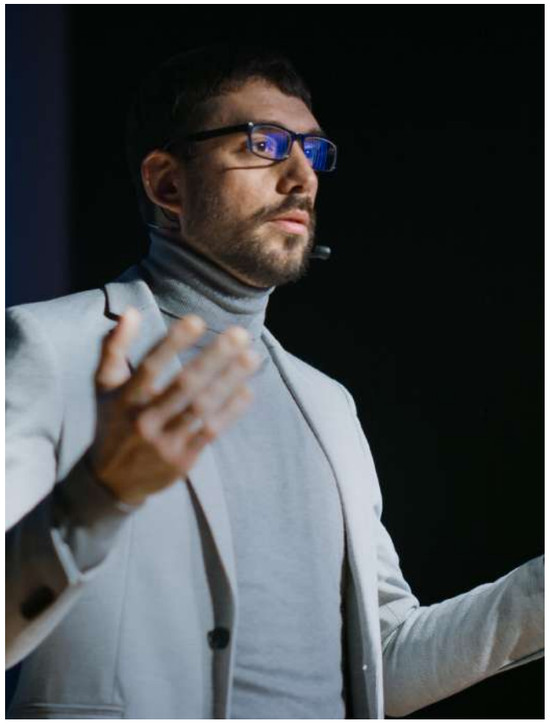

Technical Advisories & Vulnerability Research

Cross-Site Scripting in Umbraco Rich Text Display
Liyun Li

12th February 2025

This reports a stored cross-site scripting (XSS) vulnerability affecting Umbraco CMS versions 14.3.1 and below, caused by a lack of server-side input sanitisation in the rich text rendering logic. The issue arises because, although the TinyMCE rich text editor sanitises input on the client side, the backend API does not enforce equivalent filtering, allowing an authenticated user to inject arbitrary HTML or scripts directly via HTTP requests to the document management endpoints. When this malicious content is published and viewed, the injected script executes in visitors' browsers, demonstrating typical stored XSS behaviour.

An attacker who can modify or publish rich text content in a vulnerable Umbraco instance could leverage this vulnerability to execute scripts in the context of another user's session, potentially harvesting credentials, performing actions with the user's privileges, or compromising sensitive data. Umbraco's decision not to include server-side sanitization by default means administrators must implement their own mitigation using interfaces like IHtmlSanitizer or choose alternative rendering components. The advisory includes recommendations for organisations relying on rich text features to adopt appropriate sanitization frameworks to mitigate this class of attack.

Hash Denial-of-Service Attack in Multiple QUIC Implementations
Paul Bottinelli

19th February 2025

This identifies a class of vulnerabilities in several open-source QUIC protocol libraries caused by weak hash table implementations used to track QUIC connection identifiers (SCIDs). Researchers first confirmed the issue in xquic and then reviewed other QUIC libraries, finding the vulnerability in five maintained projects (including kwik, lsquic, picoquic, xquic, and Apache Traffic Server's QUIC support), as well as in additional experimental and unmaintained implementations. The weakness stems from using simple hash functions that allow an attacker to generate many colliding SCIDs, causing hash table operations to degrade from expected constant time to much slower performance, significantly slowing or stalling a server under attack.

The advisory explains how Hash Denial-of-Service (DoS) attacks exploit algorithmic complexity by forcing a server to spend disproportionate resources handling colliding entries, with experimental results showing up to ~300× slowdown under malicious connection attempts.

Because QUIC is a foundational protocol for modern web traffic (e.g., HTTP/3), the presence of Hash DoS issues in widely used implementations presents a serious availability risk if exploited. The findings underline the importance of using secure hash functions and robust data structures in protocol implementations to prevent attackers from degrading service or causing outages through relatively low-effort attacks.

Multiple Vulnerabilities in TCPDF
Neil Bergman

14th March 2025

TCPDF versions prior to 6.8.0 contain multiple critical vulnerabilities, including Local File Inclusion (LFI) via untrusted SVG input and PHP code injection through crafted font files. Other issues include weak SSL certificate validation and type-juggling flaws. These vulnerabilities can lead to remote code execution in applications that generate PDFs from user-controlled content. Users are advised to upgrade to TCPDF 6.8.0 to mitigate the risks.

Condeon CMS
Juan Marco Sanchez

28th July 2025

Multiple vulnerabilities in Condeon CMS (versions ≤ 1.9.1) allow attackers to compromise hosted sites by exposing sensitive server memory and abusing access controls. A publicly accessible memory dump file leaks critical information – such as file paths, SQL statements, password hashes, and valid session tokens-enabling attackers to hijack authenticated sessions by reusing leaked cookies. Additionally, a mass-assignment flaw lets an authenticated user manipulate form parameters to change a user's CustomerID and gain administrative access to other customers' sites.

Espressif Systems - ESP32 BluFi Reference Application Vulnerabilities
James Chambers

22nd August 2025

Several critical flaws were identified in the BluFi reference application used with Espressif's ESP32 IoT platform, affecting ESP-IDF versions such as 5.0.7, 5.1.5, 5.2.3, and 5.3.1. These include memory corruption vulnerabilities and cryptographic weaknesses in how Wi-Fi credentials and key exchange are handled over Bluetooth, which could be exploited to achieve arbitrary code execution on an ESP32 device or expose sensitive data like Wi-Fi network credentials over the wireless channel.

Tesla Telematics Control Unit - ADB Auth Bypass
Alex Plaskett and McCaulay Hudson

29th September 2025

The advisory details a vulnerability affecting Tesla's Telematics Control Unit (TCU) firmware prior to version 2025.14, where the Android Debug Bridge (adbd) lockdown mechanism failed to fully restrict privileged functionality. Although Tesla attempted to block interactive shell access via the TCU's exposed Micro-USB port, researchers found that adbd still permitted adb push, adb pull, and adb forward operations - each running with root privileges.

By exploiting this incomplete lockdown, an attacker with physical access could write arbitrary files to the device and overwrite kernel parameters such as uevent_helper or hotplug settings, causing their payload to execute with full root access. This allowed complete compromise of the TCU, tested and demonstrated on Tesla firmware v12 (2025.2.6), leading Tesla to address the issue in patch release 2025.14.

Public Security Assessments

Public Report - VeChainThor Galactica Security Assessment
Kevin Henry and Parnian Alimi

8th May 2025

In April 2025, VeChain engaged NCC Group's Cryptography Services practice to perform an implementation review of several updates culminating in the "Galactica" fork of the VeChain Thor blockchain. VeChainThor is an EVM-compatible layer 1 blockchain which caters to enterprise users, and is intended to serve as a foundation for a sustainable and scalable enterprise blockchain ecosystem.

Public Report: VetKeys Cryptography Review
Paul Bottinelli, Marie-Sarah Lacharité, and Javed Samuel

7th October 2025

During the spring of 2025, NCC Group was engaged to perform a cryptography review of DFINITY USA Research's verifiably encrypted threshold key derivation (vetKD) protocol implementation. The use of vetKD in the Internet Computer allows users to securely and privately derive keys (vetKeys) tied to their identity. The protocol itself uses nodes' BLS signature threshold key shares (vetKD master key shares) to sign (a hash of) the user's identity, then ElGamal-encrypt this signature with an ephemeral transport key generated by the user.

Public Report - Google Confidential Space Security Assessment
Carles Pey, Viktor Gazdag, and Daniele Costa

28th July 2025

During the spring of 2025, Google engaged NCC Group to conduct the security assessment of Confidential Space. Confidential Space is a Cloud-based system designed to provide isolated execution environments for sensitive workloads. In this latest revision, Confidential Space has been integrated with Intel TDX confidential computing technology, Intel Tiber Trust Authority (an independent attestation verifier service), and AWS Identity and Access Management (IAM).

Public Report: Meta Whatsapp Message Summarization Service
Elena Bakos Lang, Paul Bottinelli, David Brauchler, Marie-Sarah Lacharité, Carles Pey, Thomas Pornin, Javed Samuel, and Eric Schorn

27th August 2025

In late January 2025, Meta engaged several of NCC Group's specialty practices to conduct a security and privacy assessment of the WhatsApp Message Summarization Service, which is part of a broader Private Processing system.

This service allows WhatsApp users to send a batch of messages to a Meta-operated Large Language Model (LLM), which returns a summarization of the message contents. Several third parties play key roles in the service: Cloudflare maintains transparency logs of signed artifacts, such as Trusted Execution Environment (TEE) images and hashes of LLM prompts and quickly acts as an Oblivious Relay between WhatsApp users and Meta. The overall system is intended to ensure user data remains private, inaccessible to Meta, not persisted, and not used for any other purposes.

Public Report: Google Private AI Compute Review
Evan Anderson, Elena Bakos Lang, David Brauchler, Thomas Cannon, Gérald Doussot, Aaron Kondziela, Giacomo Pope, Thomas Pornin, Domen Puncer Kugler, Javed Samuel and Eric Schorn

28th July 2025

Starting in the spring of 2025, Google engaged NCC Group to conduct a series of reviews involving selected aspects of their Private AI Compute in the cloud system. The purpose of the Private AI Compute system is to extend the AI capabilities of a mobile device with more powerful cloud computing resources while aiming for the user data to still have the same privacy guarantees as with local-only computations.

Public Report: AWS EKS Security Claims
Divya Natesan and Jonathan Leadbeater

12th November 2025

In the second quarter of 2025, Amazon Web Services (AWS) engaged NCC Group to conduct an architecture review of the AWS-managed Kubernetes service: Amazon Elastic Kubernetes Service (Amazon EKS), to evaluate how the system is designed in order to prevent AWS employees (Operators) from accessing Customer Content stored or processed by Amazon EKS, with a specific focus on validating a number of Data Security claims asserted by AWS.

Collaboration with Industry & Academia

As a member of the Coalition for Content Provenance and Authenticity (C2PA), we have continued to support the development of security standards as a part of the conformance task force. Such standards are critical to support the development of a trusted ecosystem, and huge progress has been made with the publication of claim generator product security requirements, which defines the security levels a product must demonstrate to achieve certification, and the certificate policy, which defines the requirements for Certificate Authorities (CA). But the real measure of success is adoption – and increasing numbers of software, devices and CAs certified through the conformance programme since its official launch in summer, point to a real desire and momentum to improve the provenance of media.

NCC Group continues to build relationships with academia to support the latest cyber security research and develop the workforce of the future. This year we supported the Centre for Doctoral Training (CDT) in Cyber Secure Everywhere by setting research challenges and hosting their current student cohort in our offices in Cheltenham, where they had presentations from NCC Group consultants on topics as diverse as the security of AI, malware analysis, mesh radio networking, incident response and privilege escalation. We are looking forward to continuing to build this relationship and others in 2026.

Tools

The following tools were uploaded or updated in 2025:

libslub is a python library to examine the SLUB managements structures and object allocations (the Linux kernel heap implementation). It is currently designed for use with GDB but could easily be adapted to work with other debuggers. It helps with understanding SLUB internals and developing Linux kernel exploits.

https://github.com/nccgroup/libslub

A proof-of-concept framework and example code that demonstrates how to write and build FreeBSD kernel modules using the Rust programming language instead of C

https://github.com/nccgroup/freebsd-kernel-module-rust

Blackbox Protobuf is a set of tools for working with encoded Protocol Buffers (protobuf) without the matching protobuf definition.

https://github.com/nccgroup/blackboxprotobuf

A TypeScript port of kramdown, the fast, pure-Ruby Markdown-superset converter. This implementation is designed to maintain bug-for-bug compatibility with the original kramdown parser. As a port, kramdown-ts was written referencing the original kramdown code, meaning effectively all of the abstractions available within kramdown are also available to kramdown-ts users.

https://github.com/nccgroup/kramdown-ts

A collection of tools for reverse engineering and interacting with the PowerG radio protocol. This release accompanies the research James Chambers and Sultan Qasim Khan presented at REcon and hardwear.io.

https://github.com/nccgroup/powerg-tools

Scout Suite is an open source multi-cloud security-auditing tool, which enables security posture assessment of cloud environments. Using the APIs exposed by cloud providers, Scout Suite gathers configuration data for manual inspection and highlights risk areas. Rather than going through dozens of pages on the web consoles, Scout Suite presents a clear view of the attack surface automatically.

https://github.com/nccgroup/ScoutSuite

Sniffle is a sniffer for Bluetooth 5 and 4.x (LE) using TI CC1352/CC26x2 hardware.

https://github.com/nccgroup/Sniffle

PhanTap is an 'invisible' network tap aimed at red teams. With limited physical access to a target building, this tap can be installed in line between a network device and the corporate network. PhanTap is silent in the network and does not affect the victim's traffic, even in networks having NAC (Network Access Control 802.1X - 2004).

https://github.com/nccgroup/phantap

A nanoMIPS module for Ghidra.

https://github.com/nccgroup/ghidra-nanomips

This project implements an HTTP server that acts as a bridge between HTTP/1.1 requests and a remote MCP server, using the mcp python library (GitHub). The main purpose of this initiative is to be able to use HTTP security tools to test remote MCP servers using the remote transport mechanisms (HTTP+SSE or Streamable HTTP).

https://github.com/nccgroup/http-mcp-bridge

Scripts for recovering string definitions in Go binaries with P-Code analysis. Tested with x86, x86-64, ARM, and ARM64.

https://github.com/nccgroup/ghostrings

Chrome release 142 added Local Network Access (LNA), which prevents DNS rebinding attack techniques implemented in Singularity. Singularity of Origin's LNA-from-Non-Secure-Contexts branch implements support for Trial for Local Network Access from Non-Secure Contexts. This temporarily allows access to resources on local networks from non-secure contexts, thus letting you experiment with DNS rebinding attacks a bit longer (until May 18, 2026) when using Chrome.

https://github.com/nccgroup/singularity

MediaTek BP firmware tools

https://github.com/nccgroup/mtk_bp

A hub for Red Team activity to aid in record keeping, situational awareness and reporting. Stepping Stones provides a web-based UI for the team to log activity and generate report snippets. The UI is intended to be rapid enough to be used throughout the engagement, not just in the reporting phase.

https://github.com/nccgroup/SteppingStones

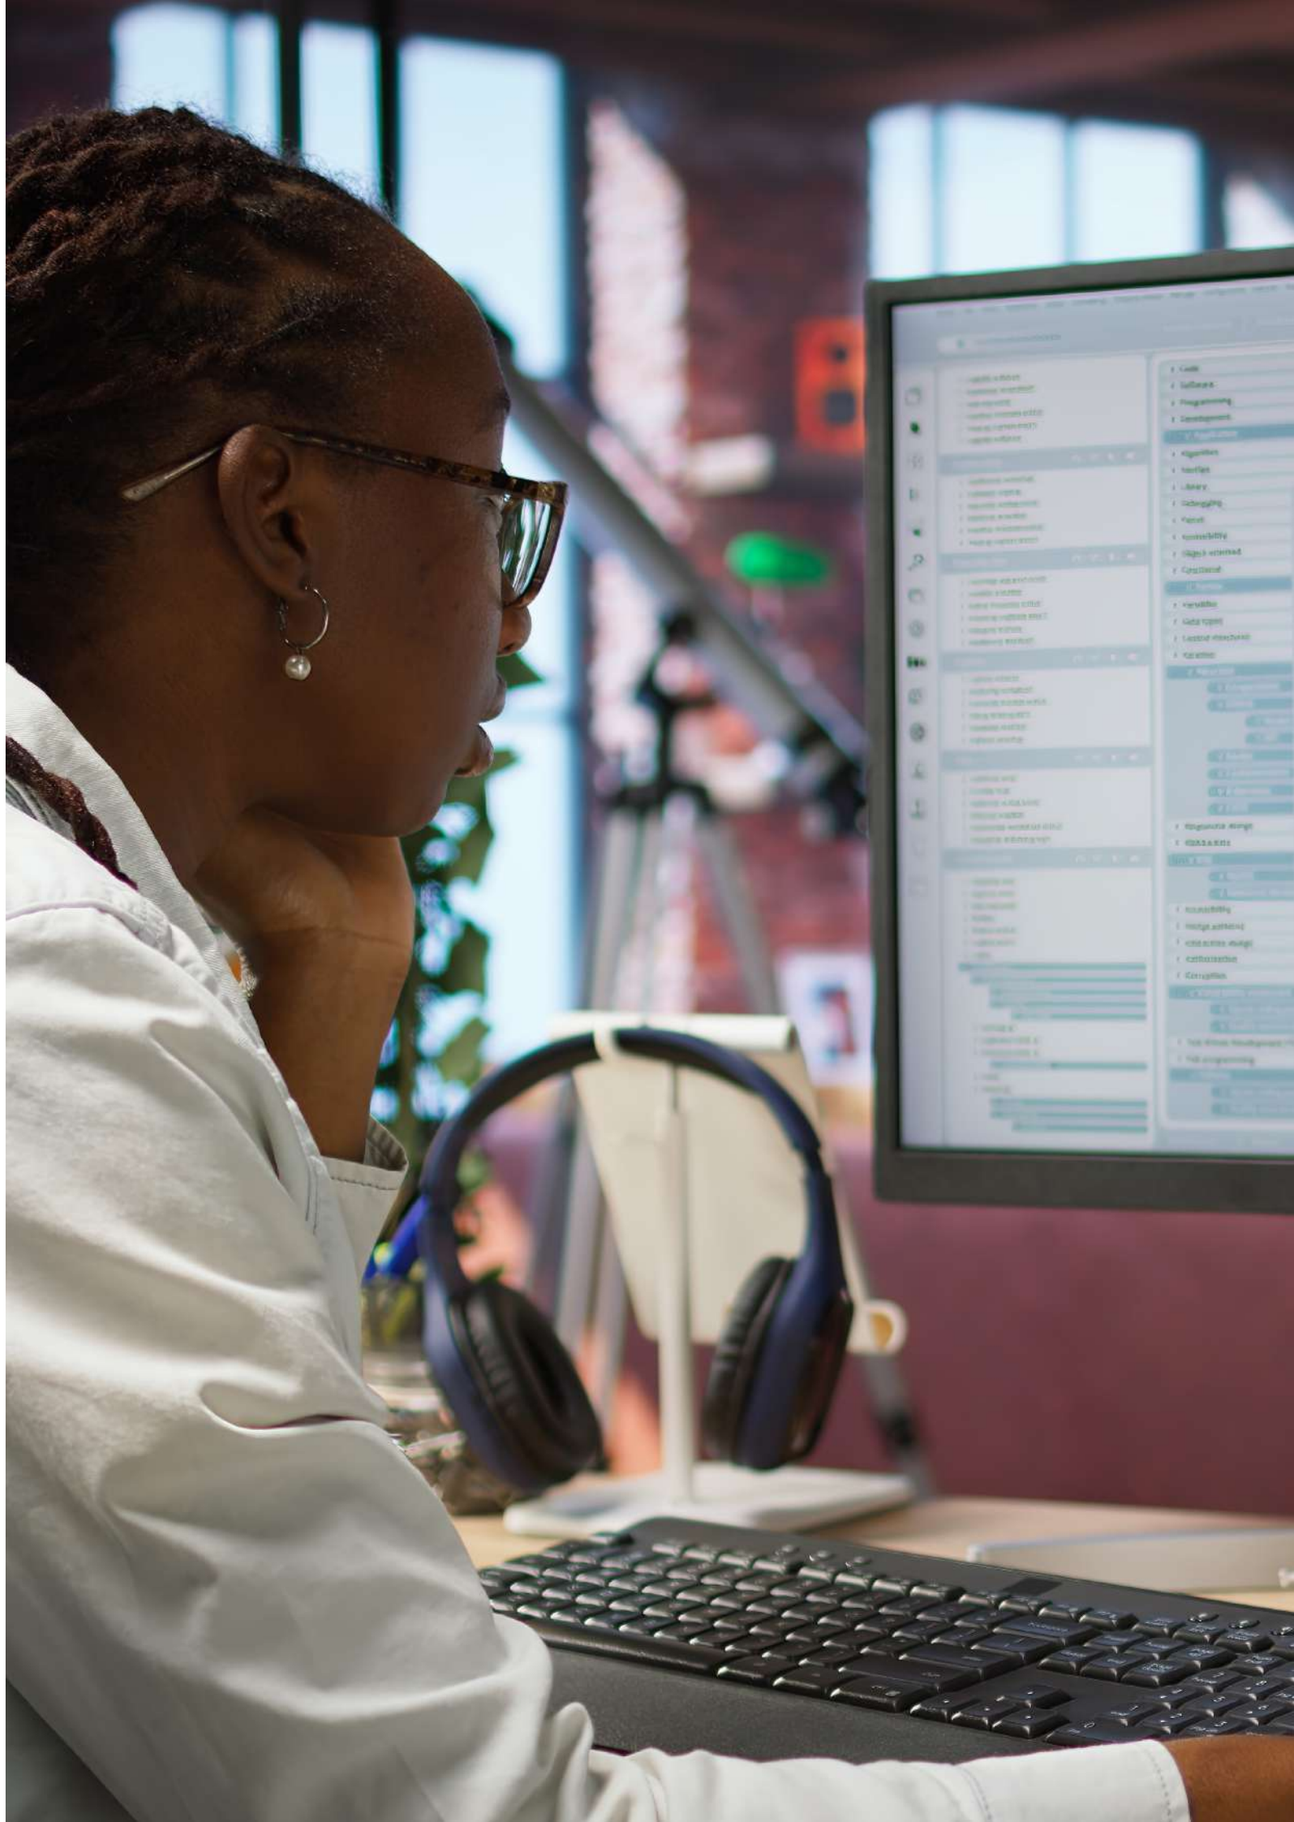

Authors & Contributors

Andy Davis
Global Research Director

Andy Davis is the Global Research Director. With over 30 years' experience across UK Government, consultancy and professional services, he leads NCC Group's global research programme – setting strategy, ensuring technical excellence, and governing research output. Previously Chief Research Officer at IRM and Head of Security Research at KPMG, Andy joined NCC Group in 2010 and later built its global Transport Practice. He is known for turning complex cyber risk into clear, actionable guidance for engineers and boards.

Javed Samuel
Practice Director,
Cryptography Services

Javed Samuel leads NCC Group's specialised Cryptography Services Division. Javed and his team specialise in novel cryptographic implementation and design assessments across a range of areas such as open-source cryptography projects, post-quantum cryptography, blockchain platforms, embedded devices, authentication mechanisms, encryption tools and custom protocol reviews. He devotes time to cryptography research across multiple areas and regularly presents at various security conferences. Prior to NCC Group, Javed worked as a developer with Oracle's Database Security group and obtained an MEng and BSc in Computer Science MIT and an MSc in Applied and Computational Mathematics at Oxford University as a Rhodes scholar.

Catalin Visinescu
Technical Director

Catalin Visinescu is a Technical Director at NCC Group, based in Waterloo, Ontario. He specialises in hardware and embedded systems security, including bootloaders, secure execution environments, Android platform hardening, microcontroller firmware and secure device manufacturing.

With a background in C/C++ and deep code-review expertise across RTOS, Linux kernels and TrustZone, he previously spent nine years at BlackBerry leading the Security Firmware team, and six years at Alcatel/Thales Transportation on secure over-the-air systems. He holds computer science degrees from the Polytechnic University of Bucharest and the University of Waterloo.

Jon Renshaw
Director of Security Services

Jon has consulted for 20 years on the design and implementation of cyber secure communication systems across the public sector, defence, transport and manufacturing industries. He began his career in cyber security, after completing a bachelor's degree in electrical and electronic engineering at Cardiff University. He subsequently took on a senior engineering role at a network security consultancy startup, where he managed the company's innovation processes, before moving on to setup his own consultancy business.

Jon joined NCC Group in 2022, and in his current role as Director Security Research Services, he is responsible for delivering fundamental and applied cyber security research in collaboration with, and on behalf of, NCC Group's customers and partners.

Jon is a Chartered Engineer (CEng) with the UK Institute of Engineering and Technology (IET) and has achieved qualifications in network and enterprise architecture, information security risk management, cloud technologies, leadership and management.

Alex Plaskett
Associate Director of
Exploit Development

Alex Plaskett is an Associate Director in NCC Group's Exploit Development Group (EDG). With more than 18 years in vulnerability research and exploitation, he has won Pwn2Own five times across desktop, mobile, embedded and automotive targets. Alex frequently speaks at conferences such as Black Hat, OffensiveCon, Hexacon, HITB, BlueHat, POC and Troopers. He has previously led security teams in fintech, mobile security and research, and his work has helped drive timely vendor patching across high-profile products.

Liz James
Managing Security Consultant

Dr Liz James is a Managing Security Consultant and Research & Innovation Lead within NCC Group's Transport practice. Her focus spans cyber-physical and intelligent mobility systems, bringing applied research and assurance to automotive, rail and wider OT environments. She has served as NCC's principal investigator in a HORIZON Europe consortium on edge/IoT security and conducted EVSE security testing at the 2024 Pwn2Own Automotive challenge. A former Vice-Chair for Intelligent Mobility & Transport at techUK, Liz combines deep technical rigour with cross-industry collaboration to raise security standards across next-generation transport.

Chris Anley
Chief Scientist, NCC Group

Chris Anley is Chief Scientist at NCC Group. Since 1996, he has delivered thousands of security assessments – penetration tests, and code and design reviews – across diverse platforms and architectures for many of the world's largest organisations. As Chief Scientist, he guides NCC research programs and advises customers and partners on security issues. Alongside this, he conducts independent research into new and emerging threats and the technologies shaping them.

Peter Bishop
Chief Scientist, Adelard

Peter Bishop is a Chief Scientist at Adelard (part of NCC Group pls) and a professor at City St Georges, University of London. He has extensive experience in the assurance of critical systems in a wide range of industries, including nuclear, defence, air traffic control, rail, automotive, finance and medical devices. He has also been involved in the development of standards and guidance for different industry sectors and undertaken security and safety research for government bodies including the UK Civil Aviation Authority, UK National Cyber Security Centre (NCSC), and both the US and UK nuclear regulators.

Peter has contributed to national and international research projects on software fault tolerance, software diversity, testing, static analysis, formal methods and safety case design methodology, and has authored or co-authored over 80 different publications in journals, conference proceedings and books.

Ristin Rivera
Research Manager

Ristin Rivera is Research Manager at NCC Group. They help build and design research projects and tools to help fill the gaps in security knowledge and capabilities.

Acknowledgements

We would like to sincerely thank the researchers, engineers, and collaborators whose dedication and ingenuity defined our work this year. Your technical excellence, curiosity, and commitment to advancing cyber security have been central to the progress and impact of our research.

We would also like to recognise the outstanding efforts of our new AI Services team, whose work this year has helped shape our understanding of the rapidly evolving AI security landscape. By exploring emerging risks, developing practical methodologies, and contributing clear, thoughtful research in a fast-moving space, the team has played a key role in ensuring our research remains relevant, responsible, and forward-looking.

Our gratitude also extends to our research leads, whose leadership is critical to the success of our programme. Through mentorship, guidance, and active engagement, they help bring new researchers into the community, foster collaboration, and support individuals in shaping impactful research. Their ability to balance technical depth with encouragement and direction ensures that strong ideas are developed into meaningful contributions.

We are especially grateful to our industry and academic partners for their continued collaboration, thoughtful feedback, and shared expertise, all of which have strengthened the quality and reach of our outputs. We also thank our clients and stakeholders for their trust and ongoing engagement, which motivates us to pursue rigorous, high-impact research and practical solutions.

Lastly, we acknowledge the wider cyber security community for its openness, collaboration, and willingness to share knowledge. Through this collective effort, we continue to strengthen the security of the digital ecosystem. Thank you for your contributions and for being part of this ongoing journey.

About Research at NCC Group

“NCC Group’s research in areas like ML and threat intelligence shapes its resilience and offensive security, helping clients anticipate new threat tactics and techniques.”

- The Forrester Wave™ Cybersecurity Consulting Services in Europe, Q4 2025

NCC Group employs some of the most talented security consultants and researchers on the planet, serving 15,000 clients worldwide and uncovering countless vulnerabilities per year through both client work and independent vulnerability research. With hundreds of specialised consultants, our technical security research areas extend into almost every area of security, as well as global standards bodies, including CIS Benchmarks. We perform offensive and defensive research across a vast range of targets, including blackbox and whitebox testing of previously unanalysed emerging technologies and computational architectures.

We publish research in a variety of subfields including applied cryptography, hardware and embedded systems, secure coding and programming languages, browser and client-side security, cyber-physical systems, operating systems and their internals, mobile security and privacy, application security, privacy enhancing technologies, distributed systems, network and protocol security, cloud, containerization, and virtualisation, and both offensive attacks on – and defensive uses of – machine learning and AI systems.

You can find samples of some of our recent public-facing work, including blog posts, whitepapers, conference talk listings, and technical advisories on our Research site, alongside our technical Twitter (X) account and our public GitHub, which hosts over 300 open-source tools and datasets authored by NCC Group researchers. We also have deep academic research partnerships with several leading universities, as evidenced across several of our research publications. NCC Group also regularly conducts publicly reported security audits across a range of high-impact and security-critical technologies.

Our technical capabilities extend beyond our public-facing work to include our internal-only groups and resources, including our world-class Exploit Development Group (EDG), Threat Intelligence Team and Full Spectrum Attack Simulation (FSAS) group, as well as several technical specialty practices and hundreds of pieces of unpublished proprietary tooling.

Our research program delivers thousands of research days annually, by researchers at all levels from across our global business. We support our researchers through a full-time technical research leadership team, mentorship and coaching, incentives and awards, and collaboration within and across several internal research groups.
We regularly present our work in top research venues including Black Hat USA, Shmoocon, Hardwear.io, REcon, Appsec USA, Toorcon, BSidesLV, Chaos Communication Congress, Microsoft BlueHat, HITB Amsterdam, RSA Conference, CanSecWest, OffensiveCon, DEF CON, and countless others.

Our research is regularly covered by publications including Wired, Forbes, The New York Times, Politico, DarkReading, TechCrunch, Fast Company, the Wall Street Journal, The Register, SC Magazine, and The Hacker News, Bleeping Computer, Trend Micro, Security Week, alongside other mainstream and trade publications globally.

research@nccgroup.com
www.nccgroup.com/research
@nccgroupinfosec

Cyber attack?

Call our 24/7 Hack Hotline now.

UK & Europe
+44 331 630 0690

U.S. & Canada
(855) 684-1212

Australia
1800 975 310

Singapore
+61 2 8379 7870

Fox-IT - Benelux
0800 369 23 78 (NL)
+32 2304 22 16 (BE)
+31 (0) 88 369 23 78 (Int)

We are here for you

Contact us today to learn more about global cyber security regulations.

UK & Europe
+44 (0) 161 209 5200

U.S. & Canada
+1 (800) 813 3523

Australia
+61 (0) 2 9552 4451

Singapore
+65 6800 0950

Fox-IT - Benelux
+31 (0)85 799 0680 (NL)
+32 2304 22 19 (BE)

Philippines
+63 2 8540 9450