2021 NCC Group Annual Research Report

Written by Jennifer Fernick

SVP & Global Head of Research, NCC Group

research.nccgroup.com

Table of Contents

Executive Summary: Research at NCC Group (2021) 3

Message from the Global Head of Research 4

Open Source Security Tool & Code Releases 6

Publicly Reported Security Audits 8

Applied Cryptography 9

Attacking (& Attacking with) Machine Learning Systems 12

Misinformation, Deepfakes, & Synthetic Media 14

Reducing Vulnerabilities at Scale & Improving Open Source Security 15

Virtualization, Emulation, & Containerization 19

Hardware & Embedded Systems 21

5G Security & Smart Environments 22

Public Interest Technology 24

Cloud & CI/CD Pipeline Security 27

RIFT, Threat intelligence, CIRT, & Honeypots 28

Managed Detection & Response 31

Exploit Development Group 32

Other Research & Speaking 34

Acknowledgements 37

About Research at NCC Group 38

Contact Us 39

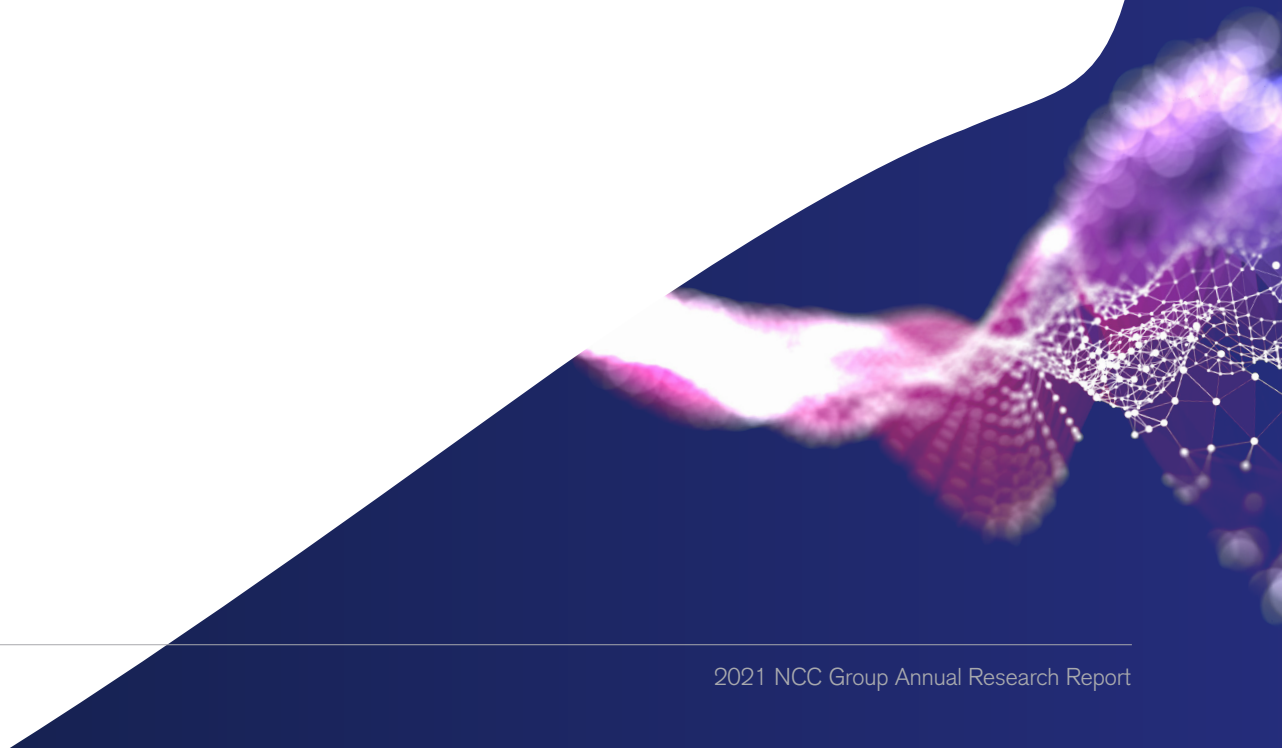

Executive Summary: Research at NCC Group (2021)

In 2021, NCC Group researchers hacked drones out of the sky; attacked real-world machine learning systems; ironically gained RCE against Metasploit; released user-centric mobile privacy analysis tooling; advised US Congressional staffers about open source and supply chain security; asked whether GPT-3 can generate exploits; reverse engineered Intel's proprietary binary translator which runs ARM binaries on x86 to write targeted malware that bypasses existing platform analysis for platforms used by hundreds of millions; shared our expertise in Kubernetes security to help improve recommendations made by the NSA and CISA; discovered new vulnerability classes, and found many critical vulnerabilities in high-impact systems.

We also delivered thousands of dedicated research days and over **237** research publications and conference presentations, pursuing research in areas including applied cryptography, hardware & embedded systems security, artificial intelligence, programming languages, operating systems, cloud & container security, exploit development, threat intelligence, vulnerability research against a range of high-impact targets, and beyond. We delivered **139** research papers, whitepapers, and technical blog posts and advisories, at least **30** new open source tools, as well as at least **68** conference presentations, in venues including Black Hat USA, Toorcon, the Linux Foundation Member Summit, Ekoparty, Black Hat Europe, Hardwear.io, the ACM/SIGAPP Symposium on Applied Computing, the Symposium on Usable Privacy and Security (SOUPS), an IETF Internet Architecture Board Workshop, POC 2021, the International Cryptographic Module Conference (ICMC), DEF CON, and many more.

This year, we also began to see our efforts to create 9 internal Research Working Groups begin to bear fruit, through dozens of publications, advisories, and conference presentations that we'll share throughout this report, as well as through increased global collaborative research across the firm, a stronger research community in which those new to research can explore ideas with experienced and world-class researchers in a radically open and inclusive environment, and a broadened and more interdisciplinary approach to security research. Our approach to research has been celebrated as a differentiator in the marketplace in the Forrester Wave™: European Cybersecurity Consulting Providers, Q3 2021 report, where it was noted that *"NCC dedicates a large proportion of staff time (up to 20%) for own research projects, culminating in a lot of specialist security research and the development and release of open-source tools, setting it apart on this dimension in a crowded field. NCC excels in its testing work and its research capabilities have made demonstrable improvements in security beyond its direct work on client projects."*

In 2021, we continued our long-running industry partnership with the Centre for Doctoral Training in Data Intensive Science at University College London (UCL), where we conducted collaborative research in artificial intelligence with UCL PhD students. Our research and standardization efforts included several contributions to the forthcoming version of the C programming language standard (C23), as well as contributions to the Center for Internet Security's CIS Benchmarks for Microsoft 365, and IETF Internet Drafts and presentations before the IETF Internet Architecture Board. We also served on a number of advisory boards including the Industrial Advisory Board at King's College London, the Governing Board and Technical Advisory Council of the Open Source Security Foundation, the Executive Steering Board for Internet of Things Security Foundation (IoTSF), the UK's National Cyber Security Centre (NCSC) Research Advisory Panel, among others. We have also had members appointed to the Science Advisory Council for the United Kingdom's Home Office.

From a Commercial Research perspective, 2021 saw us deliver thousands of person-days of paid research for clients. Besides our Public Reports, that research is **not** covered in this Annual Research Report, aside from its mention here. Those engagements included horizon scanning on future technologies and their impacts for our clients, and research of various defensive techniques to help our clients in their risk reduction endeavours. We were also part of a winning research consortium for UKRI funding (£11.6m) on the topic of the Quantum Data Centre of the Future. Over the next 3 years we will provide security advice and guidance to the project as it sets out to define a blueprint for future secure Quantum Data Centres.

NCC Group researchers Ross Bradley, Eva Esteban Molina, Philip Marsden, & Mark Tedman came in First Place at the global "5G Cyber Security Hack'' competition this summer, hosted by by the Finnish Transport and Communications Agency, Aalto University, Cisco, Ericsson, Nokia, and PwC. Researchers Dale Pavey and Guy Morley were named the winners of the Best Ethical Hacker/ Pentester Award at the Security Serious Unsung Heroes Awards for their work with independent UK consumer body Which?, investigating the safety and security of a range of IoT devices, exposing a number of unsafe devices. Jennifer Fernick, our Global Head of Research, was named one of Canada's Top 20 Women in Cybersecurity by IT World Canada. We also had the privilege to serve as Keynote Speakers for events like the Linux Foundation Member Summit, and the SANS Pentest Hack Fest, as well as serving on the program committees of influential security research venues including the USENIX Workshop on Offensive Technologies (WOOT), Toorcon, and Black Hat USA, among others.

Message from the Global Head of Research

As our world undergoes rapid social change and new norms for our relationship with information and with emerging technologies are set, I believe it is essential that as technologists, we connect with - and communicate about - our values, intentions, and the big questions that inspire us. It is my hope that through this report, readers can begin to see and interrogate themes that emerged across our wider security research program, reflective of the things that matter most to individuals across NCC Group.

Jennifer Fernick
SVP & Global Head of Research at NCC Group

From my own view, in the past year, we've seen a number of themes:

- **Time to exploitation of vulnerabilities in the wild has rapidly decreased.** Industry research aligns with what we're seeing in practice - time to exploitation of a vulnerability in the wild after a patch has been released has dropped from several weeks as of a few years ago, to only a few days now. This is challenging because it requires a seriously retooled approach to vulnerability management, and because it puts tremendous pressure to patch on key intermediaries in the software supply chain - many of whom are open-source maintainers.
- **Misinformation kills.** Much in the way that 2020 taught us that election security was hardly about the hardware at all, 2021 showed that the health of populations depends at least as much on a population's scientific and media literacy as it does on the core scientific advancements themselves, and that social media companies may be playing potentially even more of a role in the health and safety of citizens than those individuals' own democratically-elected leaders.
- **Software supply chain security matters.** It is no longer in the realm of the unusually paranoid to care not only about the security of our own code, but also about its upstream dependencies and its downstream deployment pipeline. Numerous high profile incidents in late 2020 through 2021 emphasized this point - so much so, that U.S. President Biden issued an Executive Order largely aimed at remediating it.
- **Ransomware remains an unsolved problem.** Despite it being a relatively uninteresting topic for most security researchers, ransomware is one of the greatest threats to digital security as perceived by both the general public and countless CISOs alike. Why then, is something this impactful being largely overlooked by the research community outside of threat intelligence teams? It seems like the best advice we have is for IT teams to patch everything, faster - but perfection is not a strategy. I would like to understand: In a world of limited resources, which interventions are the most effective at preventing ransomware infections? Sanctions against ransomware operators are likely the greatest tool we have, but that doesn't erase the question of what technical solutions to this problem look like.
- **Geopolitical power increasingly depends on information access, and can swing faster than ever before.** What does it mean when we put internet-enabled sensors into the things that keep us alive? What does it mean to put internet-enabled sensors into all of the things that make life worth living? Does our acceptance of that risk change when we imagine highly-skilled foreign adversaries a few keystrokes away?

In 2022 and beyond, we're going to face a number of challenges, new and old. These include:

- **Affective computing, brain-computer interfaces, and the broader question of where we draw the line as the physical distance between computers and human bodies grows increasingly small.** With the increasing adoption of smart cities, smart buildings, and other sensor-rich IoT-connected environments, avoiding surveillance in daily life and in physical space becomes increasingly difficult, even on one's own property. Compounded with the coming cultural normalization of wearable and even implantable sensors for gaming, entertainment, and the "metaverse" - not to mention the active commercial development of brain-computer interfaces - computers keep getting closer and closer to our bodies, behaviours, and emotions. I often worry that the last bastion of human freedom - one's own mind - is being eroded quickly and quietly in irreparable ways.

- **Full-stack security auditing and defense of real-world machine learning systems.** The ubiquitous deployment of machine learning systems represents a significant and poorly-understood real-world attack surface. This is problematic because these systems are too frequently granted undue and unchecked autonomy over safety-critical and socially-critical decisions, and also because machine learning models are code. Every enterprise AI system that as an industry we fail to secure can put the sensitive (and often irrevocable) data of millions or even billions of people at risk. Furthermore, insufficiently secure systems can open up both individuals and society to unjustified machine learning "decisions" that are often trusted implicitly - even if the machine that computed them is not trustworthy at all. Developing rigorous, full-stack threat models and testing frameworks is critical and urgent work.

- **The rapid advancement of large language models in artificial intelligence.** Large-scale natural language models such as GPT-3 and others have been scaling at 1-2 orders of magnitude per year, and the generalizable performance of transformer models seems to improve - even over domain-specific models - with increasing scale. While this may not necessarily mean that artificial general intelligence is imminent, we would be unwise to underestimate the impact that the weaponization of large natural language processing models trained on a corpus of the world's source code and vulnerability databases could have on software security and the safety of the internet. For example, we might ask, could large language models like GPT-3 or its' successors generate exploits of software vulnerabilities? Can low-code A.I. pair-programming tools create or propagate unsafe code? To answer these questions completely requires further research, but there are compelling reasons to believe that the answer is yes.

- **The relative monoculture of the public cloud's attack surface making the exploitation of singular vulnerabilities more socially and economically catastrophic.** The public cloud made incredible computing possible at scale - but with software infrastructure of such homogeneity, it also made vulnerability exploitation possible at scale, raising the intrinsic cost or risk of the existence of individual vulnerabilities in these high-value targets. The question then becomes: what does this mean for the value of vulnerabilities in the global "marketplace," and how might that change how vendors (and researchers) approach the security of these systems?

- **Decentralized finance ("DeFi") coupling code and money more tightly than ever before, forcing those who believe "code is law" to reckon with the fact that all software has flaws.** Large financial institutions each spend hundreds of millions - sometimes even billions - per year on cybersecurity, because they know that failing to do so may cost them even more. Are proportional investments being made in decentralized finance applications and cryptocurrencies? Not usually. While improvements in cyber resilience today are largely driven by regulatory interventions, in the sphere of decentralized finance (DeFi), where value is exclusively stored digitally and mediated directly by code, a threat actor could directly and immediately remove value from a company through attacking the underlying infrastructure, protocols, or cryptographic implementations. One leaked cryptographic key or a single software flaw could lead to the collapse of entire organizations.Time will tell whether we will see a bottom-up, market-driven push for higher assurance systems for serious decentralized finance companies, or whether they will become the highest-risk value stores on the planet.

- **As the COVID-19 pandemic rages on, how do we maintain privacy in a world in which we are accountable to one another?** In our research this year, we conducted wide-ranging research into the security & privacy aspects of various jurisdictions' vaccine passports, but other new forms of attestation and surveillance - including smart buildings, enhanced genomic and health testing, and mobile passports - remain under-studied by the security research community. Furthermore, the pandemic has breathed new life into large tech companies' VR and augmented reality ambitions ("the metaverse") which will affect not only the security and privacy of our devices and homes, but invites new kinds of sensors and interfaces (perhaps even brain-computer interfaces) which will be perhaps more intrinsically connected to our bodies, minds, and behaviour that anything we've seen before.

- Finally, we must ask ourselves, **what are the grand challenges for cybersecurity? What are the problems that matter the most to the security & privacy of individuals, organizations, and the internet**? As an industry, we face a reckoning in which I believe that we need to elevate ourselves toward taking a more "scientific" and rigorous approach to the study of information security cause and effect, and let go of the unspoken agreements, copycat risk-mitigations, hearsay "best practices," and other unacceptable industry norms. While the development of obscure and complex exploit chains will always have its place to enable us to understand the true frontiers of what we're up against, it is time to prioritize the problems that most meaningfully affect the world in which we live. It is my intent that we become more methodical in our study of the mitigations of specific cybersecurity threats, to move beyond a world where companies sometimes feel they have no choice but to sink money into the same expensive and at times incomplete or misconfigured security products as their peers, without a real understanding of their effectiveness, and instead to a world where they are empowered through an understanding of the actual costs of both attack and defense.

Open Source Security Tool & Code Releases

In 2021, we released around 30 open source security tools, major tool updates, implementations, or other open-source repositories. Among the security tools released by NCC Group this year are:

- **Covenant v0.7:** Covenant is an open source .NET command and control framework that supports Red Team operations, similar in many ways to the well-known Cobalt Strike threat emulation software. NCC Group's Simone Salucci and Daniel Lopez Jimenez contributed a number of features to the project that they desired while using it. Some of those features include: Disabling ETW, Dumping Snapshot Processes, Process Injection Tasks, Payload Execution Guardrails and Other Pull Requests

- **GTFOBLookup:** An offline command line lookup utility for GTFOBins, LOLBAS, and WADComs, created by James Conlan

- **KilledProcessCanary:** A prototyped a Windows Service Canary in order to target parts of the ransomware kill chain to minimize impact and overall success of operations, created by Ollie Whitehouse. The tool was described in his blog post, Deception Engineering: exploring the use of Windows Service Canaries against ransomware.

- **Libptmalloc:** Heap analysis tooling for ptmalloc (pthreads malloc), and is interesting to those seeking to exploit glibc, created by Cedric Halbronn. This was part of the work discussed in the blog post, Exploiting the Sudo Baron Samedit vulnerability (CVE-2021-3156) on VMWare vCenter Server 7.0.

- **Log4j-jndi-be-gone:** A simple mitigation for CVE-2021-44228 (Log4Shell), created to help mitigate the log4j vulnerabilities which saw widespread exploitation in December 2021, by Jeff Dileo. This tool uses the Byte Buddy bytecode manipulation library to modify the at-issue log4j class's method code and short circuit the JNDI interpolation handler. It works by effectively hooking the at-issue `JndiLookup` class' `lookup()` method that Log4Shell exploits to load remote code, and forces it to stop early without actually loading the Log4Shell payload URL.

- **ML-for-RNGs:** The Jupyter notebooks underlying research exploring the utility of deep learning to predict the sequence of the (presumably) random output numbers using previously generated numbers without the knowledge of the seed for (non-cryptographic) PRNGs xorshift128 and Mersenne Twister, by Mostafa Hassan. While this research looked at a non-cryptographic PRNGs, we are interested, generically, in understanding how deep learning-based approaches to finding latent patterns within functions presumed to be generating random output could work, as a prerequisite to attempting to use deep learning to detect previously-unknown patterns in cryptographic (P)RNGs.

- **Ndsp-discover:** An Nmap script to identify Netgear Switch Discovery Protocol (NSDP) on UDP ports 63322 and 63324, by Manuel Ginés Rodríguez. This tool was created in support of Manuel's extensive vulnerability research on the Netgear ProSAFE Plus switches.

- **NLAhoney:** Source code to deploy honeypots that can capture RDP handshakes, then crack them offline in an effort to understand which passwords are being sprayed at RDP honeypots we deploy, created by Ollie Whitehouse and Ray Lai, as a part of their project, Cracking RDP NLA Supplied Credentials for Threat Intelligence.

- **Principal Mapper (v1.1.0 Update; v1.1.4 Update):** Principal Mapper, or PMapper, is a tool and library for in-depth analysis with AWS Identity and Access Management, as well as AWS Organizations. PMapper stores data about AWS accounts and organizations, then provides options to query, visualize, and analyze that data. The library is written in Python and was created by Erik Steringer.

- **Raccoon:** is a tool that aims to identify potential misconfigurations that could expose sensitive data within Salesforce. Specifically, it reveals where access has been granted to all records for particular objects of interest, by Jerome Smith.

- **Reliably-checked String Library Binding:** is a library binding that uses static array extents to improve diagnostics that can help identify memory safety flaws, created by Robert Seacord.

- **Ruby-trace:** A Low-Level Tracer for Ruby, created by Jeff Dileo and originally released to coincide with his DEF CON 29 talk on it and parasitic tracing in general. Version 1.1.0 adds support for the newly released Ruby 3.1 and includes a number of improvements, including alternate enable/disable hooks.

- **Shouganaiyo-loader:** A cross-platform Frida-based Node.js command-line tool that forces Java processes to load a Java/JVMTI agent regardless of whether or not the JVM has disabled the agent attach API, created by Jeff Dileo.

- **Sigwhatever:** For automated exploitation of netntlm hash capture via image tags in emails signatures. The tool was described in their blog post, Sign over Your Hashes – Stealing NetNTLM Hashes via Outlook Signatures created by David Cash, Rich Warren, Julian Storr.

- **SocksOverRDP:** This tool adds the capability of a SOCKS proxy to Terminal Services (or Remote Desktop Services) and Citrix (XenApp/XenDesktop). It uses

Dynamic Virtual Channel that enables us to communicate over an open RDP/Citrix connection without the need to open a new socket, connection or a port on a firewall. The author, Balazs Bucsay, presented the tool at the Cyber Security Global Summit by Geekle.

- **Solitude:** Solitude is an open source privacy analysis tool that enables you to conduct your own privacy investigations into where your private data goes once it leaves your web browser or mobile device. Whether a curious novice or a more advanced researcher, Solitude makes the process of evaluating an app's privacy accessible for everyone, created by Dan Hastings. Solitude was featured in media outlets including KitPloit, Hacking Land, and Hacking Reviews.

- **Squeak:** The tool was described in their blog post, MSSQL Lateral Movement, created by David Cash. This tool supports the work outlined within the blogpost: namely, the automation of lateral movement via MSSQL CLR without touching disk (besides a DLL being temporarily written to disk by the SQL Server process) or requiring XP_CMDSHELL.

- **UninstalledAppCanary:** This project deploys a number of canary apps which fire when uninstalled, motivated by the idea that certain threat actors uninstall a number of products prior to dropping later stages, and was created by Ollie Whitehouse. The tool was described in his blog post, Deception Engineering: exploring the use of Windows Installer Packages against first stage payloads, and builds upon prior work discussed in the blog post Deception Engineering: exploring the use of Windows Service Canaries against ransomware.

- **Wubes:** is like Qubes (a security-focused operating system that aims to provide security through virtualization) but for Microsoft Windows. The purpose is to leverage the Windows Sandbox technology to spawn applications in isolation, so that if you browse a malicious site using Wubes, it won't be able to infect your Windows host without additional chained exploits. Currently it supports spawning a Windows Sandbox for the Firefox browser but other applications can easily be added, and was created by Cedric Halbronn.

We also released source code for a variety of cryptographic implementations, exploit proofs-of-concept, and obfuscation reverse-engineering techniques:

- In January 2021, Jeff Dileo released proof-of-concept exploit code for his vulnerability, CVE-2020-15257, found in containerd - a container runtime underpinning Docker and common Kubernetes configurations - which resulted in full root container escape for a common container configuration. This was discussed in depth in his blog post titled, ABSTRACT SHIMMER (CVE-2020-15257): Host Networking is root-Equivalent, Again.

- In January 2021, Thomas Pornin published a blog post on Double-odd Elliptic Curves, which discussed some new elliptic curves he had created for cryptographic protocols. These were published on a dedicated website, doubleodd.group. He also released a complete whitepaper on the IACR ePrint archive, full of mathematical demonstrations, and several cryptographic implementations of double-odd elliptic curves in Python, Go, C, and Assembly.

- In June 2021, Cedric Halbronn published Exploit mitigations: Keeping up with evolving and complex software/hardware, seeking to address how it has become challenging to follow when certain exploit mitigations are added in an update and/or backported to some older versions of various software and hardware, by creating his mitigations tables which track mitigations exploit mitigations available across numerous operating systems (Windows, Linux, Android, iOS, OpenBSD, FreeBSD), architectures (ARM) and applications and versions, including the glibc library, Mozilla Firefox, Microsoft Edge, Google Chrome, and Microsoft Office.

- In September 2021, Eric Schorn released his implementations of montgomery multiplication in assembly, alongside a blog post titled Optimizing Pairing-Based Cryptography: Montgomery Multiplication in Assembly, which discussed and demonstrated selected optimizations found in pairing-based cryptography, foundational to the BLS Signatures central to Ethereum 2.0, zero-knowledge arguments central to Zcash, Filecoin, and other blockchain/cryptocurrency projects relying upon zk-proofs.

- In October 2021, Thomas Pornin implemented a mathematically-impossible lossless compression algorithm, alongside a blog post titled Paradoxical Compression with Verifiable Delay Functions, and released a paper on the IACR ePrint archive.

- In October 2021, Nicolas Guigo released tools and methods for reversing real-world binary obfuscation, through a blog post titled A Look At Some Real-World Obfuscation Techniques.

- Outside of research, in late 2021 our Strategic Threat Intelligence team created a public github repository, Threat-Intelligence-Alerts, where they publish alerts about major vulnerabilities, exploits, and mitigations on an ongoing basis.

Publicly Reported Security Audits

For many years, NCC Group has published publicly-reported security audits of critical components of open source software as well as select proprietary systems, including in past years for components of important systems including OpenSSL, SecureDrop, TrueCrypt, Tor, Docker, Keybase, Zcash, and many others.

Of these reports, those labelled "Public Report" were developed as a part of a paid engagement with an NCC Group client for NCC Group to conduct and publish the findings of a security audit on in-scope components. In 2021, NCC Group delivered 8 Public Reports across a number of different cryptographic implementations, as well as for the Google One VPN and WhatsApp End-to-End Encrypted Backups, among others..

Our 2021 Public Reports include:

- Public Report – BLST Cryptographic Implementation Review (January 2021)
- Public Report – VPN by Google One: Technical Security & Privacy Assessment (April 2021)
- Public Report – Dell Secured Component Verification (May 2021)
- Public Report – Protocol Labs Groth16 Proof Aggregation: Cryptography and Implementation Review (June 2021)
- Public Report – WhatsApp End-to-End Encrypted Backups Security Assessment (October 2021)
- Public Report – Zcash NU5 Cryptography Review (November 2021)
- Public Report – Zendoo Proof Verifier Cryptography Review (November 2021)
- Public Report - Whatsapp opaque-ke Cryptographic Implementation Review (December 2021)

Applied Cryptography

- In January 2021, Thomas Pornin published a blog post about his recent work on Double-odd Elliptic Curves. Double-odd curves are the elliptic curves whose order (number of points on the curve) is the double of an odd integer. About 1/4th of all curves are double-odd curves: this is a large class of curves, and Such curves have nominally been supported by generic standards such as ANSI X9.62 (ECDSA signatures over any curve) for the last two decades. He notes that the point of double-odd elliptic curves is to provide a safe structure, amenable to building cryptographic protocols on top of the "prime order group" abstraction, but that they also offer performance improvements over generic elliptic curves, making them especially useful for small embedded systems with severe constraints on power, RAM and CPU abilities. There is also a complete whitepaper published on the IACR ePrint archive, full of mathematical demonstrations; it also specifies the use of do255e and do255s in higher-level cryptographic functionalities (key pair generation, key exchange, signature, and hash-to-curve). Thomas has also released several cryptographic implementations, as well as a geometric introduction which helps build mathematical intuition related to double-odd curves.
- In January 2021, our Cryptography Services team (virtually) attended the Real-World Cryptography conference, and published a blog post in which they shared summaries and insights from some of our favourite talks from RWC 2021.
- In January 2021, we published the NCC Group Cryptography Services Public Report on the BLST Cryptographic Implementation Review. This involved a cryptographic implementation review of the BLST library, which implements support for the draft IETF specifications on Hashing to Elliptic Curves and BLS Signatures. The latter specification uses advanced cryptographic-pairing operations to feature aggregation properties for secret keys, public keys and signatures, which is central to the emerging Ethereum 2.0 Proof-of-Stake block-validation mechanism. This report was commissioned by Supranational, Protocol Labs and the Ethereum Foundation.
- In January 2021, Gérald Doussot published Software Verification and Analysis Using Z3. This post provided a technical introduction on how to leverage the Z3 Theorem Prover to reason about the correctness of cryptographic software, protocols and otherwise, and to identify potential security vulnerabilities. In this post, he covered both: (1) Modeling and analysis of an algorithm documented in an old version of the QUIC Transport protocol IETF draft, as well as (2) Modeling of specific finite field arithmetic operations for elliptic curve cryptography, with integers represented using a uniform saturated limb schedule (four limbs of 64 bits), to prove equivalence with arbitrary-precision arithmetic, and for test case generation. In February, researchers from Galois published a blog post, "Cryptol as an SMT Frontend," referencing Gérald's research, in which they checked the implementation of part of the QUIC protocol, and built upon this work in June 2021, "Cryptographic Assurance with Cryptol," which they explore an optimized implementation of field arithmetic.
- In February 2021, Eli Sohl published Cryptopals: Exploiting CBC Padding Oracles. This post - which was the first in a new series by Eli which offers educational walkthroughs of the beloved Cryptopals cryptography challenges - explored the classic padding oracle attack on CBC-mode block ciphers, through the lens of Cryptopals challenge #17.
- In June 2021 we published the NCC Group Cryptography Services Public Report on the cryptography and implementation review of Protocol Labs Groth16 Proof Aggregation. This was a cryptography and implementation review of the Groth16 proof aggregation functionality in the bellperson and two other related GitHub repositories. This code utilizes inner product arguments to efficiently aggregate existing Groth16 proofs while reusing existing powers of tau ceremony transcripts. This report was commissioned by Protocol Labs.
- In June 2021, NCC Group's Thomas Pornin - alongside external peers Liz Steininger, isis agora lovecruft, JP Aumasson, and Taylor Hornby - spoke on a panel on Auditing Cryptography at the Zcash Foundation's conference, Zcon2lite.

- In June 2021, Parnian Alimi published On the Use of Pedersen Commitments for Confidential Payments. This blog post looked at the Zether protocol, which uses ElGamal public key encryption to hide transaction amounts and utilizes zero-knowledge proofs to demonstrate the validity of a transaction to stakeholders in a financial blockchain, namely network validators, investors, and auditors. In general, this is important to the principle of transaction confidentiality, desirable for some financial blockchains, which requires hiding investors' account balances and transaction amounts, while enforcing compliance rules and performing validity checks on all activities. The post also explains underlying cryptographic building blocks including Pedersen Commitment, ElGamal Encryption, and Zero-Knowledge Proofs, and also discusses a number of implementation considerations including privacy (vs confidentiality), the front-running problem, side-channel attacks, forward secrecy, and integrated signing & encryption.

- In June 2021, Eric Schorn published his first post in a two-part code-centric blog post series about selected optimizations found in pairing-based cryptography. For his first post, Optimizing Pairing-Based Cryptography: Montgomery Arithmetic in Rust, he discussed selected optimizations found in pairing-based cryptography, foundational to the BLS Signatures central to Ethereum 2.0, zero-knowledge arguments central to Zcash, Filecoin, and other blockchain/cryptocurrency projects relying upon zk-proofs. This post covered modular Montgomery arithmetic from start to finish, including context, alternatives, theory, and practical working code in Rust running 9X faster than a generic Big Integer implementation. His second post, published in September 2021, was about Optimizing Pairing-Based Cryptography: Montgomery Multiplication in Assembly. This second post takes the Montgomery multiplication algorithm developed in Rust even further to seek the maximum performance a modern x86-64 machine can deliver from an implementation hand-written in assembly language, resulting in a Montgomery multiplication routine running more than 15X faster than a generic Big Integer implementation, due to several specialized instructions and advanced micro-architectural features enabling increased parallelism.

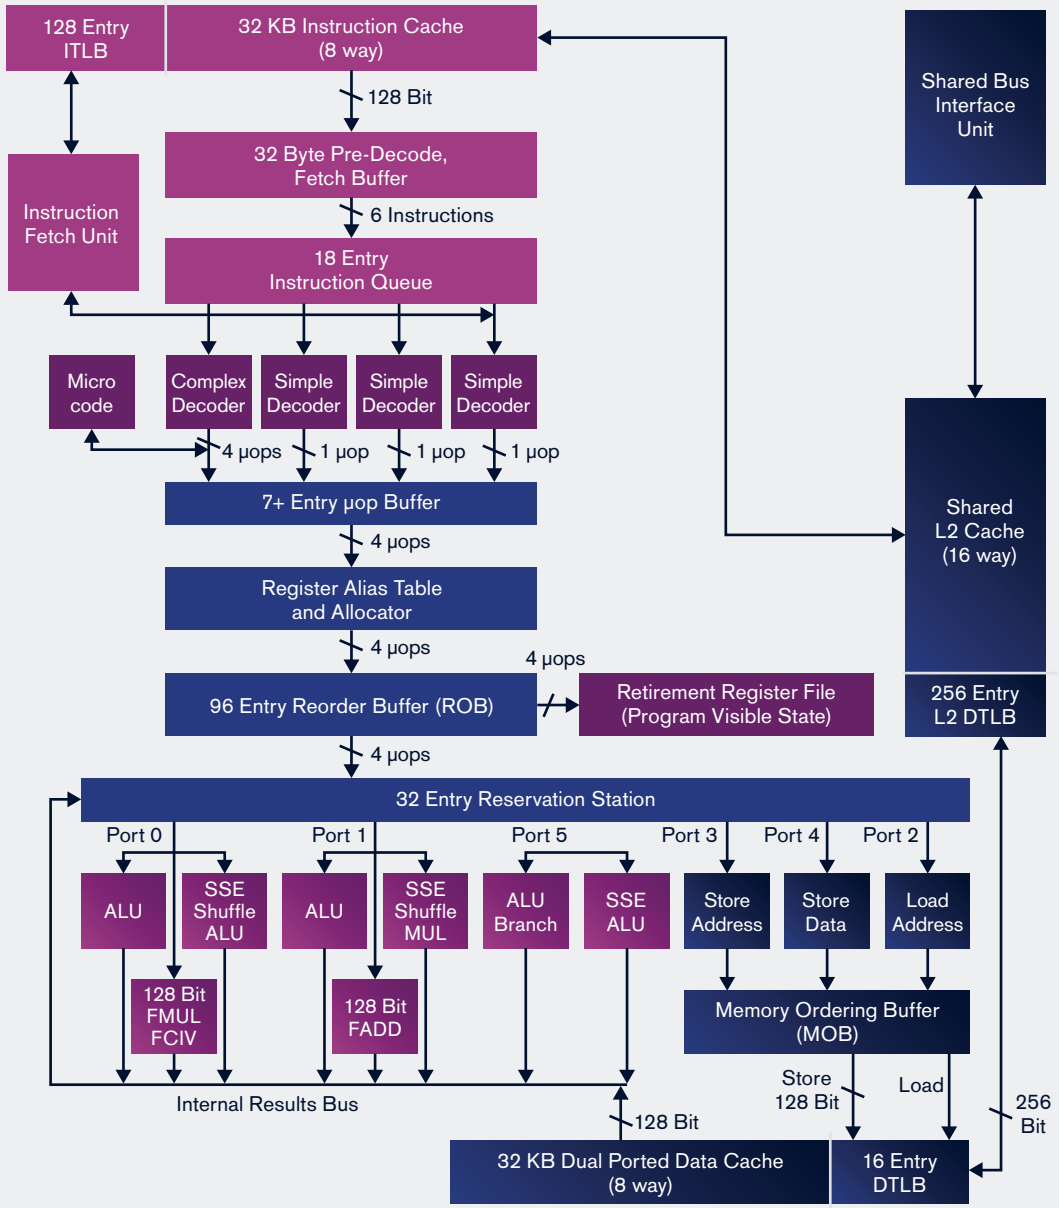

Intel Core 2 Architecture

- In September 2021, Javel Samuel presented an Overview of Open Source Cryptography Vulnerabilities at the International Cryptographic Module Conference, ICMC21. This presentation reviewed the foundations of cryptographic vulnerabilities as applicable to open-source software from a penetration tester's perspective over multiple public cryptography audit reports. This presentation also discussed what attacks in the past took advantage of these cryptography vulnerabilities and what the consequences were, as well as the success rate of various mitigations, as well as some thoughts on suggested focus areas for the future of open source cryptography.

- In October 2021, Thomas Pornin published Paradoxical Compression with Verifiable Delay Functions, which described (and implemented!) a mathematically-impossible compression algorithm, known as paradoxical compression, despite the mathematical impossibility of upholding three qualities (described in the paper/blog post) simultaneously. This is a good illustration of a fundamental concept of cryptography: namely that there is a great difference between knowing that some mathematical object exists, and being able to build it in practice.

- In October 2021, we published the NCC Group Cryptography Services Public Report on their WhatsApp End-to-End Encrypted Backups Security Assessment. This report is the result of a security assessment we performed of its End-to-End Encrypted Backups project. End-to-End Encrypted Backups is an hardware security module (HSM) based key vault solution that aims to primarily support encrypted backup of WhatsApp user data. This report was commissioned by WhatsApp.

- In November 2021, we published the Public Report – Zcash NU5 Cryptography Review. This report is the result of our review of the upcoming network protocol upgrade NU5 to the Zcash protocol (codenamed "Orchard"), and was commissioned by the Electric Coin Company.

- In November 2021, Paul Bottinelli published Technical Advisory – Arbitrary Signature Forgery in Stark Bank ECDSA Libraries (CVE-2021-43572, CVE-2021-43570, CVE-2021-43569, CVE-2021-43568, CVE-2021-43571). These critical vulnerabilities allowed arbitrary signature forgery by an attacker, and existed in several open-source cryptography libraries - one with over 7.3M downloads in the previous 90 days on PyPI, and over 16,000 weekly downloads on npm.

- In November 2021, Paul Bottinelli published An Illustrated Guide to Elliptic Curve Cryptography Validation. This blog post offered an illustrated description of the typical failures related to elliptic curve validation and how to avoid them in a clear and accessible way. This is important because Elliptic Curve Cryptography (ECC) is widely used to perform asymmetric cryptography operations, such as to establish shared secrets or for digital signatures, but insufficient validation of public keys and parameters is still a frequent cause of confusion, leading to serious vulnerabilities, such as leakage of secret keys, signature malleability or interoperability issues.

- In November 2021, Frans van Dorsselaer of Fox IT (a part of NCC Group) presented at CWI Symposium on Post-Quantum Cryptography.

- Based on the popularity of Eli Sohl's original February 2021 blog post on the Matasano Cryptopals challenges, in December 2021 he began a blog series in which he offers walkthrough tutorial videos explaining the solutions to the Cryptopals challenges, with the post Announcing NCC Group's Cryptopals Guided Tour.

- In December 2021, we published the NCC Group Cryptography Services Public Report – WhatsApp opaque-ke Cryptographic Implementation Review. In this work, we conducted a security assessment of the 'opaque-ke' library, an open source Rust implementation of the OPAQUE password authenticated key exchange protocol. The protocol is designed to allow password-based authentication in such a way that a server does not actually learn the plaintext value of the client's password, only a blinded version of the password computed using a verifiable oblivious pseudorandom function. This report was commissioned by WhatsApp.

- In December 2021, Jennifer Fernick presented Financial Post-Quantum Cryptography in Production: A CISO's Guide for the Financial Services Information Sharing and Analysis Center, FS-ISAC. In this talk, starting from the known fact that most of the public-key cryptography on the internet will be trivially broken by existing quantum algorithms, she covered strategic applied security topics to address this need for a cryptographic upgrade of the global financial cryptographic infrastructure, including: Financial services use cases for cryptography and quantum-resistance, and context-specific nuances in computing environments such as mainframes, HSMs, public cloud, CI/CD pipelines, customer-facing systems, third-party and multi-party financial protocols; Whether quantum technologies like QKD are necessary to achieve quantum-resistant security; Post-quantum cryptographic algorithms for digital signatures, key distribution, and encryption; How much confidence cryptanalysts currently have in the quantum-resistance of those ciphers, and what this may mean for cryptography standards over time, and designing extensible cryptographic architectures and deciding when to begin integrating PQC in a world of competing technology standards.

Attacking (& Attacking with) Machine Learning Systems

- In June 2021, Jennifer Fernick published a blog post titled Machine Learning for Static Malware Analysis, with University College London. For this project, we sought to determine the efficacy of various individual machine learning primitives - as well as ensemble methods of multiple algorithms - for the static classification of Windows binaries in terms of whether or not they are malicious. The full research paper, written by CDT PhD students Emily Lewis, Toni Mlinarevic, and Alex Wilkinson of University College London, ultimately demonstrated that it is possible to create number of different high-efficacy machine learning models to identify malicious executables on the basis of features which included PE headers, bytes n-grams, control-flow graphs and API call graphs, all of which performed well. Making use of ensemble methods, the researchers were able to achieve a classification accuracy of 98.9%, suggesting that the particular featureset and ensemble model used (multi-modal late fusion) is effective for the detection of malware binaries at scale. Media coverage on this work, asking "Can machine learning help detect zero day malware?" was published by The Cyber Post, SC Media, thisweekin4n6, and RedPacket Security.

- In June 2021, the NCC Group Data Science Team published Incremental Machine Learning by Example: Detecting Suspicious Activity with Zeek Data Streams, River, and JA3 Hashes. This post, which explores the incremental learning machine learning paradigm in a real-world setting, provides a simple example involving JA3 hashes showing how some of the foundational algorithms that enable incremental learning techniques can be applied to novelty detection (the first time something has happened) and outlier detection (rare events) on data streams derived from Zeek. This was also discussed in thisweekin4n6.

- As a part of an August 2021 presentation at Black Hat USA on open source security, Jennifer Fernick discussed the impact that scalable, machine-learning driven vulnerability detection, triage, and potential exploitation methods could have on open-source targets in the coming years, arguing that scalable vulnerability finding methods are dual use, in that they can benefit either defenders or attackers, and the growth of these tools may exacerbate asymmetries between attack and defense.

- In October 2021, Mostafa Hassan published a blog post series on Cracking Number Generators Using Machine Learning. In Part 1 of this series, he explored using deep learning to predict the sequence of the (presumably) random output numbers using previously generated numbers without the knowledge of the seed for the (non-cryptographic) PRNG, xorshift128, essentially to break the PRNG. In Part 2 of this series, he demonstrated the effectiveness of deep learning techniques against the so-called randomness of the (non-cryptographic) PRNG, Mersenne Twister. While this research looked at a non-cryptographic PRNGs, we are interested, generically, in understanding how deep learning-based approaches to finding latent patterns within functions presumed to be generating random output could work, as a prerequisite to attempting to use deep learning to find previously-unknown patterns in cryptographic (P)RNGs, as this could potentially serve as an interesting supplementary method for cryptanalysis of these functions. In these blog posts, we show our work in beginning to explore this space. This work was also discussed by hackaday.

- In December 2021, Margit Hazenbroek published Encryption Does Not Equal Invisibility – Detecting Anomalous TLS Certificates with the Half-Space-Trees Algorithm. This blog post outlined an approach to detecting suspicious TLS certificates using an incremental anomaly detection model. Specifically, this model utilized the Half-Space-Trees algorithm and provides NCC Group's Managed Detection and Response Security Operations Center (SOC) with the opportunity to detect suspicious behavior, in real-time, even when network traffic is encrypted.
- In December 2021, NCC Group commenced our annual research project with the Centre for Doctoral Training in Data Intensive Science (CDT in DIS) at University College London (UCL), where we will in 2022 be studying the use of Generative Adversarial Networks (GANs) for fuzzing.
- In December 2021, Jennifer Fernick published On the malicious use of large language models like GPT-3, in which she asked, "Can large language models generate exploits?" This blogpost explored the question of whether (and how) large language models like GPT-3 or their successors may be useful for exploit generation, and proposed an offensive security research agenda for large language models.

Into 2022, we will be continuing research on several A.I. security related projects including:

- Real-world, practical attacks on machine learning systems, including both a whitepaper as well as new findings and attacks
- Machine learning techniques for vulnerability finding in C source code, studying the comparative performance between graph-based methods and NLP (natural language processing)-based methods
- New work in the use of Generative Adversarial Networks (GANs) to improve fuzzer performance
- Studying the security and privacy implications underlying chatbots, and studying an attack model that treats chatbots as an oracle for private or sensitive training data

Misinformation, Deepfakes, and Synthetic Media

- In October 2021, Jennifer Fernick discussed the problem of deepfake detection in the Threatscape 2023 and Beyond panel at MapleSec, also covered by IT World Canada. In this presentation, she challenged the idea that deepfakes could ever be reliably detected, warning that the use of machine learning to both generate and detect deepfakes would result in a cyclical competition in which deepfake generation algorithms improve their output's evasive capability against a given detection model, while detectors hone feature-detection, in a never-ending cycle.
- In November 2021, Swathi Nagarajan published Vaccine Misinformation Part 1: Misinformation Attacks as a Cyber Kill Chain. In this work, she used the Cyber Kill Chain model to describe vaccine misinformation attacks online, as well as to describe interventions against vaccine misinformation at each state of the killchain.

Into 2022, we will be continuing research that takes a security researcher's toolkit and applies it to the problems of misinformation, deepfakes, and synthetic media, through ongoing projects including:

- Subverting facial recognition systems, including broadly-applicable, generic, "skeleton key" attacks on facial recognition systems (or as we're calling them - "Eigenfaces")
- Authenticating against deepfakes
- Understanding the offensive use of audio deepfakes in fraud, abuse, and account takeover
- Experimental tooling for the automated debunking of misinformation
- Further security research and digital forensics methods for combating misinformation

Reducing Vulnerabilities at Scale & Improving Open Source Security

Software Supply Chain Security & Securing the Open Source Ecosystem

- In February 2021, Jennifer Fernick of NCC Group (alongside colleagues from across OpenSSF) presented a panel at FOSS Backstage, titled Frontiers in Securing the Open Source Ecosystem. In this panel, the panelists explored challenges and opportunities in securing the open source software ecosystem against a range of threat actors through a variety of interventions at all phases of the software development lifecycle, giving a brief overview of the mission, priorities, and current work within the Open Source Security Foundation (openssf.org), including an end-to-end threat model of the open source ecosystem, and a discussion of some of the most pressing issues in the security of open source software.
- In February 2021, Jennifer Fernick was interviewed by the GitHub README Project for an article about "How InfoSec pros keep open source safe—and how you can help".

> "I think we have yet to see the true potential of techniques for finding vulnerabilities at scale," Fernick says. "Large-scale fuzzing projects, vulnerability discovery query languages such as GitHub's CodeQL, innovations in program analysis, applications of machine learning to identifying examples of particular bug classes, and recent research in automated exploit generation (AEG) have yet to, in my opinion, become fully realized, and are likely to shift the security landscape."

- In May 2021, the OpenSSF's Identifying Security Threats Working Group announced their Security Metrics initiative. This working group, which began in mid-2020, was initially led by Jennifer Fernick of NCC Group and Michael Scovetta of Microsoft, aimed to create "dashboards" which enable developers and other users of open-source codebases to make security-informed decisions about the relatively security of their components, by providing both a graphical and API dashboard which summarizes key security metrics about a repo. Work in this group has also inspired and led to things like Google's Security Scorecards this year, as well as the excellent Security Metrics project currently led by Scovetta.
- In August 2021, Jennifer Fernick of NCC Group and Christopher Robinson of Intel presented Securing Open Source Software - End-to-End, at Massive Scale, Together at Black Hat USA. In this presentation, we shared key lessons learned in our experience coordinating the industry-wide remediation of some of the most impactful vulnerabilities ever disclosed (Heartbleed, Shellshock, Rowhammer, and BlueZ), presented a threat model of the many unmitigated challenges to securing the open source ecosystem, and shared new data which illustrates just how fragile and interdependent the security our core infrastructure can be, debate the challenges to securing OSS at scale. We also spoke oft-unspoken truths of coordinated disclosure and where it can fail, and discussed research advances that are making it easier for adversaries to find and exploit vulnerabilities at scale, offering guidance for how members of the security community can get involved and contribute meaningfully to improving the security of OSS - especially through coordinated industry-wide efforts. This work was covered by media outlets including Security Boulevard, the Veracode blog, Linux Today, SiliconANGLE, eSecurityPlanet, and Dark Reading.
- In September 2021, the Open Source Security Foundation published their Guide to implementing a coordinated vulnerability disclosure process for open source projects, available here on GitHub, and described here in a blog post. Jennifer Fernick of NCC Group contributed to this guide, as well as to the broader, ongoing work of the OpenSSF Vulnerability Disclosure Working Group, which will be disseminated through conference presentations and other developer outreach in 2022.
- Since prior to OpenSSF's founding in August 2020, Jennifer has been a founding member of the prior Governing Board as well as the Technical Advisory Committee, and made considerable contributions in the first several months of 2021 to help drive improved decision-making and governance within OpenSSF, to help advise senior US Congressional staffers on supply chain security, and to contribute to the technical vision for OpenSSF. She has also begun to advocate for the creation of a security incident response helpline for under-resourced open source projects who believe their project to be under active attack and require emergency intervention.

- In November 2021, Jennifer Fernick was an invited Keynote Speaker at the Linux Foundation Member Summit, where in her presentation, Securing Open Source Software, she made a case for why coordinated efforts to secure the open source ecosystem are urgently needed to strengthen software supply chain security. She and her co-speaker, David Wheeler of the Linux Foundation, also highlighted progress made since OpenSSF's founding in August 2020. A recording of this presentation is available here.

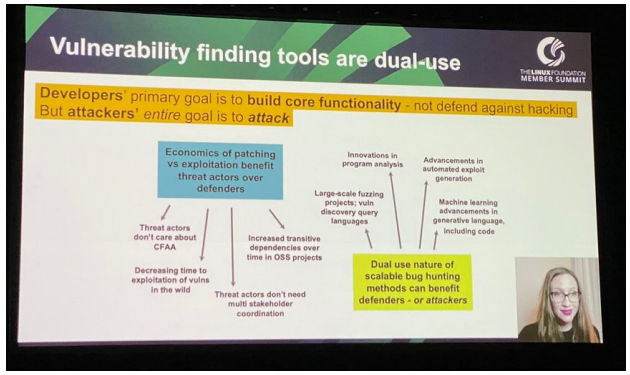

- Later that month, Jennifer was also voted onto the new Governing Board of the Open Source Security Foundation as the new General Member Representative, where joins senior technical executives from major technology companies including AWS, Cisco, Dell, GitHub, Google, Facebook, IBM, Intel, Microsoft, Oracle, Red Hat, VMWare, and more to help lead an industry-wide effort to help secure the open source ecosystem. This was discussed in detail in this press release.

- In the wake of the Log4Shell vulnerability (CVE-2021-44228, as well as CVE-2021-45046 and CVE-2021-44832) affecting the open-source project, Log4J, we responded in a number of ways: Firstly, by publishing our threat intelligence blog post, Log4Shell: Reconnaissance and post exploitation network detection which was frequently updated between December 12th and 29th to include Suricata network detection rules that can be used not only to detect exploitation attempts, but also indications of successful exploitation, as well as a list of indicators of compromise, relevant pcaps, and a number of references for prevention and mitigation. Secondly, Jeff Dileo released log4j-jndi-be-gone: A simple mitigation for CVE-2021-44228, which can be used to stop log4j from loading classes remotely over LDAP, preventing malicious inputs from triggering the "Log4Shell" vulnerability and gaining remote code execution on affected systems. We also maintained a highly-active meta thread on Reddit on the log4j 0day being exploited covering the unfolding events of Log4J, discussing advice, detection, response and remediation. Our hot patch, as well as our threat intelligence were cited in the US Cybersecurity & Infrastructure Security Agency's Alert, "Mitigating Log4Shell and Other Log4j-Related Vulnerabilities" on December 22 2021.

Secure Programming

- In March 2021, Robert Seacord (in collaboration with Jens Gustedt of INRIA) presented their paper on C language mechanism for error handling and deferred cleanup at the ACM/SIGAPP Symposium on Applied Computing. This paper introduces the implementation of a C language mechanism for error handling and deferred cleanup adapted from similar features in the Go programming language, which improves the proximity, visibility, maintainability, robustness, & security of cleanup and error handling over existing language features. This feature is currently under consideration for inclusion for forthcoming versions of the C programming language standard. A reference implementation of the features described by this paper is also available under an open source (BSD) license. This paper was also published in the conference's proceedings, SAC '21: Proceedings of the 36th Annual ACM Symposium on Applied Computing.
- In March 2021, Robert Seacord published a recording of his internal conference presentation, The Future of C Code Review, in which he discussed optimizations resulting from pointer provenance-based alias analysis in the C programming language that can modify the behavior of code with undefined behaviors, ultimately explaining what pointer provenance is, how it can introduce security flaws into your C program, and how to spot the problem and repair it.
  This presentation was based upon his Draft Technical Specification, "A Provenance-aware Memory Object Model for C," submitted to ISO TC 1/SC22/WG14, the C Standards Committee, in September 2020.
- In May 2021, Robert Seacord published a blog post in which he argues that supply chain security begins with secure software development.
- In July 2021, Robert Seacord released Reliably-checked String Library Binding, a library binding that uses static array extents to improve diagnostics that can help identify memory safety flaws. This tool is part of his work on broader initiatives in the C Standards Committee to improve bounds checking for array types.
- In August 2021, Robert Seacord presented his training, Secure Coding in C, at Black Hat USA, which offered detailed explanation of common programming errors in C and C++ and describes how these errors can lead to code that is vulnerable to exploitation, and demonstrates specific remediation techniques as well as general secure coding practices that help prevent the introduction of vulnerabilities in C and C++ codebases.
- In September 2021, Robert Seacord presented Why Can't Johnny Code Securely? at CyberEd 2021. This presentation discussed how software developer demand continues to increase faster than universities can produce qualified graduates, increasing the software development skills gap, and argued that graduates are inadequately prepared to write secure code, which is a problem because inadequately prepared programmers tend to experiment randomly until they find a combination that appears to work but produces code that works only under optimal conditions but is insecure. This talk explored this problem and offered solutions toward developing a more efficient, effective, and skilled software developer workforce, more capable of writing secure code.
- In September 2021, Robert Seacord presented Secure Coding at Auto-ISAC Analysts.
- In October 2021, Robert Seacord submitted 5 papers to the C Standards Committee (ISO/IEC WG14), including:

  – ISO/IEC WG14 - Identifier Syntax using Unicode Standard Annex 31

  – ISO/IEC WG14 - Clarifying integer terms

  – ISO/IEC WG14 - Clarifying integer terms v2

  – ISO/IEC WG14 - calloc overflow handling

  – ISO/IEC WG14 - Annex K Repairs

  He explained what each of these meant for the future security of the C programming language in his blog post, The Next C Language Standard, C23.
- In October 2021, Robert Seacord presented Volatility Ahead at NDC TechTown 2021, which explained - in the context of the C programming language - where volatile is useful, what the C and C++ standards say and how they got there, and finally suggest how the standards might be revised in the future. He also presented Secure Coding in C and C++, a two day training course that provides a detailed explanation of common programming errors in C and C++ and describes how these errors can lead to code that is vulnerable to exploitation.

Standardization & Verification

In addition to the C Standards work outlined above, NCC Group made standards-related contributions in the following ways:

- In February 2021, Matt Lewis released his paper, Investigating Potential Security Vulnerability Manifestation through Various Analyses & Inferences Regarding Internet RFCs (and how RFC Security might be Improved), to help understand how and why security vulnerabilities manifest, from design to implementation. This research parsed IETF RFCs, extracting RFC data and metadata into graph databases to explore and query relationships between different properties of RFCs. The ultimate aim of this work was to use any key observations and insights to stimulate further thought and discussion on how and where security improvements could be made to the RFC process, allowing for maximised security assurance at protocol specification and design so as to facilitate security and defence-in-depth. The research showed the value of mining large volumes of data for the purpose of gaining useful insights, and the value of techniques such as graph databases to help cut through the complexities involved with processing and interpreting large volumes of data.
- In July 2021, Ollie Whitehouse (alongside external collaborators Kirsty Paine of the UK National Cyber Security Centre and James Sellwood of Twilio) published the IETF Draft, Indicators of Compromise (IoCs) and Their Role in Attack Defence. This draft reviews the fundamentals, opportunities, operational limitations, and best practices for the use of Indicators of Compromise (IoCs) to identify, trace, and block malicious activity in networks or on endpoints. It highlights the need for IoCs to be detectable in implementations of Internet protocols, tools, and technologies - both for the IoCs' initial discovery and their use in detection - and provides a foundation for new approaches to operational challenges in network security. This work was presented at IETF 111, the presentation for which can be viewed here.
- In September 2021, Matt Lewis presented at the Safety, Security & Verification in Critical Systems Conference in Manchester.
- In November 2021, Matt Lewis of NCC Group alongside Mark McFadden of Internet Policy Advisors LTD (an expert on the development of global internet addressing standards and policies, and an active contributor to work in the IETF and ICANN), presented at the IETF Internet Architecture Board Workshop on Analyzing IETF Data. In this presentation, they sought to baseline how RFC Security Considerations - including RFC3552 (Guidelines for Writing RC Text on Security Considerations) - should be expressed and improved, seeking to use improvements to the RFC process itself to improve the security of the resulting standards.

Into 2022, we will be continuing research on reducing vulnerabilities at scale through ongoing yet-unpublished projects such as:

- Rewriting BSD kernel modules in memory-safe languages
- Continued contributions to the work of the Open Source Security Foundation, both technically and from a leadership perspective, including toward improving coordinated vulnerability disclosure and incident response for open source projects
- New contributions to programming language standards and internet architecture specs
- Creating an "entomology" of security bug types observed in Rust
- Continued experimentation with vulnerability-finding query languages, including CodeQL
- A whitepaper on Automated Exploit Generation (AEG)

Virtualization, Emulation, and Containerization

- In January 2021, Jeff Dileo released proof-of-concept exploit code for his vulnerability, CVE-2020-15257, found in containerd - a container runtime underpinning Docker and common Kubernetes configurations - which resulted in full root container escape for a common container configuration. This was a flaw we disclosed in late 2020 for which a technical deep-dive is available in a blog post entitled ABSTRACT SHIMMER (CVE-2020-15257): Host Networking is root-Equivalent, Again.

- In May 2021, Jeff Dileo and ex-NCC Group colleague Addison Amiri presented their research on dRuby Security Internals at NorthSec. This presentation discussed how dRuby (a "distributed object system" built into Ruby - think CORBA or Java's RMI) works, where its insecurities lie, and how it is much more insecure than previously understood to be, "which is a feat, considering that dRuby already provides code execution as a service." In this presentation they discussed the dRuby API, its internals, and its underlying wire protocol — and how they make dRuby fundamentally unsafe — and demonstrated several novel proof-of-concept exploits targeting dRuby servers and clients, the latter of which have not been known to be vulnerable until this research was conducted. While dRuby is arguably well-known to be a readily exploitable service enabling remote code execution, this research shows how the underlying protocol exposes a number of additional risks that enable not only alternate methods of compromising dRuby services, but also the means to compromise dRuby clients. They also critiqued some of the existing advice and documentation for "securing" dRuby and how it fails to guard against dRuby's inherent issues, as well as the researchers' own efforts to harden dRuby, including the kinds of protocol, logic, and API changes needed to negate its issues. This presentation closed with a discussion of the insecurity of existing dRuby exploits, and showed how you can penalize your pentesters for using off-the-shelf exploits, including that Metasploit's exploit for dRuby used the standard dRuby library to exploit dRuby, thus making it vulnerable to both standard and novel dRuby exploitation itself. A recording of this presentation is also available here.

- Also at NorthSec, Jeff Dileo joined a number of other researchers on a panel on Vulnerability Research.

- Throughout the year, Brian Hong presented Sleight of ARM: Demystifying Intel Houdini at a variety of conferences including Ekoparty, DEF CON 29 (Main Track), ToorCon and the Black Hat USA Briefings. This research targeted Intel's proprietary Houdini binary translator, which runs ARM binaries on x86 platforms, such as higher-end Chromebooks and desktop Android emulators. This presentation began with a high-level discussion of what we uncovered about how Houdini works and is loaded into processes, then dived into the low-level internals of the Houdini engine and memory model, including several security weaknesses it introduces into processes using it. Brian concluded by discussing methods to escape the Houdini environment, execute arbitrary ARM and x86, and write Houdini-targeted malware that bypasses existing platform analysis.

- Iain Smart also taught the training Mastering Container Security V5 - Black Hat edition at Black Hat USA. The course covered Docker and how Linux containers work, fundamental Linux security concepts, and container orchestration and clustering, looking at how Kubernetes works and the security pitfalls that can leave the clusters and cloud-based environments which use containers exposed to attack, as well as practical examples of configuring key Kubernetes security controls such as RBAC, PodSecurityPolicies and Network Policies.

- At DEF CON 29, Jeff Dileo presented Instrument and Find Out: Writing Parasitic Tracers for High(-Level) Languages. This presentation discussed the process for developing generalized parasitic tracers targeting specific programming languages and runtimes using Ruby as our case study, showing the feasibility of writing external tracers targeting a language and its runtime, challenging the notion (and quality) of the performance monitoring and introspectability features of (some) modern programming languages by writing his own implementation and instrumenting them into the language dynamically.
- In August 2021, Jeff Dileo also published Some Musings on Common (eBPF) Linux Tracing Bugs, which discussed an insecure coding pattern commonly used in system observability and program analysis tools, and several techniques that can enable one to evade observation from such tools using that pattern, especially when they are being used for security event analysis.
- In August 2021, Iain Smart was interviewed on how container vulnerabilities can put the software supply chain at risk, discussing how compromised external dependencies can enable an attacker to access and even potentially modify a build process.
- In September 2021, Iain Smart published a blog post titled NSA & CISA Kubernetes Security Guidance – A Critical Review, in which he critiqued the recent Cybersecurity Technical Report (CTR) on Kubernetes Hardening Guidance, authored by the United States' National Security Agency (NSA) and Cybersecurity and Infrastructure Security Agency (CISA). In this post, he gave an overview of the guidance, and highlighted specific places where the NSA/CISA guidance overlooked important aspects of Kubernetes security, or where the guidance was out of date at time of publication, as well as offering considerations for some of the more common complex use cases not covered by the CTR guidance, including useful audit configurations that won't require excessive levels of compute power or storage, advice on handling external dependencies, and some notes around the complex mechanisms of Kubernetes RBAC.
- In October 2021, Jeff Dileo discussed container security on the Cyberwire podcast.
- In November 2021, Ben Lister and Kane Ryans published Detection Engineering for Kubernetes clusters. This blog post detailed the collaboration between NCC Group's Detection Engineering team and Containerization & Orchestration practice in tackling detection engineering for Kubernetes, including both a description of the detection engineering team's more generic methodology around detection engineering for new/emerging technologies and how it was used when developing detections for Kubernetes-based attacks, as well as detailing the novel detection rules they have created around how privilege escalation is achieved within a Kubernetes cluster, to better enable security operations teams to monitor security-related events on Kubernetes clusters and thus to help defend them in real-world use.

Into 2022, we will be continuing research on virtualization, emulation, and container security, including but not limited to:

- Container tracing and other research pertaining to container runtimes
- Transient execution vulnerabilities
- Security audits and associated research related to key open source projects in this space

Hardware & Embedded Systems

- In March 2021, Eric Evenchick offered the training Reverse Engineering Firmware with Ghidra at Black Hat Spring Trainings. This hands-on course taught the concepts, tools, and techniques required to reverse engineer firmware and assess embedded devices, using Ghidra, a popular, extensible, and powerful open-source reverse engineering tool that supports many different processor architectures, developed by the National Security Agency. In this training, participants used binary reverse engineering techniques to find exploitable vulnerabilities in an embedded Linux device, mapped device vector tables, peripheral memory, and system calls to find exploitable vulnerabilities in a bare-metal device, and identified remotely exploitable vulnerabilities in a Bluetooth Low Energy device. In July 2021, this training was also offered at Hardwear.io USA.

- In May 2021, we published our Public Report - Dell Secured Component Verification, which was commissioned by Dell to explore supply chain security functionality and related and supportive foundational security functionality on 14th and 15th generation Dell servers.

- In May 2021, Eric Evenchick presented Building CANtact Pro: An Open Source CAN Bus Tool at NorthSec. In this talk, he discussed the design and release process for the CANtact Pro device, to help would-be tool creators interested in launching their own hardware product to understand the many things that go into bringing a hardware idea to market, from PCB design to driver development. He also discussed open source tools for designing PCBs, writing cross-platform drivers using Rust, the economics of releasing a device, and logistical challenges one can expect when building hardware.

- In July 2021, Sultan Qasim Khan and Jeremy Boone published a three-part blog series on Alternative Approaches for Fault Injection Countermeasures. In Part 1 of this series, they covered the basic principles of fault injection – types of glitches, their effects, and how an attacker can characterize hardware and firmware to achieve a successful glitch. In Part 2 of this series, they discussed various C functions, macros and programming patterns that can be used to achieve double glitch resistance within software. In Part 3 of this series, they enumerated the drawbacks of common software-based glitching countermeasures, and outlined alternative countermeasures to fault injection, including instruction duplication, memory store verification, and some forms of control flow integrity intended to provide fault detection.

- In July 2021, Rob Wood published Practical Considerations of Right-to-Repair Legislation. In this blog post, he offered OEM, device owner, and US Federal legislators' perspectives on right-to-repair, and discussed how device security requires trade-offs and compromises with usability, performance, cost, and repairability, and outlined specific implications of the proposed US legislation and how an OEM might alter their designs to comply.

- In August 2021, Jon Szymaniak presented Depthcharge: A Framework for U-Boot Hacking at DEF CON 29 Demo Labs. Depthcharge is an extensible Python 3 toolkit designed to aid security researchers when analyzing a customized, product-specific build of the U-Boot bootloader. Depthcharge was first released at a Hardwear.io session in 2020, and also presented at the Open Source Firmware Conference, in a talk titled, Guiding Engineering Teams Toward a More Secure Usage of U-Boot.

- In August 2021, Diana Dragusin published, The ABCs of NFC chip security, which explored how the complexity of the NFC technology and standardization ecosystem is contributing to security weaknesses. She also surveyed a range of NFC chips by a wide and representative range of vendors,and published findings in terms of things like user memory protections and system configuration protections, and demonstrating the importance of product threat-modelling when selecting which NFC chips to use in new hardware products.

- In December 2021, Catalin Visinescu published, Why IoT Security Matters, in which he makes a case for why IoT security matters, and walks us through a threat modelling exercise to understand the different ways internet-connected devices can be attacked.

- Later that month, Catalin Visinescu also published Choosing the Right MCU for Your Embedded Device -- Desired Security Features of Microcontrollers. This article illuminated important security criteria that must be evaluated when choosing the right MCU component for an embedded systems project, helping to establish the questions engineers should ask chip vendors before deciding what is the best microcontroller for their new product.

- In December 2021, Simon Watson published FPGAs: Security Through Obscurity?, which discussed emerging use cases of FPGAs, the changing FPGA technology landscape & how it affects security, recent FPGA vulnerabilities (including Starbleed, Thangrycat, JackHammer, & others), as well as attacks & defenses for FPGAs.

5G Security & Smart Environments

- In June 2021, Traficom – the Finnish transport and communications agency – along with the Aalto University, Cisco, Ericsson, Nokia, and PwC, organized the 5G Cyber Security Hack competition, and annual hackathon-style event in which around 130 security experts from around the world participated in hacking challenges relating to 5G technology and its use cases. For the 2021 competition, NCC Group won first-place in their challenge, through the talented efforts of Ross Bradley, Eva Esteban Molina, Phil Marsden and Mark Tedman.
- In October 2021, Mark Tedman published The Challenges of Fuzzing 5G Protocols. In this post, he discusses the specific challenges unique to fuzzing 5G telecommunications protocols using both proprietary and open source fuzzers (including Fuzzowski 5GC, Frizzer, and AFLNet), as well as the relative strengths and weaknesses of the fuzzers studied for their efficacy at 5G protocol fuzzing of the NGAP, GTPU, PFCP, and DIAMETER protocols. It also discussed details of the testing environment, a sample of the vulnerabilities found, the comparative performance of the tested fuzzers, as well as lessons learned for telecommunications protocol fuzzing in general.

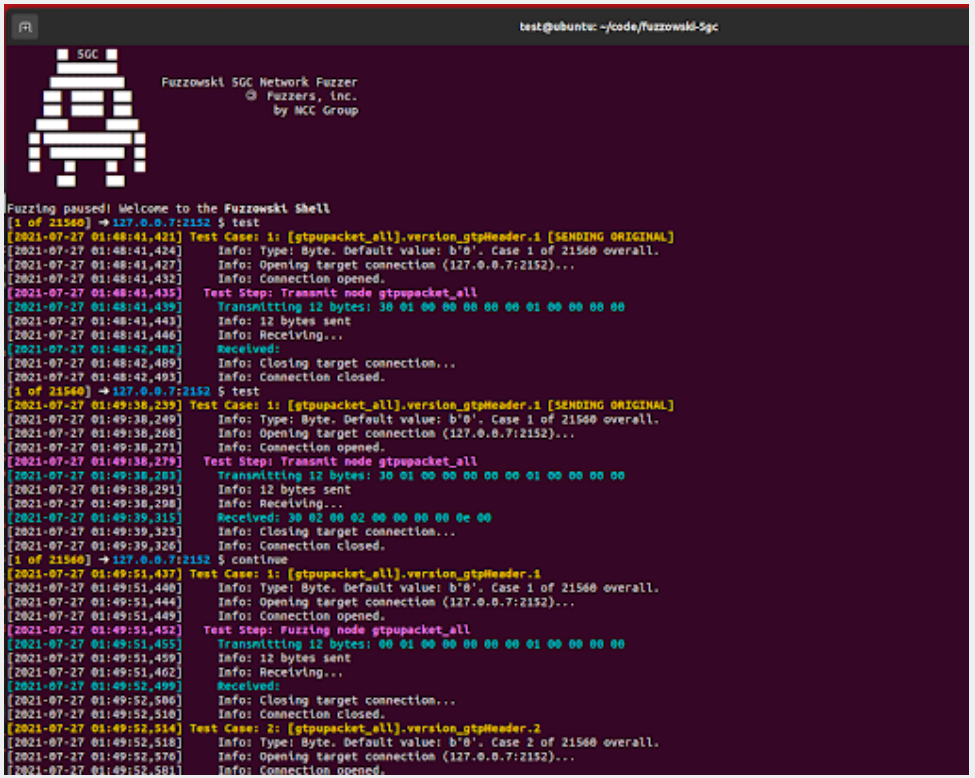

- In October 2021, Mark Tedman also published Technical Advisory – Open5GS Stack Buffer Overflow During PFCP Session Establishment on UPF (CVE-2021-41794). This denial of service vulnerability was found in the open-source implementation of 5G protocols, Open5GS, as a part of a larger research effort into protocol-level vulnerabilities in 5G communications networks.
- In November 2021, Mark Tedman published Exploit the Fuzz – Exploiting Vulnerabilities in 5G Core Networks, in which he exploited his previously-disclosed PFCP bug (CVE-2021-41794) found in earlier 5G core fuzzing efforts, using Fuzzowski 5GC against Open5GS.
- In November 2021, Philip Marsden presented at Cyber Senate Control Systems Cybersecurity Europe Conference on the 5G Threat Landscape, in which he discussed a range of threats to 5G security, including those relating to systems architecture, policy, standards, and the security of the hardware and software supply chain.
- Throughout 2021 and into 2022, Daniel Romero and colleagues have been studying implementations of the LoRaWAN protocol, initially with the intention to understand & improve the security testing of LoRaWAN networks. However, as the project progressed, it has expanded to become both a rigorous study of the LoRaWAN protocol itself, as well as an end-to-end study of the LoRaWAN network environment, from end-devices to final applications. In addition to creating security testing methodologies, one of the main outcomes this research has provided is the FLoRa framework and associated tooling. We have developed a modular tool which allows testing of the security of full LoRaWAN network implementations, and that has recently been updated with novel protocol vulnerabilities identified during a joint research with an academic research partner, to be published in 2022.
- Another ongoing project involves research into the security of the 4G/LTE Control Plane. During this research, several vulnerabilities that affect the LTE protocol and its implementation are being studied, as is the impact of their exploitation.

Into 2022, we will be continuing research on various aspect of 4G/LTE and 5G security, as well as protocols including LoRaWAN, and a range of research projects on smart environments, including:

- Continued work on LoRaWAN tooling & security testing, including collaborative work with the University of Surrey
- Vulnerability assessments of IoT components of smart environments, including smart locks and smart alarm systems, as well as smart buildings as a whole
- Continued 4G/LTE control plane research
- 5G Baseband SoC Research

Public Interest Technology

In NCC Group's Research division, we have a dedicated research working group which offers paid research time and other resources to research projects conducted in the public interest, in an effort to support security and privacy research for the greater good of society which might not otherwise have resources available to support it. It is my hope that this group will continue to grow with each passing year.

This year, our researchers focussed on a deep and broad analysis of the security & privacy implications of different vaccine passport apps around the world, the important topic of racial injustice in algorithmic decision-making, and mobile privacy from the perspectives of both users seeking to understand private data leakage from their favourite apps, as well as the true privacy impact of mobile-device-management (MDM) apps that employees are often asked to use in a bring-your-own-device (BYOD) scenario. We also continued our longstanding research partnership with independent UK consumer body Which?, where we studied a number of consumer IoT devices and reported on their security and privacy flaws to help consumers make informed choices about the devices they allow into their lives.

Vaccine Passport Security

- In October 2021, Sid Adukia published Assessing the security and privacy of Vaccine Passports. In this blog post, he offered a framework of security and privacy considerations for vaccine credential systems, through threat modelling the different ways that they could be subverted, especially in ways that are harmful to users.
- This work was also complemented by Sid's publication of Technical Advisory – New York State Excelsior Pass Vaccine Passport Credential Forgery which showed that an individual would be able to to create and store fake vaccine credentials in their NYS Excelsior Pass Wallet that might allow them to gain access to physical spaces (such as businesses and event venues) where they would not be allowed without a vaccine credential, even when they have not received a COVID-19 vaccine. The impact of this was vast, affecting all whose safety in physical event spaces in New York State depended upon the validity of people's proof-of-vaccination, in a state with a population of over 20 million people.
- Around the same time, Dan Hastings published Technical Advisory – New York State Excelsior Pass Vaccine Passport Scanner App Sends Data to a Third Party not Specified in Privacy Policy. The New York State (NYS) Excelsior Scanner App is used by businesses or event venues to scan the QR codes contained in the NYS Excelsior wallet app to verify that an individual has either a negative COVID-19 test or their vaccination status. In this work, Dan showed that some data about the businesses/ event venues using the app to scan QR codes is also sent to a third-party analytics domain, but that this was not specified in the app's privacy policy.
- Combined, Sid and Dan's work on the security & privacy impacts of some of the early vaccine credential systems received ample media coverage, including in a podcast interview with Security Ledger, as part of a story on Vox Recode and in POLITICO, and the forged credential vulnerabilities were covered in 18+ global publications including ZDNet, InfoRiskToday, Fuentitech, and DataBreachToday Europe.
- In December 2021, Drew Wade, Emily Liu, and Sid Adukia published Exploring the Security & Privacy of Canada's Digital Proof of Vaccination Programs. In this work, they studied a range of Canadian provinces' proof-of-vaccination apps to analyze their associated security and privacy properties. They also offered a deep-dive exploration of the Verify Ontario app in depth to understand its ability to correctly reject forged and malformed vaccine credentials, to avoid collecting unnecessary private data, and to manage data using security best-practices.

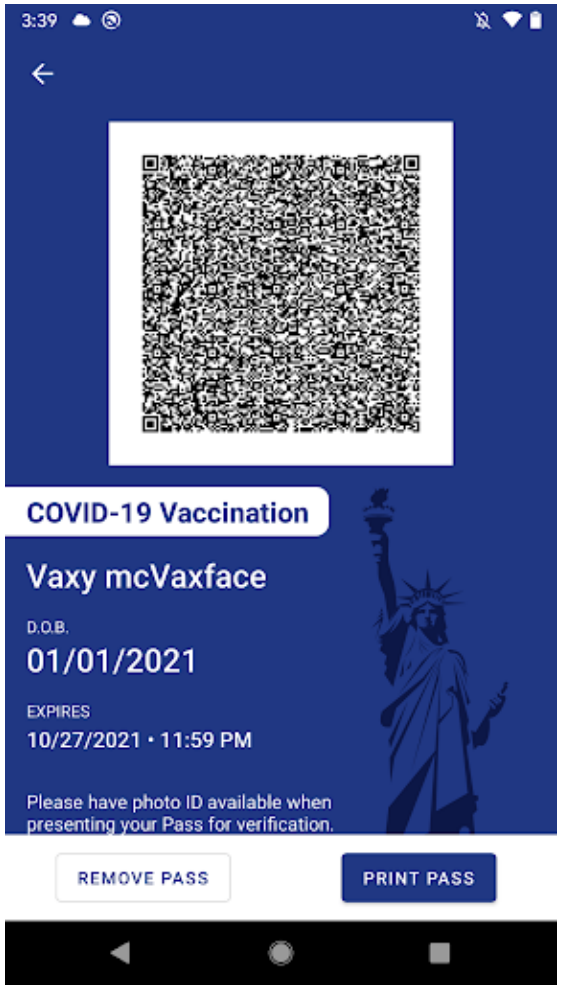

Algorithmic Bias & Racial Injustice

- In August 2021, Tennisha Martin presented How Bias and Discrimination will have members of the Black community Incarcerated or Dead at the DEF CON 29 Blacks In Cyber Village. This presentation focussed on the use of machine learning and other automated/algorithmic systems in the healthcare and criminal justice settings, and how the aggregation of data and the formulation of algorithms by a largely homogeneous population results in bias and discrimination against people of colour. She highlighted examples of real-world harms including the racial impact of predictive policing on the jail times served by people of different races, and the entrenched bias in medical diagnostic algorithms that more often fails to diagnose and treat certain illnesses in patients of colour, resulting in poorer health outcomes for this group. Overall, she shared the message that the bias and discrimination in many application areas of artificial intelligence will have members of the Black community incarcerated or dead, and that these systems must be interrogated and improved to create a more just and equitable world for all. The video recording of this presentation is available here.

Mobile Privacy for Every User

- In March 2021, Dan Hastings and Emanuel Flores officially released Solitude, an open-source privacy analysis tool. Solitude aims to help people inspect where their private data goes once it leaves their favorite mobile or web applications on iOS or Android. Whether a curious novice or a more advanced researcher, Solitude makes the process of evaluating an app's privacy accessible for everyone. Since privacy policies are often difficult to understand when trying to identify how your private data is being shared and whom it's being shared with - and privacy policies don't always tell the complete truth of what an app's actual data collection practices are - Solitude was built to help give more transparency to users of where their private data goes. Solitude makes the process of proxying HTTP traffic and searching through HTTP traffic more straightforward. Solitude can be configured to look for any type of data that you input in a mobile or web application and reveal where that data is going. The application inspects all outbound HTTP traffic, looks for various hashes of your data and recursively decodes common encoding schemes (base64,URL). In August 2021, Dan presented on Solitude at the Black Hat USA Arsenal, as well as at the USENIX Symposium on Usable Privacy and Security (SOUPS). Demos of this tool have also previously been presented at Chaos Communication Congress, and it has received media coverage by Hackin9, as well as been used in journalistic investigations conducted by reporters from a major U.S. daily paper. This work also resulted thus far in two technical advisories associated with pasteboard data leakage vulnerabilities:
  - Technical Advisory – Shop app sends pasteboard data to Shopify's servers
  - Technical Advisory – ParcelTrack sends all pasteboard data to ParcelTrack's servers on startup
- In June 2021, Nick Galloway published iOS User Enrollment and Trusted Certificates. This research involved a study of MDM (mobile device management) on iOS 13, specifically finding that the User Enrollment MDM option added with iOS 13 does not restrict MDM-deployed certificates to MDM-deployed applications. In practice, what this means is that personally installed (ie: non-work) apps will trust MDM-deployed (ie: employer) certificates, in the absence of additional controls such as certificate pinning. When using User Enrollment on the organization's Wi-Fi, it is possible for a corporate Intrusion Detection System to collect personal data by monitoring intercepted traffic, seriously compromising an individual's privacy on their own devices by virtue of trusting their employer's certificates for all activity on the device.

Defending Good-Faith Security Research

- In June 2021, NCC Group also co-signed the Electronic Frontier Foundation's Statement on DMCA Use Against Security Researchers, alongside over a dozen other security firms. This statement voiced opposition to the use of Section 1201 of the Digital Millennium Copyright Act against security researchers performing research in good faith, including when using third-party security testing tools, where we stated, *"We believe that the security of the internet and our digital world is strengthened by the work of independent security researchers who seek to discover and remediate existing security vulnerabilities before those vulnerabilities can be uncovered and exploited by threat actors, and we urge reconsideration of policy which may inhibit this important work."*

Consumer IoT Device Security

- Throughout the year, in an ongoing campaign to improve the security of consumer-oriented IoT devices in collaboration with independent UK consumer body Which?, Guy Morley and Dale Pavey conducted four research projects related to vulnerabilities and risks associated with different classes of IoT devices including a range of smart home devices, connected dash cams, and e-Scooters:
  - In their research into the security of **smart homes**, conducted in May 2021 in collaboration with the Global Cyber Alliance (GCA) and reported on by Which? in July 2021, How a smart home could be at risk from hackers, Guy and Dale built a smart home honeypot - including a collection of smart TVs, printers and wireless security cameras and Wi-Fi kettles - and detected more than 12,000 scanning or hacking attempts in a single week. They put these findings in context, offering an explanation as to why these attacks occur, and offering actionable advice for users to help better secure their smart homes. This was further discussed in their blog post, Honeypot research reveals the connected life might not be so sweet, where they highlighted both the user risk (such as leaked smart camera streams) of attacks on smart home devices, as well as threat actors' broader desire to use the compromised devices themselves to construct botnets to perform wider, more powerful hacking attempts, helping general users of smart home devices to better understand this well-understood phenomenon.
  - In their research into broad, **ecosystem-wide vulnerabilities** in the IoT and app ecosystem reported on by Which? in November 2021, Online marketplaces flooded with insecure smart products, Guy and Dale studied more than 1,800 individual smart products listed on UK online marketplaces including Amazon, eBay and AliExpress, including smart doorbells, wireless cameras and tablets. To conduct this research, they look at a range of generic and clone smart products and trawled online marketplaces for key words associated with these products. Interestingly, 1,727 of the products found used just four apps – Aiwit, CamHi, CloudEdge and Smart Life - where a number of security issues were identified. Issues including weak password security, unencrypted data transfer, vague privacy policies, lack of means for coordinated disclosure of vulnerabilities with device manufacturers, and some devices being out of support for more than 7 years. Unfortunately, these weren't low-hanging fruit, either: The identified devices had 37,129 reviews on Amazon at an average 4.1 star rating, and 15 of them featured Amazon Choice labels. This work was further discussed in their blog post, Home is where the hack is.
  - Based on the research findings above, Which also produced a video, It's this EASY to hack your smart home, in which Guy Morley joins Which? to demonstrate attacks on a smart doorbell, connected camera, printer, smart lock, and more, demonstrating how one compromised device can be used to more easily compromise other devices - in this case, unlocking the home's door via its' insecure smart lock.
  - In their **dashcam** research reported on by Which? in July 2021, How secure is the data on your dash cam? Guy and Dale studied the security of 9 dash cams – from BlackVue, Garmin, Halfords, Kitvision, MiVue, Nextbase, Road Angel, Transcend and Viofo – resulting in pervasive findings of weak encryption, lack of server protection, and weak passwords, noting that, "If a criminal were able to access the data on a dash cam, they could use it to work out where you live, where you work, what time you usually leave the house and where you go – not to mention being able to delete something incriminating from your recordings." This was further discussed in their blog post, Are dash cam users en-route to security risks?, which also explored the ongoing debate about whether GPS and dash camdata should be considered personal data.
  - In their recent research into **e-Scooters**, Guy and Dale looked at the security and safety of 10 popular brands s-Scooters, and found a number of safety issues in relation to electronic brakes and the ability to remotely tamper with the devices via Bluetooth-based attacks. They also found that a number of them could be modified quite easily to push them beyond legal speed limits. At time of writing, this work is undergoing coordinated disclosure with the affected manufacturers, for publication of findings set for early 2022.
- For the aforementioned research with independent UK consumer body Which? investigating the safety and security of a range of IoT devices, Dale Pavey and Guy Morley were named the winners of the Best Ethical Hacker/Pentester Award at the 2021 Security Serious Unsung Heroes Awards, covered here in Infosecurity Magazine.

Cloud & CI/CD Pipeline Security

- This year, Viktor Gazdag joined to the Center for Internet Security community and made technical contributions to the CIS Benchmarks for Securing Microsoft 365. The benchmark has two levels of checks and recommendations that provides guidance for creating and configuring a baseline security in their Microsoft 365 tenant. At time of writing, the CIS Microsoft 365 Foundations Benchmark was in version 1.3.0, to which Viktor has contibuted.
- In March 2021, Erik Steringer published version 1.1.0 of his tool, Principal Mapper (PMapper), a tool for quickly evaluating IAM permissions in AWS.
- In May 2021, Xavier Garceau-Aranda gave a training at NorthSec (which was offered again by NCC Group in August at Black Hat USA 2021) titled Offensive Cloud Security. This training allowed attendees to experience first-hand how security vectors that exist in ecosystems where conventional technologies integrate with cloud-based solutions present opportunities for abuse by attackers, as well as the detection and mitigation of these attacks. The training was structured as a sequence of scenarios, which mix theory and practical exercises in multi-cloud environments, exploring topics like leveraging CI/CD systems to gain a foothold into cloud environments, lateral movement & privilege escalation, abusing containers & clusters, hybrid networks & moving from the management plane to the resources plane, Azure AD synchronization mechanisms and pitfalls, and more. Ultimately, while this training took an offensive perspective on cloud security, it aimed to show how to defend and mitigate against a range of attacks against public and hybrid cloud systems.
- In June 2021, Jerome Smith published a blog post titled "Are you oversharing (in Salesforce)? Our new tool could sniff it out!" This post introduces Jerome's new open-source tool, Raccoon, which aims to identify potential misconfigurations that could expose sensitive data within Salesforce, revealing where access has been granted to all records for particular objects of interest. This helps combat the "complex relationship between role hierarchies, user permissions, sharing rules, and exceptions for certain situations" within Salesforce, helping to mitigate the common concern about potential unauthorized access to data among clients who commission security assessments on their instancesof Salesforce.
- In August 2021, Tim Rawlins discussed the technical debt incurred by a rapid shift to cloud to support remote working during the COVID-19 pandemic with Communications of the ACM.
- In August 2021, Erik Steringer presented his open source tool, Principal Mapper (PMapper), at Black Hat USA Arsenal and DEF CON 29 Demo Labs. Principal Mapper is a script and library for identifying risks in the configuration of AWS Identity and Access Management (IAM) for an AWS account or an AWS organization. It models the different IAM Users and Roles in an account as a directed graph, which enables checks for privilege escalation and for alternate paths an attacker could take to gain access to a resource or action in AWS.
- In September 2021, Erik Steringer presented Automating AWS Privilege Escalation Risk Detection With Principal Mapper at fwd:cloudsec at the Marriott City Center in Salt Lake City, in which he discussed how to use his open source tool, Principal Mapper (PMapper), for in-depth evaluation of AWS IAM Authorization Risks, as well as how to extend it to automate finding risks (continuous monitoring) and test for resource isolation. The video of this presentation is available here.
- In September 2021, Dirk-Jan Mollema presented Breaking Azure AD joined endpoints in zero-trust environments at RomHack 2021. In this presentation, he asked, *"how much trust is zero trust anyway?"* and discussed how as more security controls are added to protect cloud accounts, much of that trust is concentrated at a user's endpoint, where long-term credentials are stored which comply with strict security policies, such as the use of 2FA. However, he noted that while hardware protection and use of Trusted Platform Modules (TPMs) would ideally offer high assurance of these credentials' security, that in practice, he has been able through the past year of research to break a number of security controls pertaining to Azure AD device security, and demonstrated in his presentation both how this can be done, as well as what the consequences are of those attacks. The video for this presentation is available here.

RIFT, Threat Intelligence, CIRT & Honeypots

NCC Group's Research & Intelligence Fusion Team (RIFT) leverages our strategic analysis, data science, and threat hunting capabilities to create actionable threat intelligence, ranging from indicators of compromise and detection capabilities to strategic reports on tomorrow's threat landscape. To ensure that our managed services remain effective against the latest threats, NCC Group operates a Global Fusion Center - a multidisciplinary team that converts our cyber threat intelligence into powerful detection strategies. The work described in this section includes the work of RIFT, as well as complementary research and insights from NCC Group's Global Incident Response/CIRT teams.

- In January 2021, Ollie Whitehouse wanted to build a mechanism to capture all the passwords used (successful or not) against RDP in order to provide blue teams with a source of high-signal intelligence around potential credential compromises. His blog post, Building an RDP Credential Catcher for Threat Intelligence, provides the background on an approach and the steps to build such a system. This post described how to build an RDP credential catcher for threat intelligence, with the initial caveat that he had to disable NLA (Network Level Authentication - an authentication method where a user is initially authenticated via a hash of their credentials over RDP) in order to receive the password in cleartext. However, RDP clients of course may be attempting to connect with NLA enabled, which is a stronger form of authentication that does not send passwords in cleartext. In October 2021, this research was extended by Ray Lai and Ollie Whitehouse, as discussed in their blog post Cracking RDP NLA Supplied Credentials for Threat Intelligence. In that second post, they discuss their work in capturing and cracking the hashed passwords being sent over NLA connections to ascertain those supplied by threat actors, and release the associated source code.

- In January 2021, Wouter Jansen published Abusing cloud services to fly under the radar, to provide the wider defensive infosec community with information, intelligence, and historic data that can be used to hunt for the threat actor threat in historic data and improve detections for intrusions by this intrusion set. The APT group, known as "Chimera," has been best known for abusing cloud services from Google and Microsoft to seek everything from intellectual property (IP) from victims in the semiconductor industry, through to data from the airline industry.

- On January 23rd 2021, NCC Group's Research and Intelligence Fusion Team (RIFT) published a blog post titled Analysing a Lazarus Shellcode Execution Method. This post begins with a discussion of techniques used by Lazarus Group used to execute shellcode from VBA (initially identified in a malware sample shared by Check Point Research two days prior), as well as the known modus operandi of the threat actor of phishing using macro documents disguised as job descriptions in LinkedIn (as documented in research by ESET). In their blog post, NCC Group's RIFT extended this research by analyzing sample macro-enabled documents, and pivoting on the macro, to identify a number of other similar documents for further analysis. In these documents they came across a new technique being used to execute shellcode from VBA without the use of common "suspicious" APIs, whereby they abuse (otherwise-presumed) benign features of the Windows API to achieve code execution. This work was covered by Hacki9.

- In March 2021, Rich Warren and Sander Laarhoven published a blog post titled RIFT: Detection capabilities for recent F5 BIG-IP/BIG-IQ iControl REST API vulnerabilities CVE-2021-22986. In this post, they discuss the wild exploitation attempts and detection logic for the F5 BIG-IP/BIG-IQ iControl REST API vulnerabilities CVE-2021-22986.

- On April 23rd 2021, our Research and Intelligence Fusion Team (RIFT) published some statistics around deployment of Pulse Connect Secure versions in the wild. The hope was that by releasing these statistics, RIFT could help to highlight the risk to enterprises around outdated versions of Pulse Connect Secure, which are being actively exploited by state-level threat actors. This message further amplifies the concerns shared in this joint advisory from CISA, the NSA, and the FBI from April 15th, which discussed CVE-2019-11510, in addition to subsequently-disclosed vulnerabilities including CVE-2021-22893.

- In April 2021, Michael Gough presented "Incident Response Fails – What we see with our clients, and their fails, preparation will save you a ton of $$$, heartache, maybe your sanity and job," and Aaron Crawford presented, "Exploring the Hacker Mentality for Positive Solutions," both at BSides Oklahoma. In June 2021, Michael Gough also presented his Incident Response Fails talk at BSides SATX.

- In May 2021, our Research and Intelligence Fusion Team (including fumik0) published RM3 – Curiosities of the wildest banking malware - a blog post resulting from over 2 years of their work tracking the Gozi malware variant, RM3. In this post, they provided detailed history, analysis and observations on the most pernicious family of banking malware targeting Oceania, the UK, Germany and Italy, giving an overview of Gozi's origins and current operations, and then offering a deep dive technical analysis of the RM3 variant, which RIFT has observed to have targeted at least 130 financial institutions.

- In May 2021, Aaron Greetham published Detecting Rclone – An Effective Tool for Exfiltration. This post discusses NCC Group's Cyber Incident Response Team's (CIRT) experience responding to a large number of ransomware cases where the open source tool Rclone is being used for data exfiltration by threat actors. In this post, he builds on the defensive and threat intelligence work by others and provides additional methods of detection, including Sigma rules to assist with hunting in your own environment.

- In June 2021, Michael Matthews and William Backhouse published a blog post, Handy guide to a new Fivehands ransomware variant, describing their observation of a new variant of the FiveHands ransomware, deployed by a threat actor matching several characteristics shown through the campaign which suggested a link to UNC2447. This blog post aims to help defenders through offering a description of the developments in the ransomware variant as well as the techniques and attack toolkit used by the ransomware operator. This research was covered by ITPro Today and ThisWeekin4n6.

- In July 2021, Michael Gough published a blog post titled Detecting and Hunting for the Malicious NetFilter Driver. This work discussed recent allegations by security researchers from G Data that a driver for Microsoft Windows named "netfilter.sys" had a backdoor added by a 3rd party that was subsequently signed by Microsoft as a part of the Microsoft OEM program. This post offered details (including IoCs) about the malicious driver, the attack, and the post-exploitation process to help defenders with threat hunting and detection. It was covered by ThisWeekIn4n6.

- In August 2021, Sanne Maasakkers presented Phish like an APT: Phenomenal pretexting for persuasive phishing at the DEF CON Adversary Village and later in the year, at Ekoparty. In this talk, Sanne presented an analysis of hundreds of phishing emails that were used in real campaigns. All characteristics of an email, like the method of influence, tone of speech and used technologies are classified. By comparing and measuring the state of these phishing emails, she shared how we can learn more about how certain groups operate, and critically, how real-world APT phishing strategies differ from red team exercises. You can watch a recording of her DEF CON version of the presentation here.

- In August 2021, Michael Gough wrote a blog post on the importance of disabling office macros, where he offers methods for detection and prevention of malicious macros.

- In September 2021, Michael Gough published a blog post titled Detecting and Hunting for the PetitPotam NTLM Relay Attack. This post discussed the proof of concept tool named "PetitPotam" released by security researchers a few weeks earlier, which exploits a flaw in Microsoft Windows Active Directory Certificate Servers with an NTLM relay attack. The flaw allows an attacker to gain administrative privileges of an Active Directory Certificate Server once on the network with another exploit or malware infecting a system. This post provided details to assist organizations in detecting and threat hunting for this and other similar types of threats.

- In October 2021, the Research and Intelligence Fusion Team published SnapMC skips ransomware, steals data, which discussed NCC Group's observation in late 2021 of an increase in so-called data breach extortion cases, whereby an attacker steals data and threatens to publish said data online if the victim decides not to pay, all without use of ransomware. Specifically, RIFT has observed an adversary with a consistent pattern of behaviours whom they track as "SnapMC" - so-named because of both the speed of the actor's attacks, which are generally completed in under 30 minutes, as well as the exfiltration tool mc.exe it uses. This adversary has not yet been linked to any known threat actors. This post details both the psychological tactics used on victims of SnapMC, as well as techniques used and potential mitigations.

- In October 2021, Michael Gough published Detecting and Protecting when Remote Desktop Protocol (RDP) is open to the Internet, in which he explains the dangers of exposing RDP to the internet (including that 42 percent of ransomware cases in Q2 2021 leveraged RDP compromise as an attack vector), and provides details to assist organizations in detecting, threat hunting, and reducing malicious RDP attempts.

- In November 2021, RIFT published TA505 exploits SolarWinds Serv-U vulnerability (CVE-2021-35211) for initial access, in which they discuss how an observed increase in Clop ransomware victims enabled them to trace the surge back to a vulnerability in SolarWinds Serv-U that is being abused by the TA505 threat actor, known for extortion attacks using the Clop ransomware. RIFT believes that exploiting such vulnerabilities is a recent initial access technique for TA505, deviating from the actor's usual phishing-based approach, and published this information both as a call to action for organisations using SolarWinds Serv-U software, as well as to inform incident responders currently dealing with Clop ransomware.

- In November 2021, Pepijn Hack & Zong-Yu Wu presented "We Wait, Because We Know You" - Inside the Ransomware Negotiation Economics at Black Hat Europe 2021. This session explored three main topics. First, can we explain how adversaries use economic models to maximize their profits? Second, what does this tell us about the position of the victim during the negotiation phase? And third, what strategies can ransomware victims leverage to even the playing field? To answer these questions, the researchers analyzed over seven hundred attacker-victim negotiations between 2019 and 2020. Also at BHEU, Pepjin Hack (NCC Group), Kelly Jackson Higgins (Dark Reading), & Rik Turner (Omdia) hosted a panel titled Ransomware as the New Normal, in which they discuss why ransomware actors are hard to stop, and what organizations can do to improve their defenses against these debilitating attacks. Later that month, they published, "We wait, because we know you." Inside the ransomware negotiation economics, which complemented their BHEU presentation, covering both the economics of ransomware attacks, as well as negotiation strategies to use with threat actors when a ransom is being demanded.

- In December 2021, Michael Gough presented a talk titled ARTHIR: ATT&CK Remote Threat Hunting Incident Response Windows Tool at the Open Source Digital Forensics Conference. In this talk Michael discussed ArTHIR, a modular framework that can be used remotely against one, or many target systems to perform threat hunting, incident response, compromise assessments, configuration, containment, and any other activities one can conjure up utilizing built-in PowerShell (any version) and Windows Remote Management (WinRM), and enables you to map your threat hunting and incident response modules to the MITRE ATT&CK Framework.

- In December 2021, RIFT continued their exposition on threat actor TA505 in the blog post, Tracking a P2P network related to TA505, written by Nikolaos Pantazopoulos and Michael Sandee. As outlined in some of the items above, for much of 2021 RIFT had been closely tracking the operations of TA505 and the development of their various projects (e.g. Clop). During the research in question, the researchers encountered a number of binary files that we have attributed to the developer(s) of 'Grace' (i.e. FlawedGrace). These included a remote administration tool (RAT) used exclusively by the thread actor, TA505. The identified binary files are capable of communicating with each other through a peer-to-peer (P2P) network via UDP. While there does not appear to be a direct interaction between the identified samples and a host infected by 'Grace', they believe with medium to high confidence that there is a connection to the developer(s) of 'Grace' and the identified binaries. In this post, they offered a history of TA505 beginning as early as 2014, and offered technical analysis of the downloader, signed driver, and node tool aspects of the execution chain.

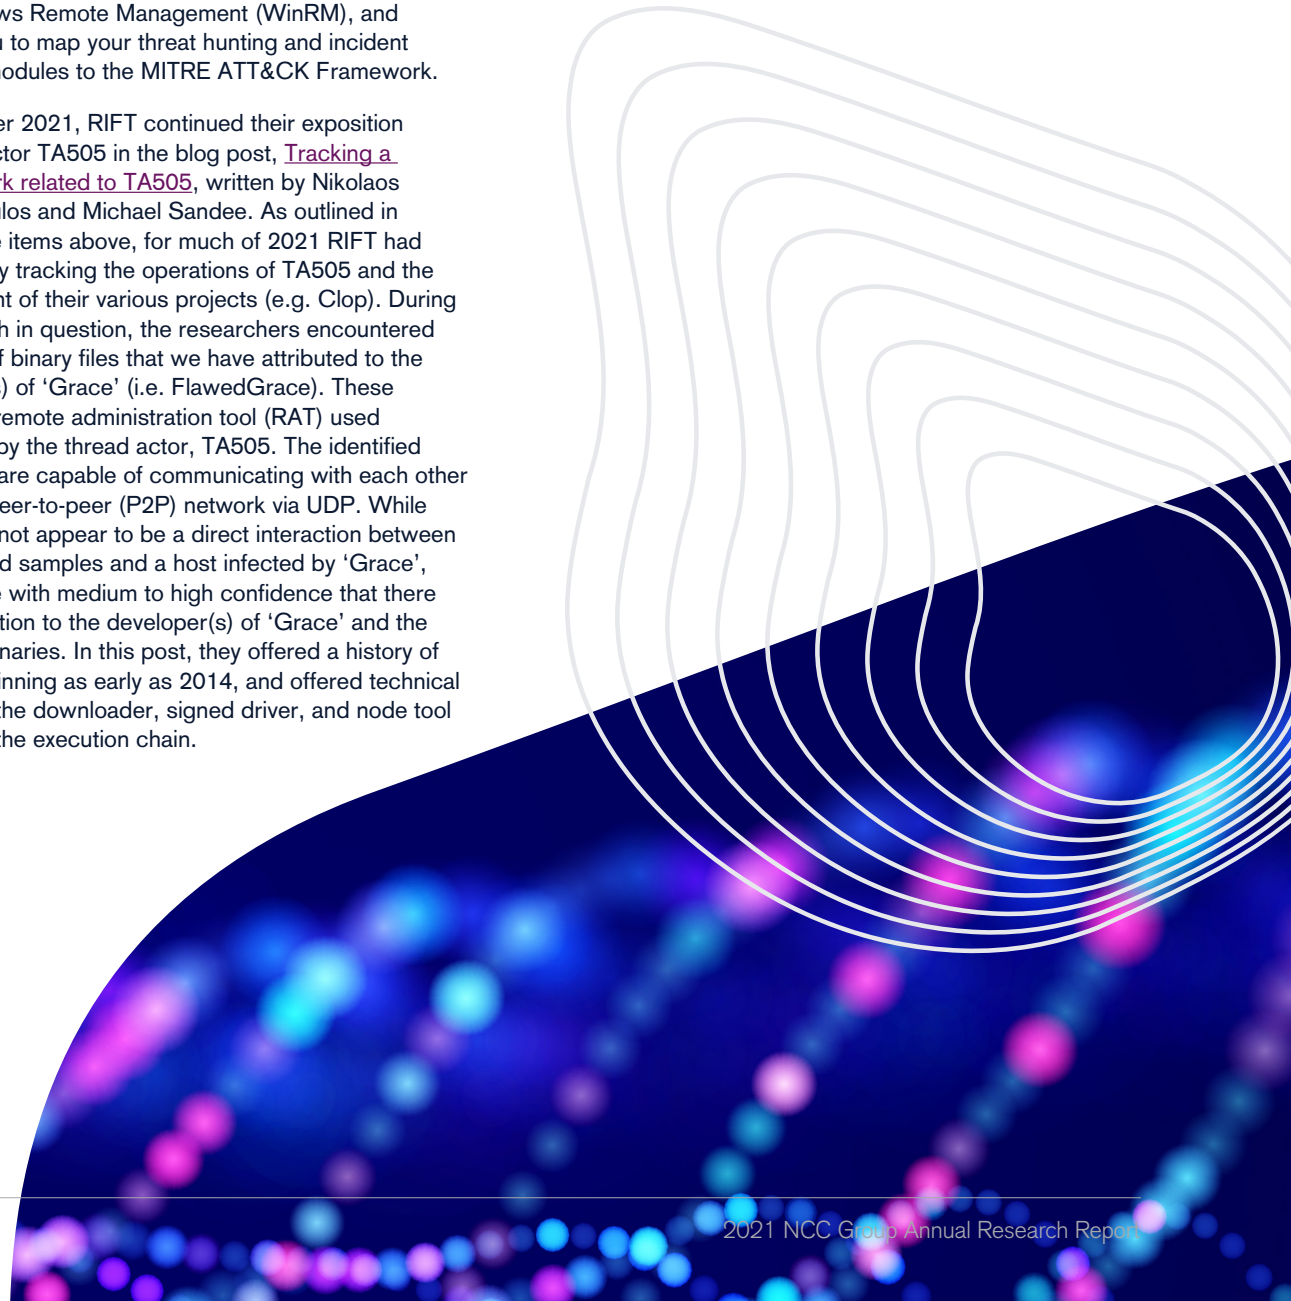

Managed Detection & Response (MDR)

- In January 2021, Liam Stevenson published Using AWS and Azure for Cost Effective Log Ingestion with Data Processing Pipelines for SIEMs. In this post, he showed how to use smart log ingestion via data pre-processing pipelines and modern cloud services to enable SOCs to only store what they really need, ultimately enabling SOCs to reduce data volumes to the SIEM (and associated cost) without losing the residual value and accessibility of the underlying data.

- In March 2021, Group CTO Ollie Whitehouse published a two-part series on Deception Engineering. In part 1, Deception Engineering: exploring the use of Windows Service Canaries against ransomware, in which he prototyped a Windows Service Canary to target parts of the ransomware kill chain to minimize impact and overall success of ransomware operations, based on the idea that it is possible to intercept a ransomware operation prior to files being encrypted, inspired by study of Ryuk ransomware tradecraft. To do this, one installs multiple instances of the canaries which masquerade as common Windows services that are targeted by threat actors prior to encryption. If multiple instances of these services are stopped, then a Canary token is triggered and the host hibernates, alerting us of the presence of ransomware operators, allowing defenders to minimize the impact and give the best possible chance of recovery. In part 2 of this series, Deception Engineering: exploring the use of Windows Installer Packages against first stage payloads, he built upon the previous post, this time focussing on the concept of targeting tradecraft, specifically that used by the cryptomining actor known as LemonDuck who exploited the Microsoft Exchange vulnerabilities in 2021. The thesis here was that, given that a number of packages are uninstalled as a part of infection, there exists the opportunity for detection of malicious actors with a very high signal. To do this, we can deploy a number of Canary Windows Installation Packages with various names that are likely to be targeted. If these packages are uninstalled we can use custom actions to fire a canary token to provide the alert.

- In June 2021, the NCC Group & Fox-IT Data Science Team published Incremental Machine Learning by Example: Detecting Suspicious Activity with Zeek Data Streams, River, and JA3 Hashes. This post, which explores the incremental learning machine learning paradigm in a real-world setting, provides a simple example involving JA3 hashes showing how some of the foundational algorithms that enable incremental learning techniques can be applied to novelty detection (the first time something has happened) and outlier detection (rare events) on data streams derived from Zeek.

- In September 2021, Peter Scopes blogged about CertPortal which allows users to create and manage S/MIME certificates, automating certificate registration and renewal to allow enterprise scale deployment.

- In October 2021, Philipp Schaefer published a blog post describing some of NCC Group's new Managed Detection and Response services in Azure, titled Enterprise-scale seamless onboarding and deployment of Azure Sentinel using Lighthouse for multi-tenant environments. This gave a behind-the-scenes view of some of the technology we use in this service, such as the automated processes involved in setting up new Azure Sentinel environments and managing custom analytics for each customer, including details about our scripting and automated build and release pipelines, which are deployed as infrastructure-as-code.

- In November 2021, Ben Lister and Kane Ryans published Detection Engineering for Kubernetes clusters. This blog post detailed the collaboration between NCC Group's Detection Engineering team and Containerization & Orchestration practice in tackling detection engineering for Kubernetes, including both a description of the detection engineering team's more generic methodology around detection engineering for new/emerging technologies and how it was used when developing detections for Kubernetes-based attacks, as well as detailing the novel detection rules they have created around how privilege escalation is achieved within a Kubernetes cluster, to better enable security operations teams to monitor security-related events on Kubernetes clusters and thus to help defend them in real-world use.

- In December 2021, Margit Hazenbroek published Encryption Does Not Equal Invisibility – Detecting Anomalous TLS Certificates with the Half-Space-Trees Algorithm. This blog post outlined an approach to detecting suspicious TLS certificates using an incremental anomaly detection model. Specifically, this model utilized the Half-Space-Trees algorithm and provides NCC Group's Managed Detection and Response Security Operations Center (SOC) with the opportunity to detect suspicious behavior, in real-time, even when network traffic is encrypted.

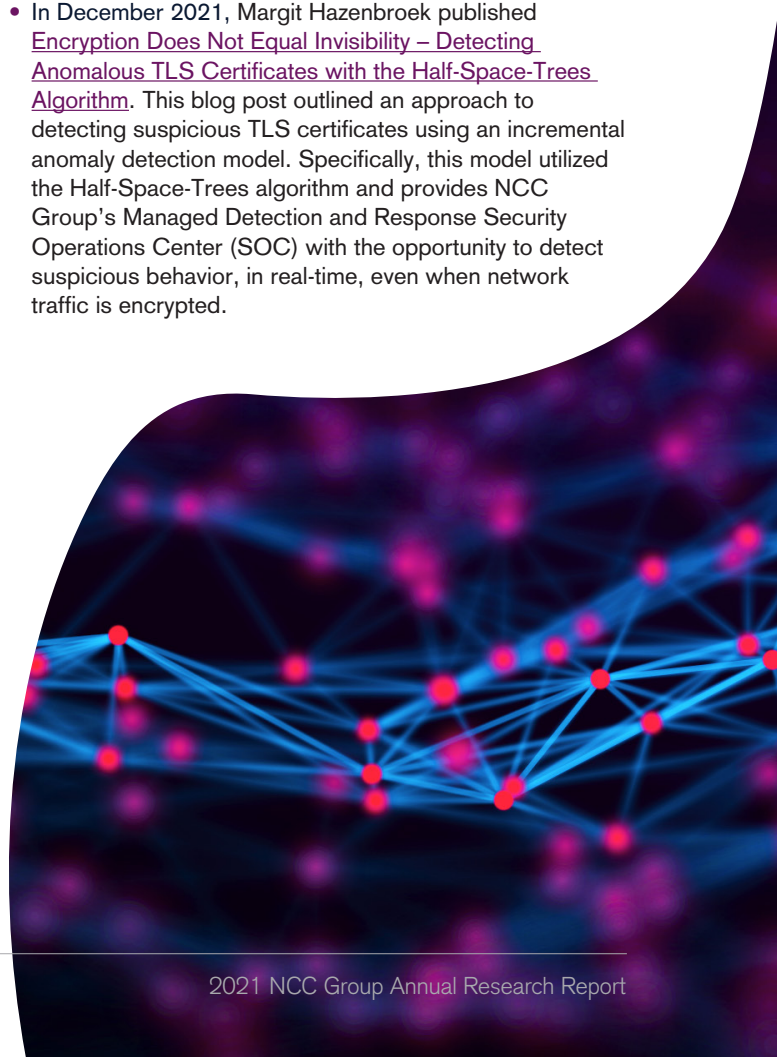

Exploit Development Group

NCC Group's Exploit Development Group (EDG) is a small team of full-time exploit developers who write custom exploits exclusively for the purpose of helping our clients test their own infrastructure and systems against real-world attacks of contemporary vulnerabilities and exploits in the wild, to better understand their risk and resilience. This team reports directly into Group CTO, Ollie Whitehouse. Sometimes, this team presents some of their research externally, and occasionally will speak publicly about consensual, proof-of-concept exploitation of vulnerabilities on our clients' infrastructure, such as in our previous discussion of how in 2017, we unleashed our version of NotPetya on global commodities trading firm, Trafigura.

- In March 2021, Cedric Halbronn published Wubes: Leveraging the Windows 10 Sandbox for Arbitrary Processes. Through this post, he released his new Wubes, which offers Qubes-like containerization but for Microsoft Windows. The idea is to leverage the Windows Sandbox technology to spawn applications in isolation. The initial release of Wubes supported spawning a Windows Sandbox for the Firefox browser but other applications could be easily added The rationale for creating Wubes was that if you browse a malicious site using Wubes, it won't be able to infect your Windows host without additional chained exploits. Specifically, this means attackers need 1, 2, 3 and 4 below instead of just 1 and 2 in the case of Firefox:

  1. Browser remote code execution (RCE) exploit
  2. Local privilege exploit (LPE)
  3. Bypass of Code Integrity (CI)
  4. HyperV (HV) elevation of privilege (EoP)

- In June 2021, Cedric Halbronn published Exploit mitigations: Keeping up with evolving and complex software/hardware, a knowledge base of exploit mitigations available across numerous operating systems (Windows, Linux, Android, iOS, OpenBSD, FreeBSD), architectures (ARM) and applications and versions, including the glibc library, Mozilla Firefox, Microsoft Edge, Google Chrome, and Microsoft Office.

- In June 2021, Cedric Halbronn published Exploiting the Sudo Baron Samedit vulnerability (CVE-2021-3156) on VMWare vCenter Server 7.0. This post detailed the technique of abusing `defaults` structures to exploit CVE-2021-3156 (which was originally made public by researcher Worawit, but not fully explained), and allows privilege elevation from the regular user `vsphere-ui` to `root`. As a part of this post, Cedric also released a new version of the libptmalloc tool, which is heap analysis tooling for ptmalloc (pthreads malloc), and is interesting to those seeking to exploit glibc. In this post he also made public an updated version of the exploit that is more robust and works on vCenter Server.

- Beginning in July 2021, Alex Plaskett published a series of blog posts on CVE-2021-31956, Exploiting the Windows Kernel (NTFS with WNF). These posts looked at CVE-2021-31956 (NTFS Paged Pool Memory corruption), a local privilege escalation within Windows due to a kernel memory corruption bug which was patched within the June 2021 Patch Tuesday. He attempted to exploit this vulnerability on Windows 10 20H2 to determine the ease of exploitation, and to understand the exploit mitigation challenges attackers face when writing a modern kernel pool exploits for Windows 10 20H2 and onwards. The first blog post in the series described the vulnerability, the initial constraints from an exploit development perspective and finally how Windows Notification Framework can be abused to obtain a number of exploit primitives. The second blog post in the series, CVE-2021-31956 Exploiting the Windows Kernel (NTFS with WNF) – Part 2, described improvements which can be made to an exploit to enhance reliability, stability and clean-up afterwards, including exploitation without the CVE-2021-31955 information disclosure, enabling better exploit primitives through `PreviousMode`, and some thoughts on detection.

- In November 2021, Alex Plaskett presented Pwning the Windows 10 Kernel with NTFS and WNF at POC (Power of Community) 2021, based on his earlier blog posts. His slides for that presentation are available here.

- In November 2021, NCC Group's Exploit Development Group (Aaron Adams, Cedric Halbronn, and Alex Plaskett) participated in Pwn2Own Austin 2021, where they successfully exploited the Lexmark MC3224i printer with a file write bug, as well as gaining code execution on the Western Digital PR4100 NAS. This work was supported by Catalin Visinescu to retrieve firmware from one of the devices and Matt Lewis and Matt Trueman with equipment procurement and logistics support.

Other Research & Speaking

- In January 2021, Manuel Ginés Rodríguez & Diego Gómez Marañón published Technical Advisory – Linksys WRT160NL – Authenticated Command Injection (CVE-2021-25310). Successful exploitation of this vulnerability on the Linksys WRT160NL switch can lead to remote code execution on the affected device.
- In February 2021, Manuel Ginés Rodríguez published Technical Advisory - Administrative Passcode Recovery and Authenticated Remote Buffer Overflow Vulnerabilities in Gigaset DX600A Handset (CVE-2021-25309, CVE-2021-25306), where he shared that there were two vulnerabilities that allowed unauthenticated users to retrieve the administrative password for the Gigaset DX600A (a high-end ISDN desktop phone) due to a weak authentication mechanism or compromise its availability through low traffic Denial of Service attacks, which could result in as dramatic an impact as the attacker being able to route the device's traffic through an attacker's controlled machine.
- In March 2021, Adam Roberts published a blog post describing SAML XML Injection. The blog post described a novel class of vulnerability that was detected in several SSO services assessed by NCC Group, specifically affecting Security Assertion Markup Language (SAML) implementations. The flaw could allow an attacker to modify SAML responses generated by an Identity Provider, and thereby gain unauthorized access to arbitrary user accounts, or to escalate privileges within an application.
- In March 2021, Manuel Ginés Rodríguez published advisories, Technical Advisory – Multiple Vulnerabilities in Netgear ProSAFE Plus JGS516PE / GS116Ev2 Switches, for 15 vulnerabilities that he found in the Netgear ProSAFE Plus JGS516PE / GS116Ev2 switches, the most critical of which could allow unauthenticated users to gain arbitrary code execution.
- In March 2021, Richard Warren published a Technical Advisory - Dell SupportAssist Local Privilege Escalation (CVE-2021-21518) to disclose a vulnerability he found that when running PC-Doctor modules, the Dell SupportAssist service attempted to load DLLs from a world-writable directory. Furthermore, it did not validate the signature of libraries loaded from this directory, leading to a "DLL Hijacking" vulnerability, which would allow a low privileged user to execute arbitrary code with system privileges.
- In May 2021, Tanner Prynn discussed how using UUIDs for authorization is dangerous, even if they're cryptographically random.
- In June 2021, Dirk-jan Mollema presented "Walking Your Dog In Multiple Forests - Breaking AD Trust Boundaries Through Kerberos Vulnerabilities" as a Black Hat Webcast. This presentation began with Dirk-Jan describing how Kerberos works over Active Directory forest trusts and how the security boundary is normally enforced, to enable discussion of a flaw he had found in how AD forest trusts operate, which can be combined with a vulnerability in the Windows implementation of Kerberos to take over systems in a different forest (from a compromised trusted forest), demonstrated as a proof-of-concept as a part of the talk.
- In June 2021, aschmitz published Testing Two-Factor Authentication, in which they provided a whirlwind tour of common 2FA mechanisms, as well as detailed information on performing effective security testing against those systems.
- In July 2021, Liam Glanfield published Technical Advisory – Sunhillo SureLine Unauthenticated OS Command Injection (CVE-2021-36380), in which he shared his discovery that the Sunhillo SureLine application contained an unauthenticated operating system (OS) command injection vulnerability that allowed an attacker to execute arbitrary commands with root privileges.

- In July 2021, Stephen Tomkinson published Technical Advisory – Arbitrary File Read in Dell Wyse Management Suite (CVE-2021-21586, CVE-2021-21587). This disclosed his finding that an attacker with physical access to Dell's Wyse thin client and its network connection can exploit a Arbitrary File Read vulnerability to gain access to the management interface of the whole thin client estate. The management interface includes features such as resetting BIOS passwords and remotely shadowing terminal screens via VNC.

- In July 2021, Liew Hock Lai published Technical Advisory - Stored and Reflected XSS Vulnerability in Nagios Log Server (CVE-2021-35478,CVE-2021-35479), which could facilitate attackers in executing malicious JavaScript on victim machines such as stealing cookies or redirecting users.

- In July 2021, Derek Stoeckenius published Technical Advisory – ICTFAX 7-4 – Indirect Object Reference, in which he describes an IDOR vulnerability in the internet-based fax program ICTFAX that can allow allows a user of any privilege level to change the password of any other user within the application – including administrators, enabling a low-privilege user to access both administrative functions and user data from arbitrary users within the application.

- On August 8 2021, Richard Warren published Technical Advisory: Pulse Connect Secure – RCE via Uncontrolled Archive Extraction – CVE-2021-22937 (Patch Bypass), disclosing his finding that the Pulse Connect Secure VPN appliance suffers from an uncontrolled archive extraction vulnerability which allows an attacker to overwrite arbitrary files, resulting in Remote Code Execution as root, which is a bypass of the patch for CVE-2020-8260. This was covered in a variety of media outlets including The Hacker News, SecurityWeek, VM Virtual Machine, Security Affairs, TechNadu, TimesNewsExpress, Hackaday, and others.

- In the summer of 2021, NCC Group researchers gave 10 technical presentations at Black Hat USA, and 7 technical presentations at DEF CON 29. This included 2 Black Hat Briefings, 2 Black Hat Arsenal tools, and 6 Black Hat Training sessions, as well as 2 main track DEF CON talks, 3 DEF CON Demo Labs presentations, and 2 DEF CON Village talks. Many of those have been discussed elsewhere in this report, but the remaining ones will be outlined here.

- Among those presentations was a Black Hat USA Training titled "Bad Active Directory (BAD)" created by Dhruv Verma, Michael Roberts, and Xiang Wen Kuan. BAD is a beginner-to-intermediate level training for hacking Windows Active Directory.

- On August 12th 2021, Richard Warren observed threat actors' active attempts to exploit Exchange ProxyShell vulnerabilities, which was picked up by thehackernews, Bleeping Computer, NationalCybersecurity, Theultramods, Archyde, Threatpost, and Redmond Magazine.

- In September 2021, Richard Warren released Yara rules that detect a remote code execution vulnerability in MSHTML (CVE-2021-40444), which depends upon carefully-crafted Microsoft Office documents as an exploitation vector. It was covered in the media by Hurricane Labs Blog and Security Boulevard. Richard also developed a technique for exploiting CVE-2021-40444, which he demoed on Twitter and was subsequently written about on Decipher.

- In September 2021, Ian Robertson & Javed Samuel presented Castles Built on Sand - a pen-testers view on integrations at CornCon. This talk reviewed the foundations of cryptographic vulnerabilities as applicable to open-source software from a penetration tester's perspective over multiple public cryptography audit reports, discussing the past attacks that took advantage of these cryptographic vulnerabilities, and what the consequences were. The talk also examines ways that open-source software has been updated over time to mitigate these cryptography flaws and how successful these mitigations may have been.

- In September 2021, Jesus Olmos published Technical Advisory – Garuda Linux Insecure User Creation (CVE-2021-3784), on a vulnerability he uncovered in Garuda

Linux that would allow local attacker to impersonate a user account while it is being created, installing a backdoor to access that user account at any moment in the future.

- In September 2021, Liyun Li published Technical Advisory - PDFTron JavaScript URLs Allowed in WebViewer UI (CVE-2021-39307), where he shared his finding that the PDFTron WebViewer renders dangerous URLs as hyperlinks in supported documents, including JavaScript URLs, allowing the execution of arbitrary JavaScript code, which could be used to steal a victim's session tokens, log their keystrokes, steal private data, or perform privileged actions in the context of a victim's session.
- In September 2021, Duane Reeves presented Telephony: The Forgotten Network Threat at GSX: Global Security Exchange.
- In October 2021, Jelle Vergeer published Reverse engineering and decrypting CyberArk vault credential files, discussing his discovery that it was possible to reverse engineer the encryption and key generation algorithms and decrypt the encrypted vault password for CyberArks's vault encryption.
- In October 2021, Balaz Bucsay published Technical Advisory – NULL Pointer Derefence in McAfee Drive Encryption (CVE-2021-23893), which demonstrated a privilege escalation vulnerability in a Windows system driver for McAfee Drive Encryption
- In October 2021, Nicolas Guigo published A Look At Some Real-World Obfuscation Techniques, whereby he shared tools and methods for reversing real-world binary obfuscation.
- In the Fall of 2021, Damon Small presented "Beyond the Scan: The Value Proposition of Vulnerability Assessment" at UTINFOSEC. This talk explored how vulnerability assessment can be leveraged "beyond the scan" to provide tangible value to not only the security team, but to the entire business that it supports.
- Throughout 2021, Sourya Biswas presented "Security from Scratch: Reminiscing Being the 2nd Security Employee at a Startup" at both InfoSec World and BSides Oklahoma. He also presented "Cybersecurity is War: Lessons from Historical Conflicts" at Secure360, and "Psychology of the Phish: Leveraging the Seven Principles of Influence" at both ISACA Conference North America as well as at the Great Lakes Security Conference.
- In October 2021, Richard Warren published Technical Advisory – Apple XAR – Arbitrary File Write (CVE-2021-30833), which disclosed a new bug in MacOS that would allow an attacker to create a malicious .xar file, which when extracted by the user would result in files being written to a location of the attacker's choosing, which could be abused to gain Remote Code Execution. It was covered by The Mac Observer.
- In November 2021, Tennisha Martin gave a keynote titled "The Hacker's Guide to Mentorship: Fostering the Diverse Workforce of the Future" at SANS Pentest HackFest.
- In November 2021, Nicolas Bidron found 3 new vulnerabilities in Victure WR1200 WiFi router, one of which resulted in OS command injection, published in Technical Advisory – Multiple Vulnerabilities in Victure WR1200 WiFi Router (CVE-2021-43282, CVE-2021-43283, CVE-2021-43284).
- In December 2021, Anthony Ferrillo published Technical Advisory – Authenticated SQL Injection in SOAP Request in Broadcom CA Network Flow Analysis (CVE-2021-44050), disclosing a vulnerability which would allow an authenticated user to inject SQL into a SOAP request leading to enumeration of the back end database of the Network Flow Analysis web application.
- In December 2021, Rick Veldhoven published Technical Advisory – Lenovo ImController Local Privilege Escalation (CVE-2021-3922, CVE-2021-3969). Among the vulnerabilities found was a bug that could allow an attacker to elevate their privileges to that of the system user from a user that is able to write files to the filesystem.
- In early December, Richard Warren published 6 Technical Advisories for the SonicWall SMA 100 Series applications, including Unauthenticated Arbitrary File Deletion, Unauthenticated Stored XSS, Multiple Unauthenticated Heap-based and Stack-based Buffer Overflows (CVE-2021-20045), Post-Authentication Remote Command Execution (CVE-2021-20044), Heap-Based Buffer Overflow (CVE-2021-20043), and Unauthenticated File Upload Path Traversal (CVE-2021-20040).

Acknowledgements

Research is an intrinsic part of many of our technical security consultants' daily lives, and almost all of the research that you see in this report was delivered by dozens of consultants from NCC Group offices around the world, seconded into Research to work on their passion projects, empowered by a few thousand dedicated research days across the Group, led by just 1-2 full time staff overseeing end-to-end the over 237 independent research projects delivered in 2021. I'm really proud of this, both because it is a radically inclusive model where every interested consultant can take part and grow as a researcher, and also because we are a lean team that accomplishes a tremendous amount of meaningful things and wastes nothing.

My first acknowledgement goes out to all of our consultants who spent part of their time in the Research division this year, without whose talent, curiosity, time, and courage our research simply would not exist.

I would also like to thank Aaron Haymore, our Research Program Coordinator, for his contributions this year to our research program, including help with countless coordinated vulnerability disclosures, administrative and program management support, his assistance in compiling some of this information for this report, and the joy he brings to everything we do.

I am also grateful to our Research Working Group Leads and RDs - each a talented researcher and consultant in their own right - who generously share their time to mentor others, lead technical discussions in our monthly research working group meetings, and even inspire new projects. These leaders include Daniel Romero, Jeff Dileo, Jeremy Boone, Nick Dunn, Richard Appleby, Robert Seacord, Timur Duehr, Viktor Gazdag, and William Groesbeck.

Finally, I would like to thank our US CTO, Dave Goldsmith, our Group CTO, Ollie Whitehouse, and our Commercial Research Director, Matt Lewis, for their friendship and for all of the many ways in which they help to support and advance Research at NCC Group, as well as to thank the artists at MC2 Manchester for the graphic design of this report.

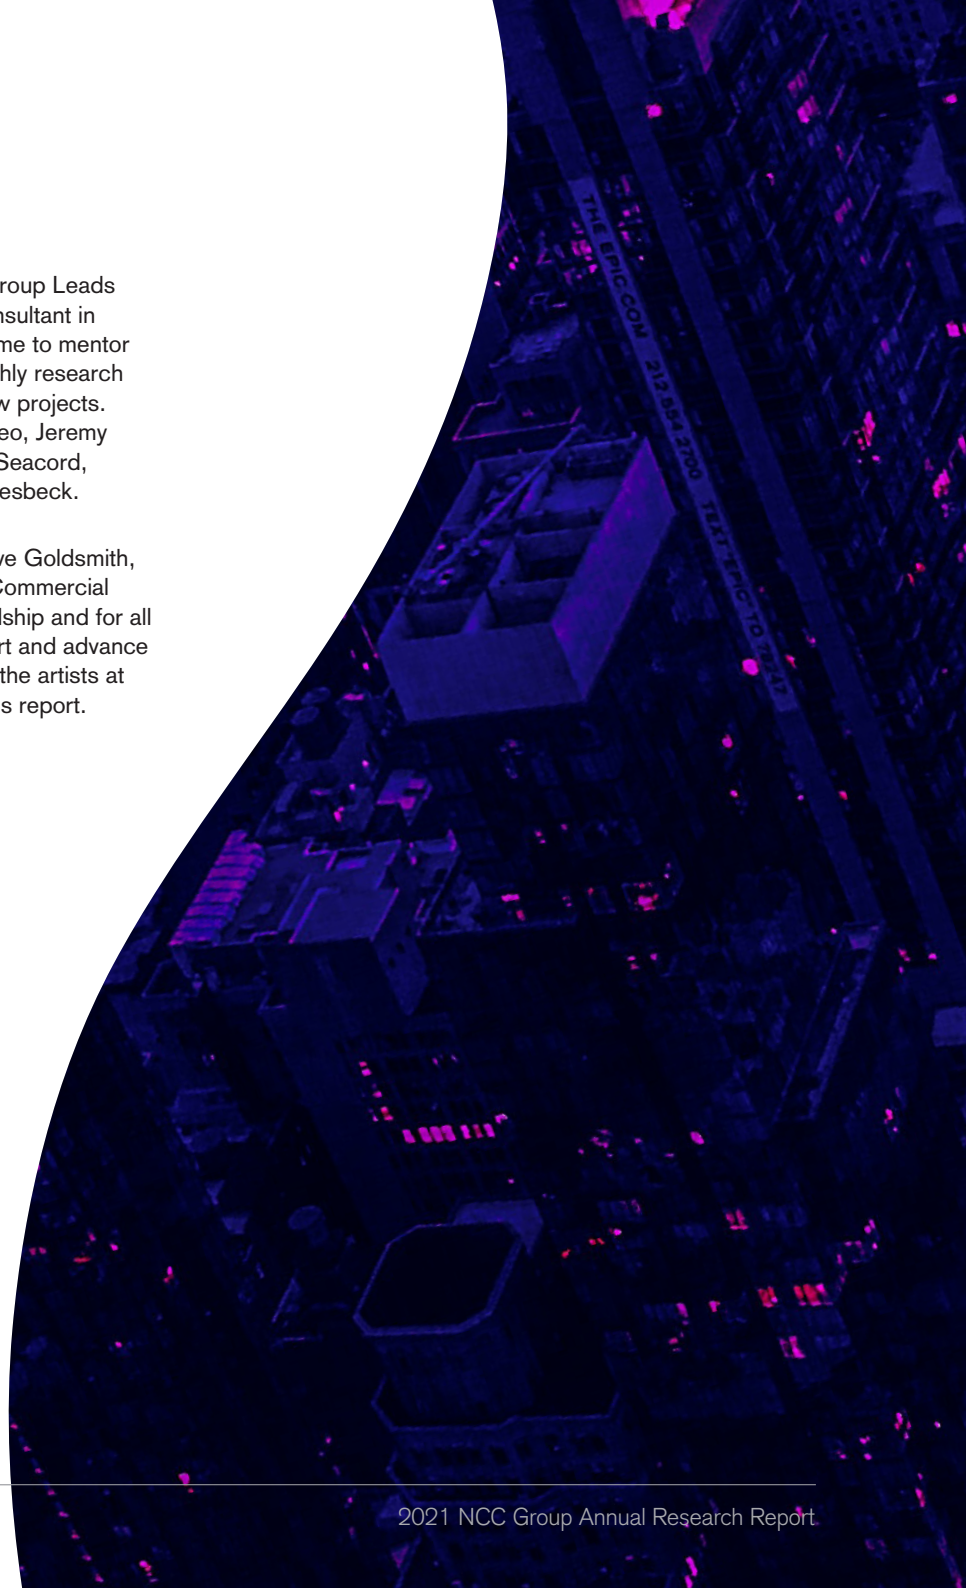

About Research at NCC Group

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. We are a research-driven firm where every researcher on our team is also an active consultant.

With hundreds of specialized consultants, our technical security research areas extend into almost every area of security, as well as global standards bodies including the C Standards Committee and CIS Benchmarks. We perform offensive and defensive research across a vast range of targets including blackbox and whitebox testing of previously unanalyzed 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 virtualization, and both offensive attacks on – and defensive uses of – machine learning and artificial intelligence 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 Blog, alongside our technical Twitter account and our public Github which hosts over 200 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. In 2020, NCC Group was the only security company which co-founded and sit on the Governing Board and Technical Advisory Committee of the Open Source Security Foundation, an industry-wide coalition within the Linux Foundation dedicated to improving the security of the open source ecosystem through a range of strategic projects. 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, Threat Intelligence Fusion Center, and Full Spectrum Attack Simulation group, as well as a number of 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, ACM CCS, Hardwear.io, REcon, Appsec USA, Toorcon, Oracle Code One, BSidesLV, O'Reilly Artificial Intelligence, Chaos Communication Congress, Microsoft BlueHat, HITB Amsterdam, RSA Conference, CanSecWest, USENIX Enigma, the Linux Foundation Member Summit, DEF CON, and countless others.
In recent years, we have served on the review boards of conferences such as IACR Cryptographic Hardware and Embedded Systems (CHES), Black Hat USA, BSidesLV, AppSec Cali, Neural Information Processing Systems (NeurIPS), USENIX Enigma, USENIX CSET, the USENIX Workshop on Offensive Technologies (WOOT) and multiple DEF CON villages. Our work 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 other mainstream and trade publications globally.

research.nccgroup.com
@nccgroupinfosec

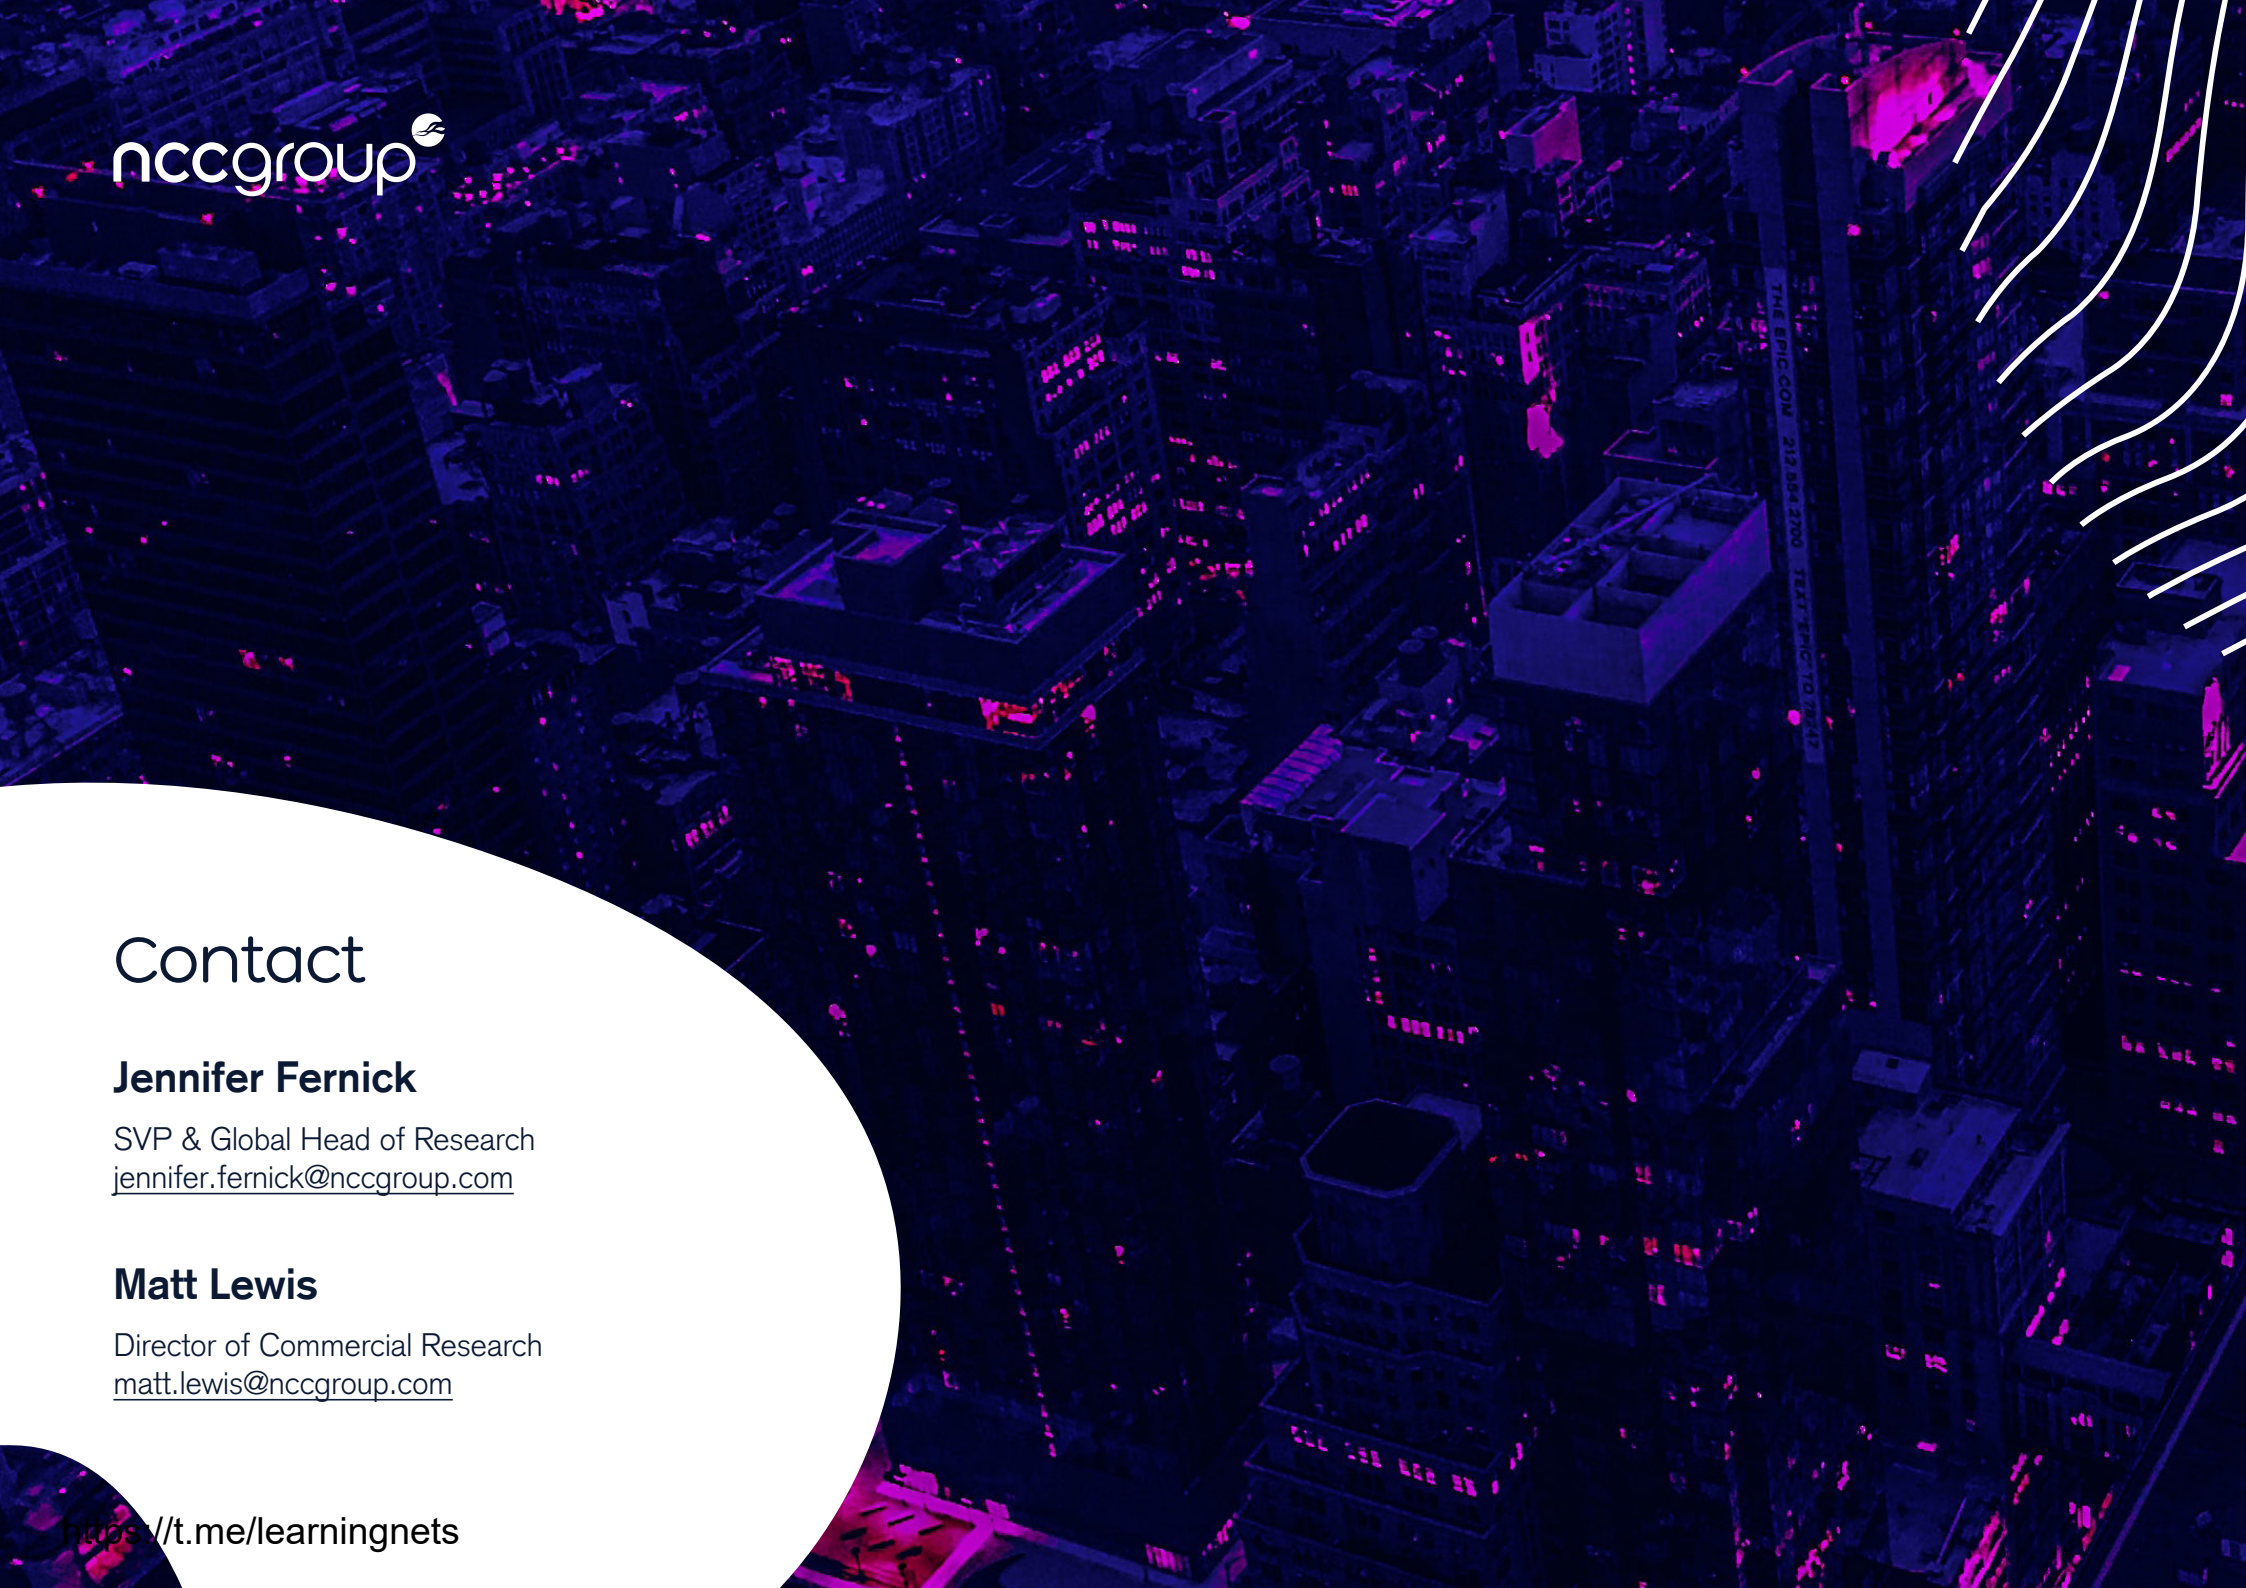

Contact

Jennifer Fernick

SVP & Global Head of Research
jennifer.fernick@nccgroup.com

Matt Lewis

Director of Commercial Research
matt.lewis@nccgroup.com