nccgroup

2020 NCC Group Annual Research Report

Written by Jennifer Fernick

SVP & Global Head of Research, NCC Group

>> research.nccgroup.com

Table of Contents

Message from the Global Head of Research: Security Research in 2020, 2021, and beyond

It has been a year that forced our priorities into sharp focus. In a world stripped of much of the ornament and distraction of daily life, we have been able to see with a rare sharpness the way that informational asymmetries, algorithmic platform-driven influence, and the privilege to manage one's exposure to risk radically change the outcomes of individual human lives and of entire nations.

Into 2021, we will need to take a hard look at the socio-technical decisions we've made in a time when our lives changed more in a week than they might ever have in a lifetime. These questions are bigger than remote work - they are about the entrenched, rapidly-deployed systems (and their accumulating security debt) rolled out at the beginning of the pandemic, which silently and often radically changed the internals of organizations around the globe.

Does remote work change enterprises' exposure to security risk? Certainly somewhat, although this doesn't feel like the security revolution it often seemed reported to be. Few things actually change in this scenario if you had already begun with the fair assumption that most networks are hostile.

We're rapidly approaching a point when companies will no longer be able to patch and insure their way out of security problems - and it is for this reason and the reasons below that I suggest that we work to both to get better at reducing vulnerabilities at scale, while simultaneously seeking to deeply understand where (and why) even the world's best tooling fails to find many of the high-impact vulnerabilities in complex codebases.

What do I see on the horizon for 2021?

- Attacking and defending full-stack artificial intelligence systems in production environments, because after several hundred academic papers and arguably not a single real-world proof of concept, it is time to partner with data-driven companies to literally run the risk before powerful adversaries do so instead.
- As an industry, a movement toward designing for usable security and no longer pushing the burden of security down onto our users only to judge them on their inevitable failures of our own design.
- After rapidly shifting all aspects of our lives online, reclaiming what we can of online privacy through the development & mainstream popularization of privacy-enhancing technologies and their underlying cryptographic protocols.
- A groundswell of policy interest in platform governance and Section 230 of the Communications Decency Act, and the potentially rapid need for related tooling and novel digital forensics capabilities, including in authenticating the origins of images to combat deepfakes and better understand the origins of misinformation efforts.
- The triviality and ubiquity of vulnerabilities in most IoT devices becoming a reason for security practitioners to pay more attention to IoT, particularly in the domain of things that can improve these products at scale, specifically: standards, regulation, security engineering, and law.
- A strive toward increased rigor to the research in our industry, where we would see more of: controlled experiments, well-designed testbeds, increased use of meaningful data science, and adoption of methods from the natural sciences to test hypotheses about the effectiveness of security tooling and risk-remediation efforts.
- A deeper commitment to reducing vulnerabilities at scale, through increased investment in research and tooling, and multi-stakeholder investments in securing core infrastructure to make high-impact reductions in dependency risks downstream where possible.
- As a society, beginning to take seriously the real risk of harm in cyber-physical systems, and people outside of our field more readily understanding that security is a prerequisite to safety, because we cannot make assurances about the safety (and while we're at it: the ethics and privacy) of systems that we do not control.

Jennifer Fernick
SVP & Global Head of Research, NCC Group

Executive Summary: Research at NCC Group (2020)

In 2020, NCC Group delivered thousands of dedicated research days, pursuing research spanning a number of areas, including hardware and embedded systems security, applied cryptography, programming languages, machine learning, mobile privacy, cloud and container security, exploit development, industrial control systems/OT security, threat intelligence, and beyond. We delivered 96 research papers, whitepapers, and technical blog posts, as well as at least 47 conference presentations, in venues including Black Hat USA, Shmoocon, OWASP Global Appsec, KubeCon + CloudNativeCon, Chaos Communication Congress (CCC), Infosec World, Black Hat Asia, RSAC, and DEF CON. We released at least 18 new open source tools, as well as our first-ever open dataset, and released 121 blog posts on research.nccgroup.com, attracting 193,000 visitors.

We also published a 20 year retrospective on research at NCC Group, highlighting some of our favourite projects from the past two decades, from over 200 whitepapers and conference presentations which included key outputs from previous cyber security acquisitions including NGS Software (2008), iSEC Partners Inc. (2010), Matasano Security (2012), Intrepidus Group (2012), FortConsult A/S (2014), Fox-IT (2015), Payment Software Company Inc (2016), and VSR Inc. (2016).

We continue our research partnerships with a variety of institutions including the University College London Centre for Doctoral Training in Data Intensive Science, and our membership within the UK Academic Centers of Excellence in Cyber Security Research via the UK CyberInvest scheme. From a standardization perspective, we continue our long-running active contributions to security-related proposals within the C standards committee, as well as contribute to the development of the Center for Internet Security's CIS Benchmarks for Kubernetes. We have also been active within the NIST post-quantum cryptography standardization process, where members of our Cryptography Services practice have proposed a post-quantum lattice-based signature scheme that was announced by NIST in July 2020 to be a round 3 finalist. We were founding members of the Open Source Security Foundation, a group within the Linux Foundation which "brings together the industry's most important open source security initiatives and the individuals and companies that support them" to help secure the open source ecosystem. We also co-authored with the UK's National Cyber Security Center an Internet Engineering Task Force (IETF) draft.

NCC Group supported academia extensively. We ran the Academic Centre of Excellence in Cyber Security Research (ACE-CSR) virtual conference (June 2020) and the NCSC PhD Winter School Event (December 2020) in partnership with the UK's National Cyber Security Center. NCC Group also sat on the research funding review board for the UK's PETRAS National Centre of Excellence for IoT Systems Cybersecurity.

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. Several NCC Group Researchers (Edward Torkington, Phillip Langlois, and Dirk-Jan Mollema) were recognized among the Microsoft Security Response Center (MSRC)'s Most Valuable Security Researchers in 2020.

Researchers across NCC Group served as leaders, mentors, and peer reviewers across the information security community. This year, we were invited to serve on the program committees and review boards of a number of high-profile computer security research venues including: USENIX Enigma, IEEE Workshop on Information Forensics and Security (WIFS), Black Hat USA, DEF CON Wall of Sheep/Packet Hacking Village, USENIX Workshop On Offensive Technologies (WOOT), and many others, and performed outreach, mentorship, and gave invited technical talks within high schools, colleges, universities and Capture the Flag (CTF) competitions in North America, Europe, the Asia-Pacific region, and around the world.

Applied Cryptography

At the beginning of the year, Thomas Pornin published parameters for a new elliptic curve for use in Elliptic Curve Cryptography, called Curve9767, in a paper titled Efficient Elliptic Curve Operations On Microcontrollers With Finite Field Extensions. In this work, he described a new and efficient elliptic curve with 128-bit security (Curve9767), and released a fully constant-time open source implementation of Curve9767 as well as a later paper describing an optimization of Schnorr signature verification. Curve9767 is particularly useful for constrained devices, and works well on small microcontrollers of the ARM Cortex-M0+ class. A summary of this work is available in his blog post, Curve9767 and Fast Signature Verification.

In January 2020, Jennifer Fernick discussed with BankInfoSecurity vulnerabilities in the Windows CryptoAPI (CVE-2020-0601) recently-disclosed by the U.S. National Security Agency.

In February 2020, Balazs Bucsay wrote about properly signed certificates on CPE devices. This work was inspired by the recent findings of two security researchers that a publicly available Netgear firmware image contains two properly signed and valid TLS certificates that contained private keys, trivially enabling Man-In-The-Middle attacks. Since the vendor's published advisory lacked remediation advice, this blog post suggested multiple potential low-cost mitigation recommendations, described their tradeoffs while noting that they are nevertheless better than the common practice of using self-signed, hard-coded, or no certificates for items of Customer Premises Equipment (CPE), such as consumer routers.

Also that month, Eric Schorn published his whitepaper, A Tour of Curve 25519 in Erlang, which offered an introduction to elliptic curve cryptography theory alongside a practical implementation in Erlang. The implementation code associated with this paper was also published.

In March 2020, Aleksandar Kircanski and Gerald Doussot published Smart Contracts Inside SGX Enclaves: Common Security Bug Patterns. In this work, they discussed nuances to using Intel SGX in practice for the purpose of Ethereum smart contracts, and outlined 9 classes of security issues that we have observed in real-world smart contract implementations.

These issues included: rogue smart contract code injection on TEE nodes; leakage of contract state ciphertext changes; smart contracts using non-constant time cryptographic implementations; the inherent non-constant nature of a smart contract's virtual machine; input, state and output length leaks; data store, operating system file system and database process query processing access pattern side channels; cross-contract calls; potential for inadequate tamper-resistance; and the range of side channel attacks possible against TEEs in general.

They concluded that for the purpose of running confidential smart contracts, using a trusted execution environment (TEE) presents a considerable attack surface to malicious parties with a presence on the host executing the contract and, to a lesser extent, to passive network observers. This was due to factors such as the existence of side channels in smart contract execution, the existence of side channels in data stores running outside of SGX, as well as ciphertext/plaintext length leaks.

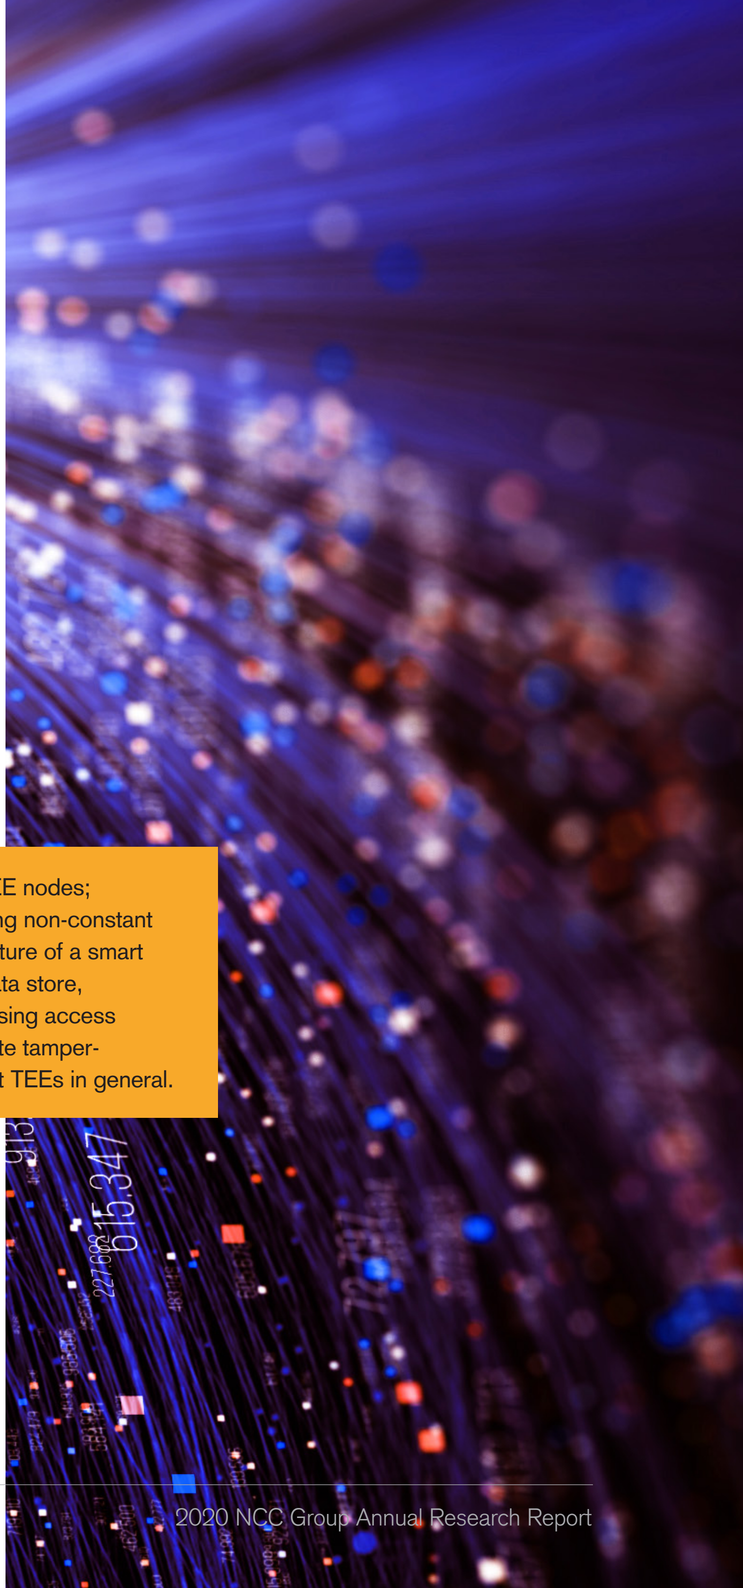

Eric Schorn published his two-part blog series on Verifiable Random Functions (VRFs), popularized due to their usefulness in blockchain applications. In Part 1 of this series he introduced VRFs in the context of other well-known cryptographic primitives, described their use cases, and highlighted over two dozen topics to consider during a cryptographic implementation review of VRFs, including questions regarding ciphersuite configuration, cryptographic key lifecycle, rigor of cryptographic implementation in source code, and code completeness. In Part 2, he published a fully self-contained Python3 reference implementation matching the ECVRF-EDWARDS25519-SHA512-Elligator2 ciphersuite configuration found in the latest IETF CRFG VRF Internet Draft.[1]

In March 2020, Aleksandar Kircanski and Terence Tarvis published their whitepaper, Coinbugs: Enumerating Common Blockchain Implementation-Level Vulnerabilities. To begin, the authors identified a gap in literature regarding implementation-level security bugs that commonly occur in basic proof-of-work blockchain node implementations, and wrote this paper to systematize the bugs discovered by themselves and across the research community during the first decade of Bitcoin's existence. This paper outlines ten broad vulnerability categories and known real-world examples of their existence and exploitation, ultimately tracing Bitcoin's introduction of several novel bug classes which are of interest to security researchers, and setting a baseline set of vulnerability types for which proof-of-work blockchain implementations should be tested.

In April 2020, Paul Bottinelli published a blog post about the cryptographic mechanisms underlying COVID-19 contract-tracing apps, ultimately discussing how cryptography is used to monitor the spread of COVID-19. **(Fig 1)** In this work, he showcased how Apple and Google's proposal for contact-tracing uses cryptographic constructions to trace individuals without sacrificing privacy. For example, he highlighted both that:

1. It should be computationally infeasible, without knowledge of the Daily Tracing Key, to link the Rolling Proximity Identifiers together
2. It should be computationally infeasible to obtain the Tracing Key of a user given a set of their Daily Tracing Keys, as may be disclosed if the user has tested positive

And that these aspects rely on the security of the underlying cryptographic primitives namely: **HKDF** (a specific Key Derivation Function that ensures that the keys obtained are cryptographically strong, provided that the underlying hash function is secure) and **HMAC** (a Message Authentication Code with SHA-256 as the underlying hash function) and on the security of the hash function they use, ultimately enabling us to understand privacy guarantees to individuals through familiar and well-studied cryptographic primitives and their computational (in)feasibility.

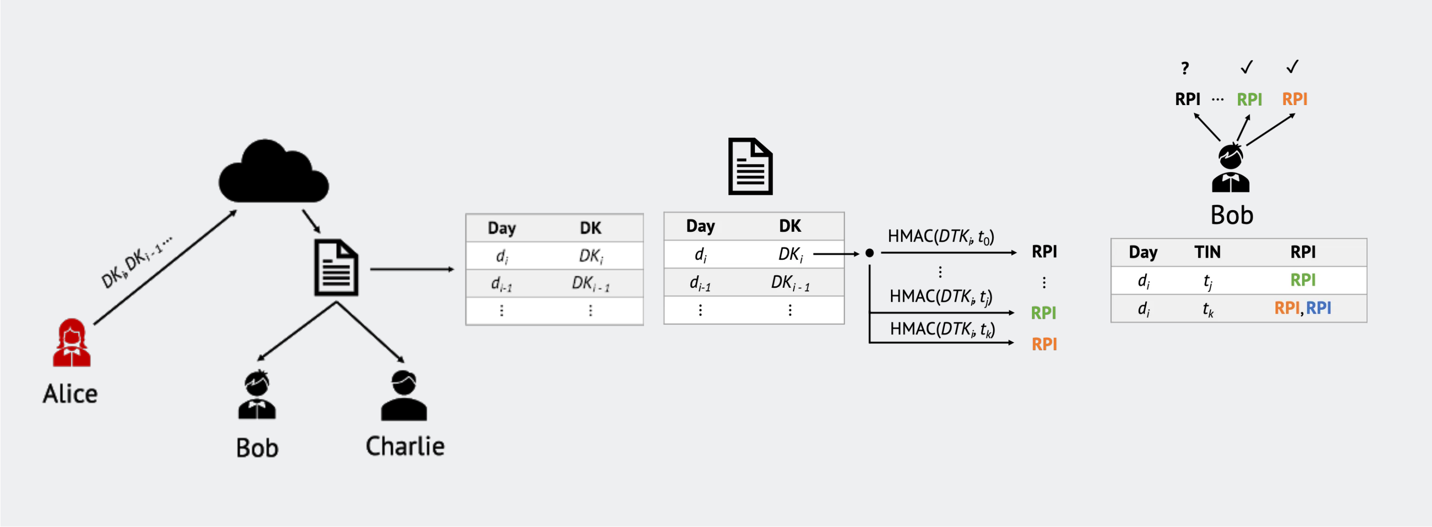

Alice is diagnosed positive. A set of her Diagnosis Keys are uploaded to the Diagnosis Server. The server then distributes these keys to all participants in the system. In a real-world scenario, the lists of Diagnosis Keys distributed by the Server would be aggregated from many infected users

Each participant iterates through the list of Diagnosis Keys and computes the corresponding Rolling Proximity Identifiers. They then check whether they have seen that RPI in the past, in which case they know they might be infected

[1] Version 6 at time of our code publication, which differs from Version 8 which is current at time of writing this report.

In June 2020, Paul Bottinelli published on the Security Considerations of zk-SNARK Parameter Multi-Party Computation. Since secure generation of parameters for zk-SNARKs is a crucial step in the trustworthiness of the resulting proof system, this work highlighted some potential pitfalls and important security considerations of the implementation of zk-SNARKs.

Beyond typical cryptographic considerations, for example, security audits for zk-SNARKs must include:

1. Secure deletion of the parameters (since the security of the multi-party computation ceremony relies entirely on the fact that at least one participant needs to securely delete their contribution for the resulting parameters to be generated honestly);
2. Adequate randomness (the use of standardized cryptographically secure pseudorandom number generator (CSPRNG) seeded with enough entropy to support the security needs of the parameters);
3. A rigorous, well-defined process for the trusted setup ceremony (including verifying all the transactions and validate all the parameters generated); and
4. Security and correctness of the source code implementing the trusted setup ceremony.

Throughout the summer, Eric Schorn published his 3-part code-centric blog series on Pairing-Based Cryptography, as well as an open-source implementation of BLS12-381 in Haskell. Pairing-based cryptography is of particular interest due to the popularity BLS12-381 pairing operations found in a variety of applications such as BLS signatures, including the proposal for use of these pairing operations in an Ethereum precompiled contract, as well as the current, real-world use of related pairing operations in precompiled contracts (EIP-196, EIP-197).

In August 2020, Thomas Pornin published a preprint of his paper "Optimized Binary GCD for Modular Inversion," on the IACR ePrint archive. In this paper, he outlined a new, practical optimization of the well-known extended binary GCD algorithm, for the purpose of computing modular inverses. He released open source implementations of this for both x86 CPUs as well as for ARM Cortex M0 and M0+ microcontrollers. This work is summarized in our blog post, "Faster Modular Inversion and Legendre Symbol, and an X25519 Speed Record".

In August 2020, Gerald Doussot published a Technical Advisory for CVE-2020-24613, a TLS 1.3 Client Man-in-the-Middle Attack for the wolfSSL cryptographic library. This advisory showed that wolfSSL had incorrectly implemented the TLS 1.3 client state machine, allowing attackers in a privileged network position to completely impersonate any TLS 1.3 servers and read or modify potentially sensitive information between clients using the wolfSSL library and these TLS servers. The patch to this bug was released on github by wolfSSL in August 2020.

The FALCON cryptosystem, among whose authors include Thomas Pornin of NCC Group, is currently under Round 3 consideration in NIST's Post-Quantum Cryptography standardization process. FALCON is a Round 3 finalist for the NIST post-quantum Digital Signature Algorithm standard, which "is based on the theoretical framework of Gentry, Peikert and Vaikuntanathan[2] for lattice-based signature schemes." This is based on the underlying hardness of the Short Integer Solution (SIS) problem over NTRU lattices, which is believed computationally infeasible for both classical and quantum computers, in the general case (that is to say, that no known efficient quantum algorithm for SIS is known).

(Fig 1)

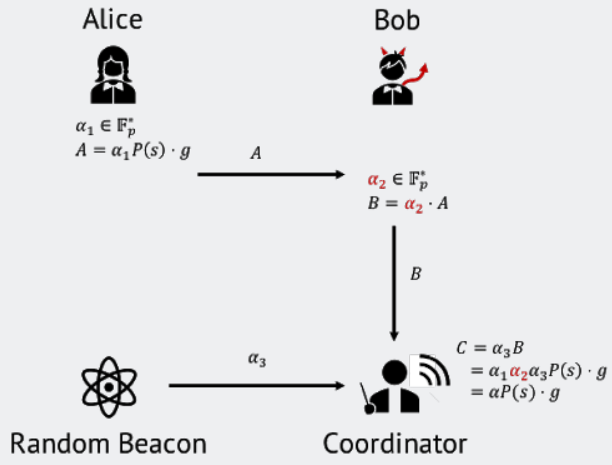
Alice
Bob
$\alpha_1 \in \mathbb{F}_p^*$
$A = \alpha_1 P(s) \cdot g$
A
$\alpha_2 \in \mathbb{F}_p^*$
$B = \alpha_2 \cdot A$
B
α_3
$C = \alpha_3 B$
$= \alpha_1 \alpha_2 \alpha_3 P(s) \cdot g$
$= \alpha P(s) \cdot g$
Random Beacon
Coordinator

2-party ceremony with Random Beacon

[2] Gentry, C., Peikert, C., and Vaikuntanathan, V. (2007). Trapdoors for Hard Lattices and New Cryptographic Constructions. https://eprint.iacr.org/2007/432

Open Source Tool Releases

Among the security tools released by NCC Group this year are:

- **Carnivore:** Microsoft External Attack Tool (a username enumeration and password spraying tool for Microsoft service such as Skype for Business, ADFS, RDWeb, Exchange and O365) created by Chris Nevin. This work was released at Black Hat USA 2020.

- **Collaborator++:** a project to extend upon the existing Collaborator functionality provided by Burp Suite - providing a number of quality of life features - and the implementation of an authentication mechanism to secure private collaborator deployments, while maintaining compatibility with all existing extensions which generate and poll Collaborator contexts, created by Corey Arthur.

- **Depthcharge:** an extensible Python 3 toolkit designed to aid security researchers when analyzing a customized, product-specific build of the U-Boot bootloader, created by Jon Szymaniak. This work was first released at a Hardwear.io session and also presented at the Open Source Firmware Conference, in a talk titled, "Guiding Engineering Teams Toward a More Secure Usage of U-Boot". This work also led to a contributed piece to Electronic Products on eliminating security flaws in open-source software for embedded systems.

- **Experiments in Extending Thinkst Canary - Part 1**, in which we extend Thinkst Canary to enable a TCP proxy to be built for honeypotting, created by Ollie Whitehouse.

- **go-pillage-registries:** which has a tool called pilreg which allows pentesters to more easily enumerate images stored in a registry in order to obtain their metadata and filesystems to assist in Docker container security assessments, created by Joshua Makinen.

- **HTTP Signatures:** A Burp Suite Extension Implementing HTTP Signatures, created by Roger Meyer. In the process of this research, Roger disclosed and helped remediate HTTP Signatures-related vulnerabilities in Oracle Cloud, PeerTube, Nextcloud, Pleroma, Friendica, and Hubzilla.

- **ICPin:** an integrity-check and anti-debug detection pintool, created by Nicolas Guigo.

- **idahunt:** a framework to analyze binaries with IDA Pro and hunt for things in IDA Pro, created by members of NCC Group's Exploit Development Group. It is a command line tool to analyse all executable files recursively from a given folder. It executes IDA in the background so you don't have to open manually each file. It supports executing external IDA Python scripts.

- **LDAPfragger:** a tool that uses the shared Active Directory component to build a communication channel between workstations joined to the same Active Directory domain but on physically segmented networks, in cases when only one network segment could connect to the internet, by Rindert Kramer of Fox-IT (a part of NCC Group).

https://github.com/nccgroup/

- **ROADtools and ROADrecon:** ROADtools is a framework to interact with Azure AD, consisting of a library (roadlib) and the ROADrecon Azure AD exploration tool, by Dirk-Jan Mollema. This work was released at Black Hat USA 2020.

- **Salesforce Policy Deviation Checker:** a tool which reveals which Profiles have become desynchronised from the Organization in Salesforce, and reviews each one's password policies and session settings to highlight any deviations from those set at the Organization level to prevent misconfiguration for certain sets of users, created by Jerome Smith. Jerome also blogged to offer further advice on securing Salesforce while remote working, and to outline common insecure practices with configuring and extending Salesforce, based upon his findings from across 35 security assessments of Salesforce customer deployments conducted by NCC Group.

- **ScoutSuite 5.8.0, ScoutSuite 5.9.0** and **ScoutSuite 5.10:** in which we made considerable extensions to NCC Group's powerful open source multi-cloud security auditing tool, ScoutSuite, which enables security posture assessment of cloud environments. Using the APIs exposed by cloud providers, Scout Suite gathers configuration data for manual inspection and highlights risk areas, and presents a clear view of the attack surface automatically.

- **Socks over RDP:** a tool that creates a virtual channel over an RDP connection and spins up a SOCKS5 proxy on a remote host, just like SSH's –D switch, created by Balazs Bucsay. This tool was later extended to work with Citrix. This tool was released at the HAVOC conference.

- **StreamDivert:** a tool for relaying (specific) network connections, created by Jelle Vergeer.

- **Whalescan:** a vulnerability scanner for Windows containers, which performs several benchmark checks, as well as checking for CVEs/vulnerable packages on the container, created by Saira Hassan.

- **Windows Executable Memory Page Delta Reporter:** an open source Microsoft Windows Service that aims to facilitate detection of anomalous executable memory, created by Ollie Whitehouse.

- **Winstrument:** an instrumentation framework for windows application assessments, created by George Osterweil.

- **WStalker:** an easy proxy to support Web API assessments, created by Jose Selvi.

We also released our first open dataset, three months of honeypot web traffic data related to the F5 CVE-2020-5902 and Citrix CVE-2020-8193, CVE-2020-8195 and CVE-2020-8196 exploitation events from earlier in 2020, with the intent to enable all threat intelligence researchers to gain further understanding and contribute back to the community.

Public Reporting on Open Source Security Audits

NCC Group has a long history of publicly-reported security audits of critical components of open source software as well as select proprietary systems. Of these reports, those labelled "Public Report" were developed as a part of a paid engagement with an NCC Group client to conduct and publish the findings of a security audit on in-scope components, whereas those labelled "Research Report" were internally funded and pursued by researchers at NCC Group. In 2020, NCC Group delivered 7 Public Reports and 1 Research Report, across a range of technologies including cryptographic implementations for high-value cryptocurrencies and the security of a popular open-source real-time operating system (RTOS), as well as of the Google Pixel 4/4XL and Pixel 4a smartphones as a part of NCC Group's role as an ioXt Authorized Lab.[3]

Our 2020 Public Reports and public Research Reports include:

- Public Report – Electric Coin Company NU3 Specification and Blossom Implementation Audit (January 2020)
- Public Report – RustCrypto AES/GCM and ChaCha20+Poly1305 Implementation Review (February 2020)
- Public Report – Coda Cryptographic Review (May 2020)
- Research Report – Zephyr and MCUboot Security Assessment (May 2020)
- Public Report – Qredo Apache Milagro MPC Cryptographic Assessment (July 2020)
- Public Report – Pixel 4/4XL and Pixel 4a ioXt Audit (August 2020)
- Public Report: Electric Coin Company NU4 Cryptographic Specification and Implementation Review (September 2020)
- Public Report: Filecoin Bellman/BLS Signatures Cryptography Review (October 2020)

[3] Notably, while our Public Reports are often performed by our researchers, it should be noted that Public Reports result from billable client engagements, while our Research Reports reflect unpaid independent research into the security properties of a target system.

Hardware and Embedded Systems

In February 2020, Sultan Qasim Khan published a whitepaper titled Microcontroller Readback Protection: Bypasses and Defenses, which described common readback protection implementation flaws, techniques that can be used to defeat readback protection, and guidance to implement effective readback protection. He also discussed this work in a bylined article with Electronic Design in May 2020, titled Implementation flaws in Microcontroller Readback Protection.

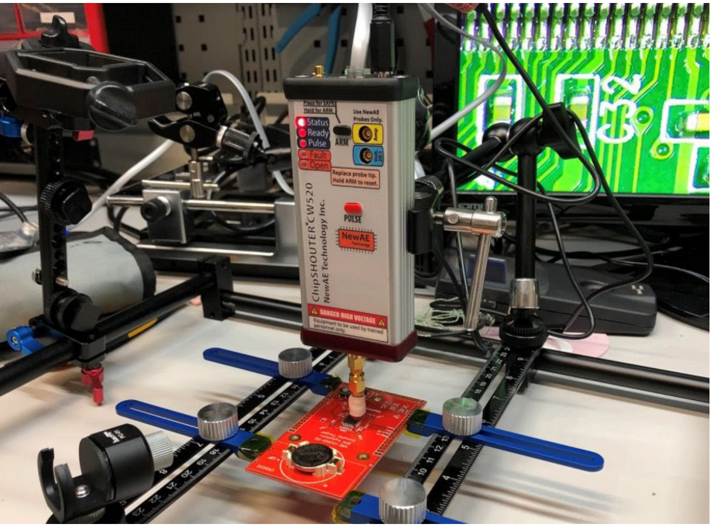

In May 2020, Jeremy Boone and Ilya Zhuravlev published a research report on their security assessment of the Zephyr open source real-time operating system (RTOS) and the MCUboot bootloader, uncovering 25 vulnerabilities affecting the Zephyr RTOS and 1 vulnerability affecting MCUboot. Many of these 26 vulnerabilities were also discussed in depth by Jennifer Fernick with Embedded Computing Design. These findings include both locally and remotely exploitable memory corruption vulnerabilities, multiple paths that allow a compromised user application to escalate privilege to kernel mode, as well as multiple weaknesses in the design of certain exploit mitigation systems that exist within the kernel.

The scope of this project included the robustness of the secure boot implementation, the kernel mode execution protection, and the kernel driver review. Our major findings included a number of both remote attack vectors and locally exploitable weaknesses, and a number of places in which insufficient argument validation in system call interfaces led to exploitable weaknesses, where in one example, a compromised userspace application was able to reveal the contents of restricted kernel memory. Furthermore, while Zephyr implemented a range of kernel hardening measures, many of them were found to contain weaknesses, and further countermeasures and mitigations are necessary to limit the impact of memory safety violations and reduce the likelihood that a single memory corruption vulnerability can result in a complete system compromise. The Governing Board Chair of the Zephyr Project responded positively to our research and coordinated disclosure. This work was also discussed in Dark Reading.

In July 2020, Jon Szymaniak released a tool, Depthcharge, which is an extensible Python 3 toolkit designed to aid security researchers when analyzing a customized, product-specific build of the U-Boot bootloader. He also published a proof-of-concept attack about a Sonos Symfonisk speaker, which leveraged NXP's High Assurance Boot (HAB), despite being one the lower-cost offerings within the Sonos product line. Unfortunately, the bootloader contained a type of vulnerability that we've previously observed in other embedded platforms – a functionality-reduced console contains an unauthenticated command that can be abused as an arbitrary memory read-write operation, leading to bypass of secure boot on the device. Through this, he demonstrated that all it takes is one extraneous enabled feature in a U-Boot configuration to undermine all the hard work involved in properly extending a hardware-backed root of trust, thus emphasizing the importance of configuring and auditing the security properties of open source security tools to meet the needs of one's individual product and threat model. This work was also presented in a webinar hosted by Hardwear.io, and discussed in a contributed piece to Electronic Products on eliminating security flaws in open-source software for embedded systems, as well as in a video interview with TFIR.

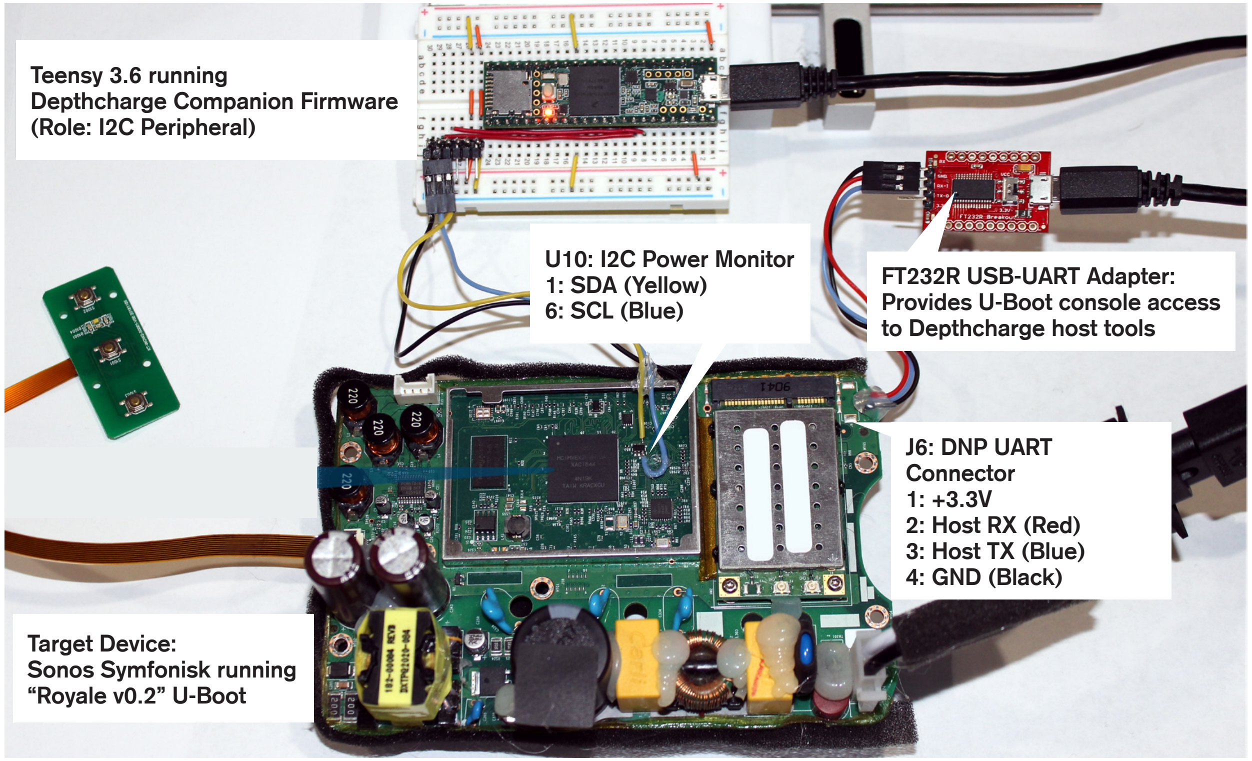

In his December 2020 release of Depthcharge v0.2.0, Jon sought to help embedded systems engineering teams proactively tackle security debt, through the introduction of a security configuration checker in Depthcharge for the popular U-boot bootloader. This work was discussed in his presentation at the Open Source Firmware Conference, in which he demonstrated how this tooling could be used to detect software supply chain risks in embedded systems, through identifying problematic default configurations in a firmware image, giving defenders actionable, real-world tools for finding weak configurations in bootloaders, thus eliminating that aspect of supply chain security risk. **(Fig 2)**

In August 2020, Rob Wood published an article on embedded.com on how a little early effort in security can return a huge payoff when designing embedded systems. He explained how early in the product development cycle of embedded systems, threat modelling can help identify key security requirements for a system that are difficult, costly, or impossible to introduce into an embedded system at a later time. He gave specific questions to ask of your component vendors when selecting the platform, processor, bootloaders and operating system of your target system, and offered other considerations in vendor selection around patch lifecycles and manufacturing and supply chain risks.

In October 2020, Rob Wood gave a presentation about building secure devices in untrusted factories. In this presentation, he outlined that since with very few exceptions, no electronics manufacturer actually designs and manufactures every single component of a device in their own factory, hardware and manufacturing supply chains introduce considerable risk that threat actors could gain an opportunity to defraud, steal, counterfeit, or otherwise undermine the security of the produced electronic devices.

In October 2020, Ilya Zhuravlev and Jeremy Boone published “There’s A Hole In Your SoC: Glitching The MediaTek BootROM”. This blog post describes NCC Group’s methodology for characterizing the boot process of the MediaTek MT8163V system-on-chip (64-bit ARM Cortex-A), as well as the design of an apparatus that is capable of reliably producing a fault injection attack against the SoC. Ultimately, our results show that the MediaTek BootROM is susceptible to glitching, allowing an adversary to bypass signature verification of the preloader. This circumvents all secure boot functionality and enables the execution of unsigned preloader images, completely undermining the hardware root of trust. This work was cited in a paper by the Microsoft Azure Sphere team on The Seven Properties of Highly Secured Devices, in their discussion of how highly-secured devices have a hardware root of trust. They cite our research when outlining real-world glitching attacks against which a highly secured device must both detect and mitigate.

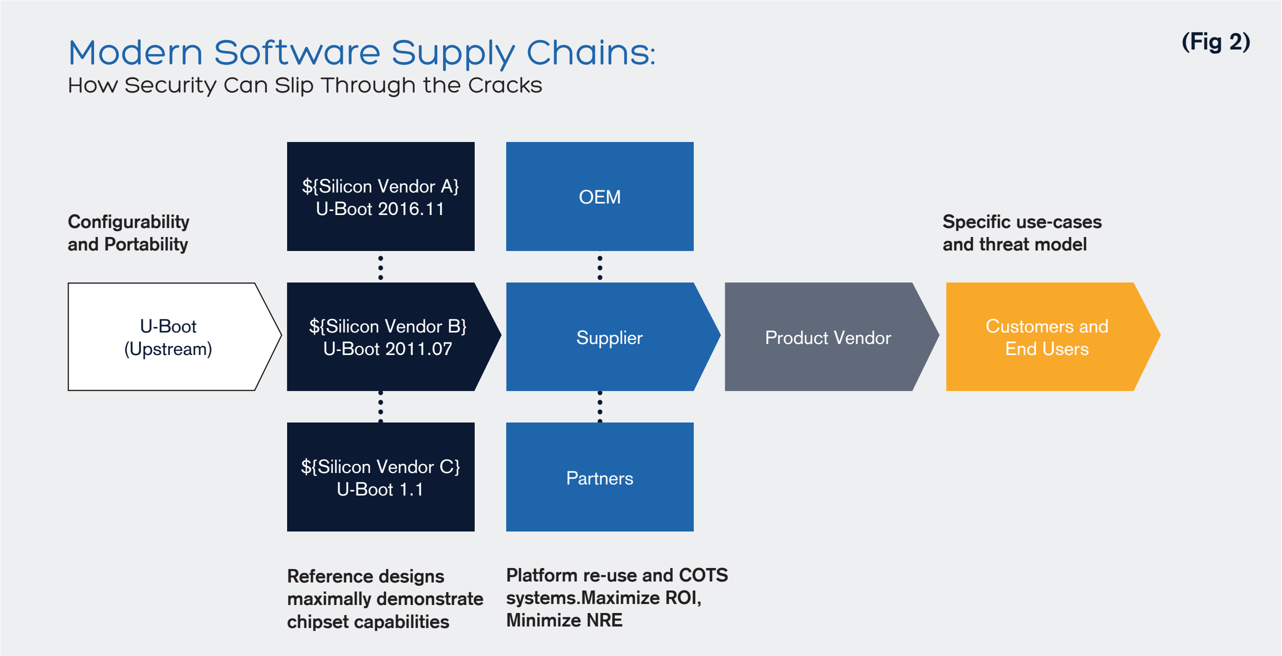
Modern Software Supply Chains:
How Security Can Slip Through the Cracks
Configurability and Portability
U-Boot (Upstream)
${Silicon Vendor A} U-Boot 2016.11
${Silicon Vendor B} U-Boot 2011.07
${Silicon Vendor C} U-Boot 1.1
OEM
Supplier
Partners
Product Vendor
Customers and End Users
Specific use-cases and threat model
Reference designs maximally demonstrate chipset capabilities
Platform re-use and COTS systems.Maximize ROI, Minimize NRE

(Fig 2)

Cloud, Virtualization, and Container Security

Container Security

Throughout the year, Rory McCune offered his popular "Mastering Container Security" training, including at Black Hat USA 2020, Black Hat Europe 2020, NorthSec, and 44CON. He discussed common container and Kubernetes vulnerabilities with Container Journal, and with Help Net Security on how our existing security environments need to be updated to account for differences in how work of containerization and cloud native computing operates.

At Shmoocon 2020, Mark Manning presented, "Command and KubeCTL: Real-World Kubernetes Security for Pentesters" (slides, demos, video here). This talk was aimed at people doing offensive security or looking at Kubernetes security from the perspective of an attacker, and was a demo-focussed talk showing an attack chain at three different kinds of companies:

1. Compromising a cluster at a company that was completely insecure… but the company accepted the risk and didn't make any changes
2. Compromising a company that was very secure for their threat model… but then pointed out that few organizations have the resources to make an environment like this, and
3. Compromising a company that was trying to do something very complex and demonstrating that compromising it required the same level of complexity.

Dark Reading covered this presentation in an article about how Kubernetes shows built-in weakness.

Following this conference presentation, Mark Manning published Deep Dive into Real-World Kubernetes Threats, in which he offered a deeper dive into one of his case studies presented at Shmoocon. In this case study, he went through an attack chain for an example organization that is trying to do per-namespace multi-tenancy, demonstrating how to go from pod compromise, to namespace compromise, to namespace tenant bypass, to node compromise, to cluster compromise. He then offered 8 specific defenses that could mitigate the risks that made this company (and many other, real-world organizations) vulnerable.

In January 2020, Joshua Makinen published his open source tool, go-pillage-registries, which allows pentesters to more easily enumerate images stored in a registry in order to obtain their metadata and filesystems to assist in Docker container security assessments.

In March 2020, Jack Leadford published "A Survey of Istio's Network Security Features," in which he took a critical look at a subset of the features of the Istio service mesh, focussing on Istio's security features that are used to control services' network interactions: restricting egress traffic, restricting ingress traffic, and requiring a TLS client certificate from a service's callers (i.e. mutual TLS). He walked through lab-based examples drawn from Istio's documentation in order to concretely illustrate the limits of each of these features, possible misconfigurations, and what can be done to ensure operators' intended security goals are met. This post foreshadowed some changes which were upcoming in the Istio 1.5 release, such as changes to the AuthorizationPolicy which will allow operators to express mesh-level authorization rules that cannot be overridden at the namespace or service-level.

In March 2020, Derek Hinch presented Smashing Containers for Food and Profit at BSides Atlanta. This presentation discussed how microservice container-based architectures have introduced several layers of abstraction between successful exploitation of an application and ultimately compromising the host system/corporate network, with the intent to demonstrate how penetration testing methodologies must adapt to the challenge these layers of abstraction present, including container breakouts via misconfiguration, and post exploitation horizontal/vertical attacks (including orchestration privilege escalation attacks).

Also at BSides Atlanta, Rory McCune presented Compromising Containers and Clusters in which he explored new techniques for breaking out of containerized systems, attacking standalone Docker daemons with mere curl and ssh, and going from existing in a Kubernetes cluster to owning the whole environment.

In November 2020, NCC Group's Rory McCune spoke on the Keynote panel at CloudNativeCon/KubeCon, SIG-Honk AMA Panel: Hacking and Hardening in the Cloud Native Garden (video) which explored their experiences with securing, attacking, and deploying cloud native infrastructure, and their formation of "sig-HONK," an unofficial Special Interest Group focused on changing the way we think about and practice security in distributed systems. He also wrote a contributed article outlining some key concerns around container security and the cloud native ecosystem.

Rory McCune also continued to make contributions to the Center for Internet Security's CIS Benchmarks for Kubernetes.

On November 30 2020, Jeff Dileo published Technical Advisory: containerd – containerd-shim API Exposed to Host Network Containers (CVE-2020-15257). This vulnerability pertained to containerd, a container runtime underpinning Docker and common Kubernetes configurations, which handles abstractions related to containerization and provides APIs to manage container lifecycles. He found that the containerd shim API is exposed over an abstract namespace Unix domain socket that is accessible from the root network namespace. This results in non-user namespaced containers with host networking being able to access this API and cause containerd-shim to perform dangerous actions and spin up arbitrarily privileged containers, enabling container escapes and escalation to full root privileges on the host. An article in Dark Reading discussed these findings on how a common container manager is vulnerable to a dangerous exploit.

He elaborated on this work in his blog post, ABSTRACT SHIMMER (CVE-2020-15257): Host Networking is root-Equivalent, Again, a technical discussion of the underlying vulnerability of CVE-2020-15257, and how it can be exploited, going deeper into the process that led to finding the issue, the practicalities of exploiting the vulnerability itself, various complications around fixing the issue, and concluding that abstract namespace Unix domain sockets can be extremely dangerous when applied to containerized contexts (especially because containers will often share network namespaces with each other). He mused that it is unclear how the risks of abstract namespace sockets was not taken into account by the core infrastructure responsible for running the majority of the world's containers. It is also unclear how this behavior went unnoticed for so long. If anything, it suggests that containerd has not undergone a proper security assessment.

Securing the Public Cloud

In 2020, our open source multi-cloud security auditing tool, ScoutSuite, continued its ongoing evolution, releasing a number of key features. ScoutSuite has long supported Amazon Web Services, Google Cloud platform, and Microsoft Azure, but in 2020 we also introduced alpha support for Alibaba Cloud and Oracle Cloud infrastructure. This work is led by Xavier Garceau-Aranda (and features a number of contributors, visible on the project's github page). Throughout the year we released ScoutSuite 5.8.0, ScoutSuite 5.9.0, and ScoutSuite 5.10.

Xavier Garceau-Aranda taught his Offensive Cloud Security Training at OWASP Global AppSec in October 2020. In this applied, hands-on training, attendees experienced first-hand how security vectors that exist in cloud ecosystems present opportunities for abuse, as well as detection and mitigation of the attacks covered in the course.

Throughout the year, Rami McCarthy shared a number of projects related to AWS security for end-users. In An offensive guide to the Authorization Code grant blog which introduced, broke down the observable vulnerabilities, and explained the exploitation of each of several aspects of the OAuth 2.0 Authorization Code flow. In The Extended AWS Security Ramp-Up Guide he outlines a number of learning resources to improve one's security knowledge as applied to AWS. In his BSides Boston presentation, AWS Security: Easy Wins and Enterprise Scale, he discusses both the easy security wins that almost anyone or any organization can identify and apply to their AWS environment, as well as the big picture security problems to consider as an organization's security maturity or AWS usage grows.

Erik Steringer is working on significant extensions to PMapper, his graph-based, open source tool for evaluating and visualization IAM permissions in AWS. We will officially release the more featureful, extended version of this tool in 2021.

Jerome Smith continued his efforts to secure popular PaaS systems, publishing a blog post on common insecure practices in configuring and extending Salesforce in June 2020. In October 2020, he published Salesforce Security with Remote Working, where he outlined methods for increased enterprise monitoring of Salesforce deployments to increase assurance with a remote-first workforce, and emphasizing the importance of proactive anomaly detection and the retroactive availability of logging for investigations, reminding us that not all successful authentication attempts are legitimate. He also published his tool, Salesforce Policy Deviation Checker.

In November 2020, Juan Garrido gave a presentation at the CCN-CERT Annual Conference on Bypassing Security Controls in Office 365. In this talk, he demonstrated multiple techniques for bypassing existing Office 365 application security controls, showing how data can be exfiltrated from highly secure Office 365 tenants which employ strict security policies in order to restrict access to a range of predefined IP addresses or subnets, or configured with Conditional Access Policies, which are used to control access to cloud applications.

Data Science, Machine Learning, and A.I.

Throughout 2020, we continued our ongoing research into circumventing real-world A.I.-driven systems through adversarial examples, and have built and trained a number of models and developed demos of high-efficacy adversarial examples. Into 2021 we have renewed our commitment and increased our support for research into offensive security against A.I. systems and how to consequently defend (and formalize the limits to defenses of) real-world systems against an increasingly diverse set of attack types. Beyond the field of adversarial A.I., we have also been studying many aspects of machine learning, not only as an attack target but also to enhance the efficacy of both offensive and defensive security techniques.

In April 2020, Liam Stevenson of NCC Group's Managed Detection and Response group authored a blog post on practical machine learning for random filename detection, which explored the efficacy of using specific machine learning algorithms to analyze the so-called random filenames generated by malware or implants when creating files on a disk to avoid traditional Indicator of Compromise signatures. In this work, he looked at the use of N-gram scoring of filenames, to calculate the probability of seeing various trigrams, and using the unlikeliness of trigrams in recently-observed filenames as an indicator of "how random" the randomness of a given malware filename generated actually is, and giving advice on parameter tuning and other intervention measures which can be used to increase of decrease the rate of false positives.

In May/June 2020, we published the whitepaper and associated blog post, Exploring DeepFake Capabilities and Mitigation Strategies with University College London, which was the result of a collaborative research partnership between NCC Group and the University College London (UCL)'s Centre for Doctoral Training in Data-Intensive Science. This paper discusses the cybersecurity implications of deepfakes, including the difficulty and efficacy of creating deepfaked content using extant open source methods, ultimately concluding that open source frameworks offer technically impressive deepfakes, which however are not as convincing as would be possible through the use of present-day GPU clusters. covered deepfake production workflows, limitations of deepfake technology, and methods for detection. This research included a proof-of-concept in which the students deepfaked NCC Group's Matt Lewis into a clip of Daniel Craig in the 2006 film, Casino Royale (pictured below), on consumer hardware and analyzed visual artifacts indicative of its' synthetic origin. This work was featured in interviews with Dark Reading, TFIR, and a guest blog post on Dark Reading by NCC Group's director of commercial research, Matt Lewis, in which he notes the many nuances to creating convincing deepfakes on high-end consumer hardware, but ultimately concludes that these technologies will inevitably continue to improve with time.

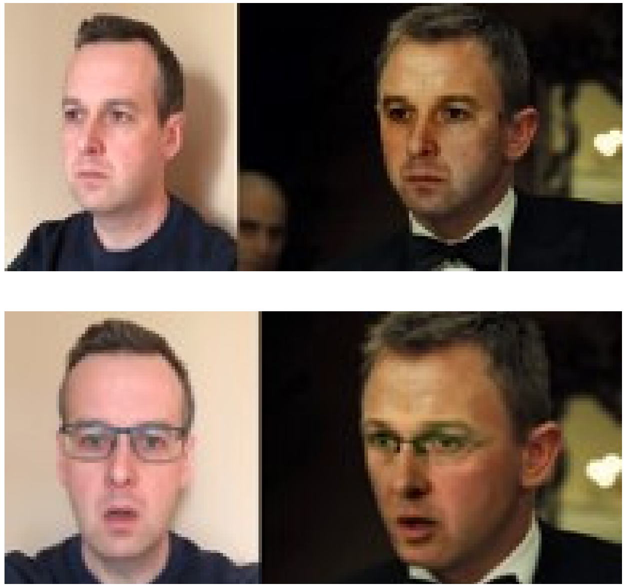

In September 2020, Jennifer Fernick discussed the arms race in machine learning-driven deepfake detection methods in an article in Dark Reading about defending against deepfakes, ultimately highlighting both the relatively poor human performance on positively-identifying deepfaked imagery, as well as the challenge intrinsic in machine-driven methods - namely, that as soon as a detection system “learns” artifacts or indicators that a video is deepfaked, the deepfake-generator will have had those properties selected against, enabling their new deepfakes to continue to evade detection with high probability. The difficult consequence of this is that it is likely for the deepfake detection problem to only become asymptotically harder over time, and counterproposing that we instead focus on cryptographically robust authentication mechanisms for audio and video streams to become better, as a society, at being confident when a video clip or audio file has not been synthetically generated.

Clinton Carpene and Peter Hannay gave a presentation titled “Let’s Play(AI) Super Metroid!” where they demonstrated the use of machine learning to automate the SNES classic Super Metroid.

In Machine learning from idea to reality: A Powershell case study Joost Jansen discussed detecting both ‘offensive’ and obfuscated PowerShell scripts in Splunk using Windows Event Log 4104, leading to detection models that can be used in a real-time environment.

During summer 2020, we also added two full-time data scientists to our team, who work on both internal- and external-facing research related to both offensive and defensive uses of machine learning, with research publications set throughout 2021.

We have also renewed our commitment to collaborative research with the University College London’s Center for Doctoral Training in Data Intensive Science for the 2020-2021 academic year, in which we will mentor students on a new project related to machine learning-driven enhancements to the detection of malicious software binaries.

In October 2020, Gene Meltser presented on the State of Disinformation on Social Media at Connecticut College, and in our 2021 fiscal year we committed to a fully-funded, dedicated research program understanding and combating technology-mediated disinformation campaigns.

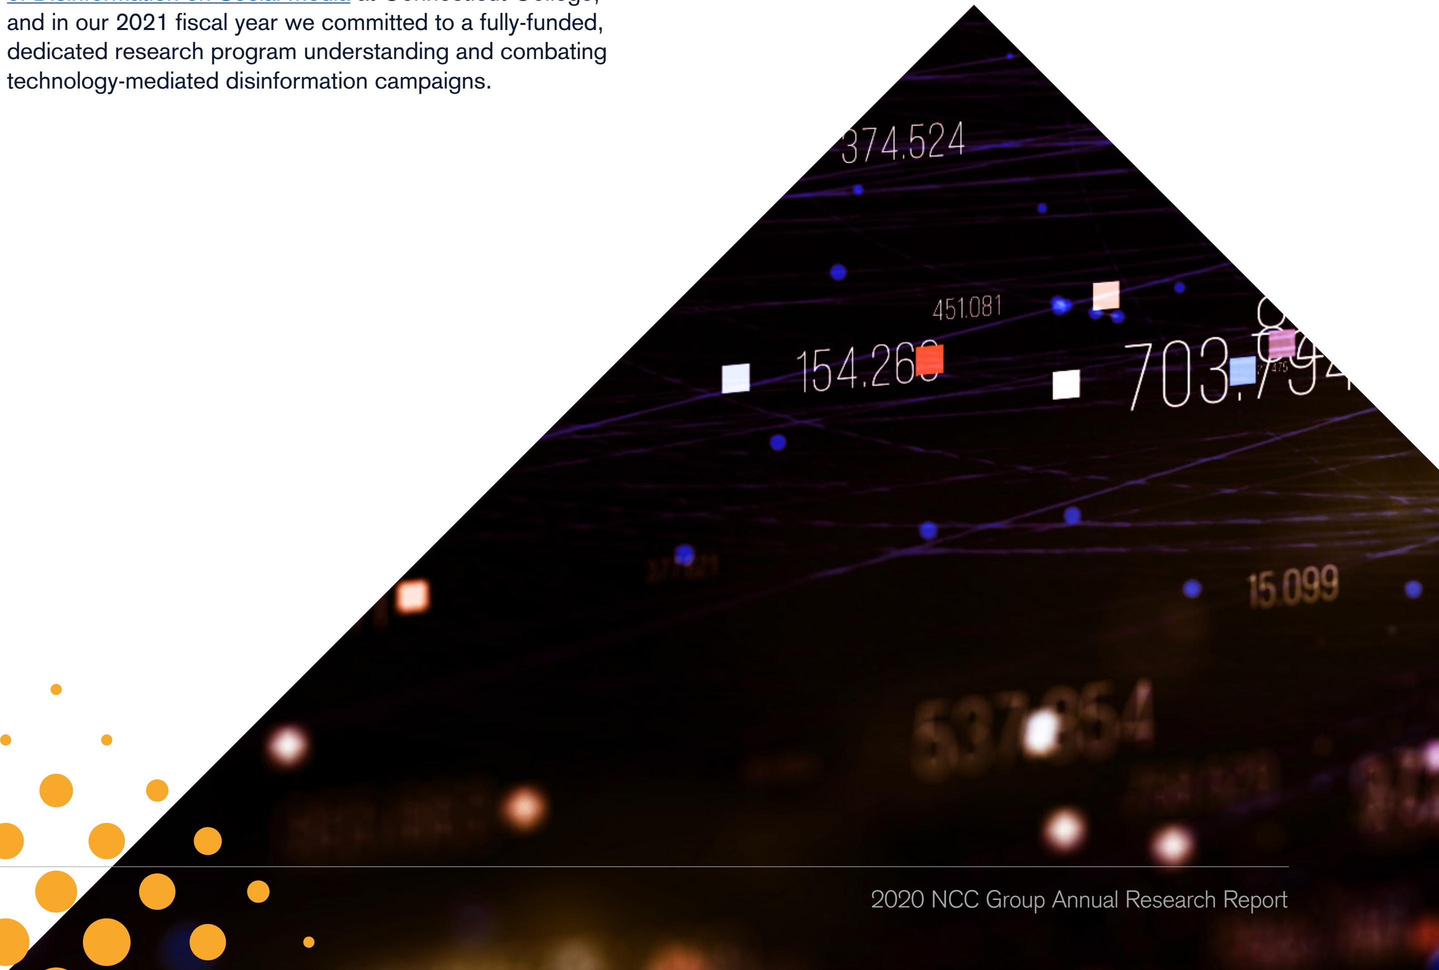

Public Interest Technology

For many technologists at NCC Group, making the world a safer and more secure place is their life's work. We do this not only through our consulting work with clients, but also through the significant research and other NCC Group-supported work that we do in the public interest. This year, we had a strong focus on enhancing individual privacy amidst the coronavirus pandemic, ensuring privacy protections for users of lower-cost mobile devices, and using threat intelligence and OSINT techniques for social good.

Protecting Privacy in a Pandemic

Paul Bottinelli wrote a blog post on the privacy-preserving design properties of Google and Apple contact-tracing apps, in response to public concern around the potential privacy impact of contact-tracing applications. While it is definitely the case that some contact-tracing apps released by governments around the world have had privacy-violating properties, in this post Paul demonstrate how the design of this specific contact-tracing app, and its thoughtful use of specific cryptographic primitives, mitigated against several key privacy concerns present in some of the other application designs.

NCC Group CTO Ollie Whitehouse also discussed contact-tracing apps with Communications Daily, emphasizing that systems design should be deliberate in not collecting excessive data, and validating the idea that it is not necessary to identify individuals outright in order to provide effective COVID-19 exposure notifications.

Robert Wessen sought to better understand the security and privacy properties of the popular open-source video conferencing software, Jitsi, publishing a number of vulnerabilities in Jitsi Meet Electron including arbitrary client remote code execution (CVE-2020-27162, CVE-2020-27161), following coordinated disclosure. Our broader contributions to the improvement in security of video conferencing platforms in a remote-working world was also noted in The Wall Street Journal, Infosecurity Magazine, and ZDNet, in their discussion of security experts hired by Zoom to perform assurance services of components of their platform, given the meteoric growth in popularity of video-conferencing services as many knowledge workers around the world rapidly transitioned to working from home in March 2020.

Damon Small's presentation at Shellcon 2020 sought to offer actionable advice to keep your home office secure. Through the year, we also offered advice on cybersecurity in the era of remote work in Security Today, SecurityInfoWatch, Small Business Computing, Security Informed, CIO Dive, Crain's New York Business, and others.

In June 2020, our threat intelligence research into Evil Corp was covered in The New York Times, where they discussed our observations on how Russian criminal groups are targeting Americans working from home. NCC Group's Maarten van Dantzig specifically discussed the evolution of the associated threat actors' techniques over the past few years, emphasizing an increasing capability and increased ransom demands over time.

Mobile Privacy

In August 2020, Neil Bergan published a whitepaper [Exploring the Security of KaiOS Mobile Applications](). KaiOS is a mobile operating system forked from the discontinued Firefox OS, in which all the mobile applications running on a KaiOS-based mobile device are built using web technologies, such as HTML, JavaScript, and CSS. In this paper, he scrutinized pre-installed mobile applications on four different phones running KaiOS, resulting in both a collection of remediated vulnerabilities, as well as constructive product security recommendations for the next generation of KaiOS operating system and its applications.

He discussed his findings of remote, and local, HTML injection attacks, which when combined with bypasses in the Content Security Policy can result in the abuse of privileged JavaScript APIs resulting in remote file disclosure or local privilege escalation, as well as explored the security implications of both documented and undocumented JavaScript APIs in the platform and general security risks of the mobile platform.

Mobile Device	SELinux	Disk Encryption	Verified Boot	Signing Keys	ADB Security
Alcatel Flip 2	Disabled	Unencrypted	Disabled	Signed with test-keys	Enabled by default. No secure USB debugging used.
Doro 7050	Permissive – not enforcing	Unencrypted	Disabled	Signed with release-keys	Disabled by default. No secure USB debugging used.
Nokia 8110	Permissive – not enforcing	Unencrypted	Disabled	Signed with test-keys in firmware 12. Signed with dev-keys in firmware 16.	Disabled by default. No secure USB debugging used.

Among the topics discussed in our KaiOS whitepaper was the general failure within this generation of KaiOS devices to make use of existing Android Platform Security features

At Rightscon 2020, the "world's leading summit on human rights in the digital age," NCC Group researcher Dan Hastings presented his work on the privacy-violating behavior of robocall-blocking mobile apps, titled, Ironically, iOS robocall-blocking apps are violating your privacy.

This project, an extension of his work presented at DEF CON 27, looked at mobile privacy through a human rights and targeted population's perspective, seeking to illuminate the methods through which technologists can perform deep privacy analysis of mobile applications in use, and demonstrating several new findings of apps promoted to be privacy-promoting that were in fact the exact opposite. Dan also showed that in many cases, mobile apps are not conforming with their data and privacy practices as stated in End User License Agreements, meaning that a deep technical analysis is often required for users to truly understand the privacy impacts of the apps on their phones.

Inspired by his findings from the robocall-blocking apps project, in the final days of 2020, Dan Hastings presented his new privacy analysis tool at Chaos Communication Congress. This tool, Solitude, is an open source privacy analysis tool that enables anyone to conduct their own privacy investigations on iOS and Android applications. Oftentimes the only way for the end user to figure out where their private data goes once they enter it into a web application or mobile device is through the app's privacy policy.

Privacy policies not only have a notorious history of being difficult to understand but don't always tell the truth about an application's data collection practices. Solitude was built to make proxying one's web and mobile traffic easier, and to thus make the process of conducting privacy investigations of one's favourite apps more streamlined and straightforward. Solitude can be configured to look for any data input into a mobile or web application, and reveal where that data is going. Whether a curious novice or a more advanced researcher, Solitude makes the process of evaluating an app's privacy accessible for everyone.

Security, OSINT, and Threat Intelligence for social good

In April 2020, our Research and Intelligence Fusion Team, driven by researchers from NCC Group's Netherlands-based Fox-IT offered free cyber threat intelligence to any interested global healthcare providers in light of increased cyberattacks targeting healthcare institutions in the wake of the global Coronavirus pandemic. Institutions including national Computer Emergency Response Teams, hospitals, and national Institutes of Public Health received threat intelligence including executive briefings, threat actor descriptions, technical Indicators of Compromise, as well as guidance for applying threat intelligence to their systems and a chronology of targeted ransomware in hospitals and health clinics, to help these institutions to detect cyber criminals at an early stage and take actionable steps to improve their resilience.

In October 2020, NCC Group APAC's Richard Appleby, Peter Hannay, Dean Hardcastle, and Clinton Carpene participated in the 2020 National Missing Persons Hackathon, which welcomed over 100 teams to a 6 hour event where teams were tasked with discovering information about actual missing person's cases where the authorities have run out of leads. The Hackathon is run by Trace Labs, a non-profit organisation whose mission is to reunite missing persons with their families while training members in Open Source Intelligence (OSINT).

We have continued and rekindled our long-running work with the Open Technology Fund to security audit third-party open source projects related to internet freedom and censorship resistance, to help identify and remediate security weaknesses in these codebases.

Securing Connected Environments

A great deal of work within cybersecurity research presently involves securing connected environments. On one hand, it is easy as researchers to dismiss pursuing further research into compromising internet-of-things (IoT) devices because it is so often trivial to do. On the other hand, we must acknowledge that technologies are not simply abandoned because they are too easy to hack - rather, we must understand the world in which we live and what it is becoming, and endeavour to secure it as well as possible, perhaps even especially in cases where the devices are ubiquitous, the market is rapidly growing, the impacts can affect human safety, and the vulnerabilities are readily found and exploits are often relatively simple to construct.

Fortunately, in 2020 we finally began to see some positive developments towards meaningful improvement of IoT security, likely driven in no small part by countless examples of security research illustrating the overall inadequate state of security across a range of consumer and industrial IoT devices.

Progress toward improved IoT security in 2020 perhaps was most notably observed in the passing of the IoT Cybersecurity Improvement Act in the United States, which directs the National Institute of Standards and Technology (NIST) to develop standards and guidelines on how federal government agencies should appropriately use and manage IoT devices connected to information systems. We also saw numerous global examples of additional progress in this space, including Singapore's announced intention for Cybersecurity Labelling Scheme for IoT devices, he UK's Department for Digital, Culture, Media and Sport (DCMS) announcing new legislation making manufacturers accountable for building robust security standards in early, from the very design stage of their products, the European ETSI Technical Committee on Cybersecurity releasing a globally applicable standard for IoT security, and the introduction of IoT security laws and frameworks in Australia, California, and Oregon.

The Internet of (Secure) Things

From an **industry** perspective, in 2020, NCC Group was one of the early members of ioXt, which aims to build confidence in IoT products through international, harmonised, and standardised security and privacy requirements, product compliance programs, and public transparency of those requirements and programs.

NCC Group is an IoXT Authorized Lab (discussed further here), providing third-party validation and testing of relevant products for the Alliance's Certification Program. One of the most notable aspects of IoXT is their security pledge:

- The ioXt Security Pledge, sets out 8 security standards "that bring security, upgradability and transparency to the market and directly into the hands of consumers".[4]
  - No universal passwords
  - Secured interfaces
  - Proven cryptography
  - Security by default
  - Verified software
  - Automatic security updates
  - Vulnerability reporting program
  - Security expiration date

From a **government and policy** perspective, we contributed to some of the discourse around government procurement of IoT devices, the generally woeful state of IoT device security, and the need for higher cybersecurity standards in IoT devices around the world:

- NCC Group's Peter Hannay co-authored a book chapter on the security of IoT devices in military networks, together with Ph.D student Imran Malik. In this publication, they discuss networked devices ("Network of Things," or "NoTs"), and how the security of NoT has not kept pace with the rapid advances in innovation and deployment of NoT devices, thus creating significant risks to all associated information assets. They argue that this security risk is prevalent in government networks due to pressures toward speed and cost-effectiveness of technology deployment, thus risking reliability and efficacy of military operations. The paper stresses the need for improving security posture when NoTs are employed within military systems and proposes measures to address the ever-evolving cyber security challenges in NoTs.

- In December 2020, Rob Wood was interviewed in an article about the IoT Cybersecurity Improvement Act, which was signed into law in the United States on December 4 2020, and seeks to ensure that the U.S. government only buys secure devices and remediates known vulnerabilities in existing devices. In this interview, he advocated for meaningful and advanced technical standards for IoT security to be put forward by NIST, beyond those that were introduced as a result of the existing, consumer-oriented California SB-327 and Oregon HB2395 laws, which only implemented baseline controls, primarily in response to the Mirai botnet.[5]

[4] https://www.ioxtalliance.org/the-pledge

[5] This same article proceeds to cite Nokia's 2020 Threat Intelligence Report in stating that while IoT was responsible for over 16% of infections observed in mobile networks in 2019, it was responsible for nearly 33% in 2020.

From a **consumer** perspective, we published a number of technical findings across a range of consumer IoT devices including IP cameras, smart doorbells, and personal IoT gateways, as well as offered several technical presentations.

In February 2020, Jason Kielpinski published Interfaces.d to RCE, in which he discusses some research into the Mozilla WebThings IoT gateway, which is a gateway essentially allows a user to host their own IoT cloud from a device (such as a Raspberry Pi) on their local network, creating a tunnel to a personal subdomain of mozilla-iot.org for managing a user's devices from the internet. In this post, he discussed his general excitement about users being able to self-manage their IoT devices, but also revealed input validation issues on the device which led to not only data injection, but code injection too! He noted that by inserting one of these event hooks into the file, it's possible to get command execution, and that since these scripts must run as root in order to be able to manipulate the network settings, this leads to a complete compromise of the affected gateway. Fortunately, when we reported the vulnerability to Mozilla, they implemented a fix and automatically applied it to all affected devices the very next day.

In July 2020, Dale Pavey published Lights, Camera, HACKED! An insight into the world of popular IP cameras. In this project, he created a testbed for studying the security of 5 popular IP cameras, resulting in numerous findings, including that one product is still vulnerable to Heartbleed, 6 years after its initial discovery. We also say things like the use of default credentials, weak non-existent encryption, ubiquitous data leakage about both developer or consumer.

He also published a related technical advisory, Heartbleed chained with a Pass-the-Hash attack leads to device compromise on TP-Link C200 IP Camera.

In July 2020, Jeremy Boone of NCC Group's Hardware and Embedded Systems practice joined the Embedded Insiders podcast, where they discussed how flaws at all levels of the IoT solution stack can be exploited, as well as engineering best practices that can minimize these vulnerabilities.

At the DEF CON IoT Village, Dewank Pant presented on how a wide range of IoT devices may be used as a point of compromise from which to directly exfiltrate data from a victim's device, or as a pivot point into a network from which an attacker can perform data exfiltration, as well as suggested attack mitigation techniques aimed at product designers and developers.

In September 2020, our printer research findings of 35+ vulnerabilities across 6 major vendors presented at DEF CON in 2019 was revisited by Dark Reading, where we further discussed how an uptick in remote work and internet-connected consumer printers results in an increased attack surface for enterprises with remote-working employees.

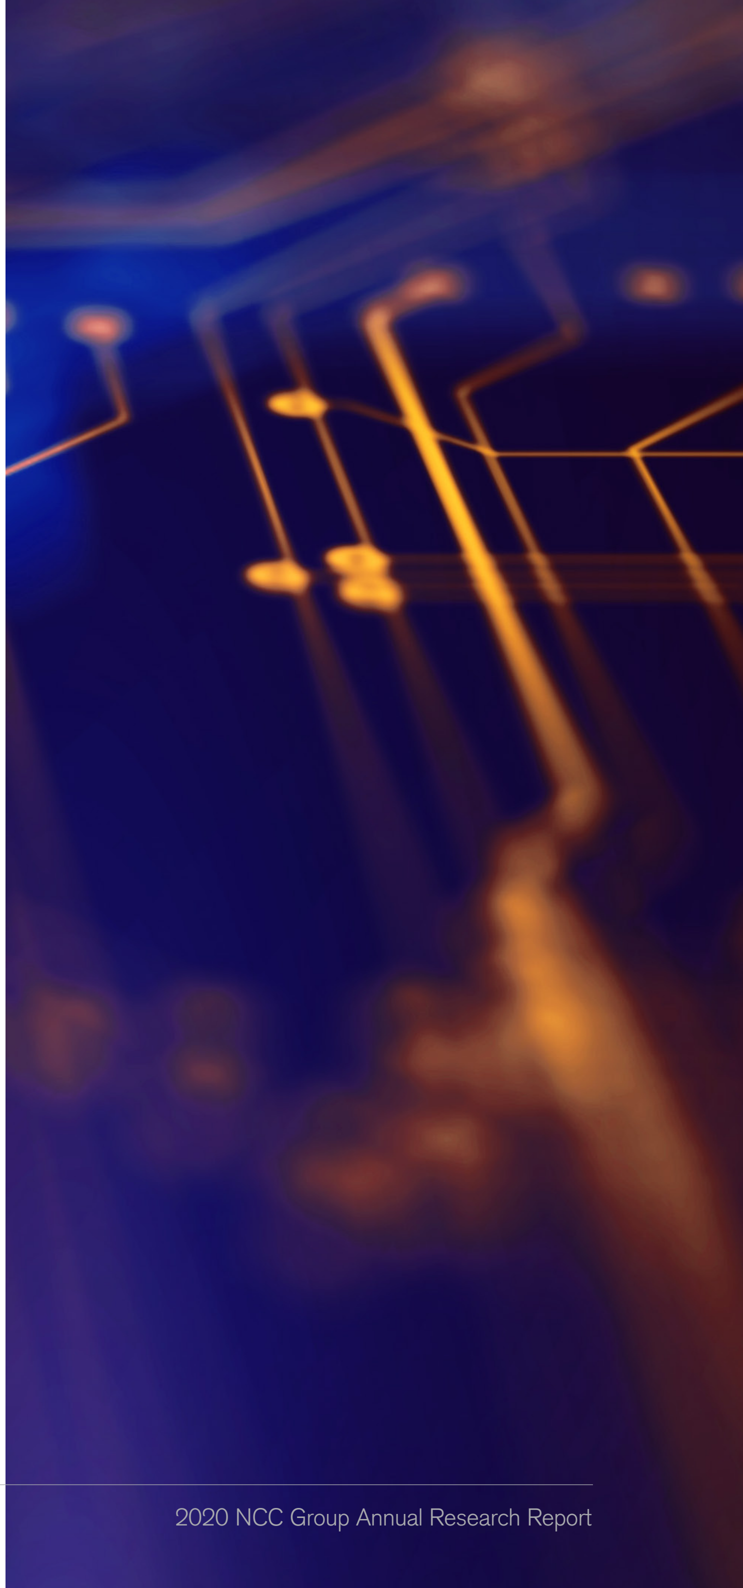

Partway through 2020, UK independent consumer champion Which? magazine reached out to us and asked if we could assist investigating the security of a series of popular domestic IoT devices and to perform a vulnerability assessment of each device. In late 2020, we published the first of this series of research projects, investigating the security of a number of Smart/Connected Doorbells.

For this project, we selected 11 different models of smart doorbells for sale on eBay and Amazon, including the #1 bestseller on Amazon, as well as several devices which had a large number of 5-star reviews.

Among our findings included devices that sent your unencrypted wifi network name and password to servers in China, devices that stored unencrypted video footage of your home, devices that could be fully disabled prior to attackers physically entering your home, devices that collected the exact GPS coordinates of your home, and critical vulnerabilities from which an attacker could compromise your entire home network pivoting from your doorbell, just to name a few.

Which? published an article about this research, The smart video doorbells letting hackers into your home, in December 2020, highlighting that NCC Group found high-risk vulnerabilities in all 11 of the smart doorbells tested. This work was also covered in Newsweek, Threatpost, Dark Reading, and Cyberscoop. Dark Reading even named this project one of their “Coolest Hacks of 2020”.

We further demonstrated both the risks inherent in – and methods of security by design for - the use of connected devices for safety-critical applications including Smart Cities, Connected Health, Connected Vehicles, and SCADA/Industrial Control Systems.

Smart Cities

In April 2020, we published our whitepaper, "A Blueprint for Secure Smart Cities," authored by Matt Lewis.

Here, we present a high-level blueprint for secure smart cities which includes principles of security by design, threat modelling, secure architecture, strong governance with appropriate policies and processes and various security assurance activities that support testing of discreet IoT components, edge, cloud and backend systems and complete end-to-end systems. This work was discussed in the Cyberwire Caveat Podcast episode: "Privacy: Once you give it away, it's very hard to get back" with Matt Lewis and Jennifer Fernick,
as well as on the podcast, Down the Security Rabbithole.

Simon Watson published, "Rise of the Sensors: Securing LoRaWAN Networks," in which he differentiates between various competing technologies in the Low-Power Wide-Area Network (LPWAN) space, including Sigfox, LTE-M and NB-IOT and LoRaWAN, ultimately taking us through properties of typical LoRaWAN deployments in real-world scenarios, and enumerating key security challenges including cryptographic key management, session management, metadata privacy issues, and challenges

to the deployment of secure embedded systems for smart cities.

In addition to his LoRaWAN research, Simon Watson also offered a Nullcon Masterclass on Common Embedded Device Vulnerabilities.

Connected Health

NCC Group began research into connected health a few years ago, publishing our first Connected Health whitepaper by Katharina Sommer, Katy Winterborn, Matt Lewis, and Stuart Kurutac in July 2019. In this paper, we outlined the current state of the Connected Health landscape, discussed a range of key themes, outlined existing standards and regulation for connected health, and offered future-thinking research directions related to the security implications of the inevitable convergence between connected health and both artificial intelligence and ambient computing.

This year, our interest in connected health evolved to include building substantial research testbeds and exploring a number of protocols and standards within health informatics.

This year, Stuart Kurutac built out the first phase of our internal Connected Medical Research Labs, which will ultimately facilitate critical research and testing in a complex, clinically-inspired environment that eliminates the risk of patient harm.

The implementations within the labs include HL7 v2 (clinical messaging standard) interfaces, an EMR (Electronic Medical Record - patient data) with an associated database, 2 DICOM (medical imaging) servers and an integration engine, with a continually growing set of additional integrated components. He has also focussed on generating test data and the development and/or integration of multiple security testing tools for connected health environments, and has been developing broad testing capabilities across these environments.

Stuart also led a project that resulted in a number of vulnerabilities in open source medical applications being discovered and reported

In May 2020, we also discussed supply chain security in healthcare environments with Dark Reading.

In September 2020, Stuart Kurutac presented some of his connected health research with CENSIS (Scotland's Innovation Centre for sensing, imaging and Internet of Things (IoT) technologies), at their IoT Cyber Challenge Program Launch.

We have been seeking to more closely demonstrate the link between patient safety and cybersecurity, discussed in this introductory blog post. In December 2020, Stuart Kurutac published, "The patient safety impact of network infrastructure vulnerabilities," in partnership with external collaborators from The AbedGraham Group. This high-level whitepaper, written for healthcare IT decision-makers, discusses the use of IoT systems in healthcare settings, and explores some specific potential use cases. In this work, they advocate for hospital IT leaders to pay greater attention to clinical network infrastructure - particular for mission-critical communications and services, emphasizing that patient safety risk is no longer primarily the domain of connected medical devices.

Connected Vehicles

In June 2020, Andy Davis published a research paper titled, "Cyber Security of New Space" **(Fig 3)** in the International Journal of Information Security with collaborators from the Surrey Center for Cyber Security and Surrey Space Center.

This paper analysed past satellite security threats and incidents to assess the motivations and characteristics of adversarial threats to satellites, finding that ground and radio frequency communications were the most favoured targets, but that the boom of satellite constellations in the upcoming years may shift this focus towards the space segment whose security must be addressed by future work in research and in the industry. They also discussed emerging security challenges and key technologies which are advancing and innovating the space and satellite industry, offering a roadmap toward securing the satellite industry.

In October 2020, Eric Evenchick presented An Introduction to Automotive Security in 2020 at SecTor. In this presentation, he introduced security generalists to concepts including vehicle networks, vehicle attack surfaces, and common vulnerabilities in vehicle systems, the changing landscape of vehicle electronics, and practical advice for getting started with hacking cars.

We also discussed ransomware in maritime environments with Maritime Logistics Professional.

(Fig 3)

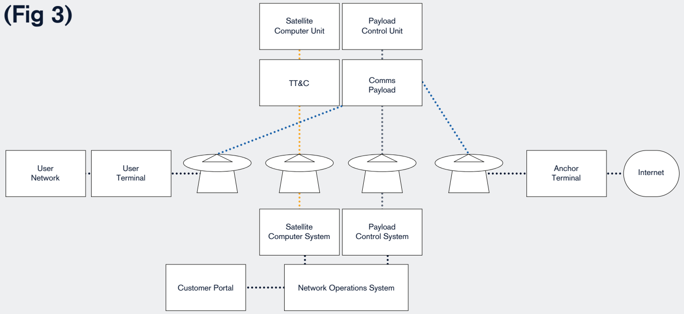

Ground segment architectures. Red lines represent satellite TT&C communications, green lines represent payload control communications, and blue lines represent the link between two users of the communication service. Black lines represent terrestrial networking links

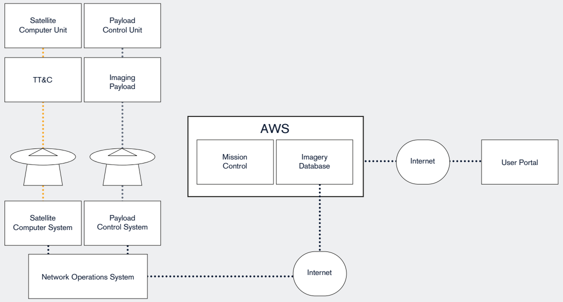

Ground segment architecture for web-based access e.g Planet [75].

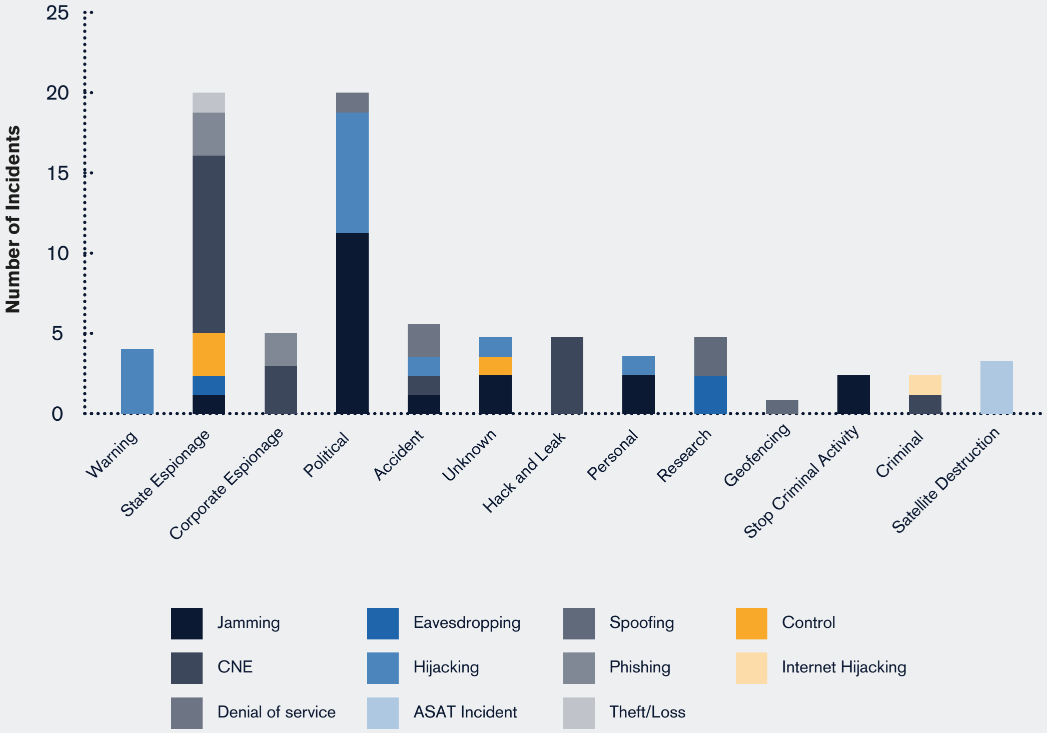

Motivation of satellite incidents and a breakdown of the techniques used

SCADA, ICS, and the Industrial Internet

For the purposes of experimentation, we created a "Nuclear Reactor on a Bench" research platform, to simulate critical connected systems within nuclear facilities to create a safe hardware testbed in which to explore cybersecurity risks in connected energy environments. This work, conducted by Barden Winkle, involved building an initial SCADA research platform using a number of different PLC manufacturers based on a real world application including HMIs. The main component manufacturers include Siemens, Schneider Electric, OMRON, Allen-Bradley, ABB and Mitsubishi (drive controller), and the platform makes use of standard industrial protocols for further research (i.e. ProfiNet, MODBUS, EtherNet/IP, etc), and provides a platform to write detailed testing methodologies. **(Fig 4)**

In Damon Small's presentation, The evolution of the Purdue Model: Integrating Industrial and Business Networks at Hou.Sec.Con, he described the evolution in recent years of the Purdue Reference Model in Industrial Control Systems (ISC) and Operational Technology (OT) security, which draws the security boundaries between users, ICS networks, and business networks, and showed the ways in which these boundaries have blurred in recent years, necessitating a new approach to thinking about ICS/OT security.

In February 2020 Threatpost webinar, Damon Small was a panelist discussing the union of IoT and Operational Technology. In this presentation, they discussed that since industrial Internet of Things (IoT) rollouts include things like networked remote sensors, mobile devices and smart machines that need advanced security and monitoring, these modernized, connected factory floors and smart cities can no longer rely on legacy OT security approaches.

They explored how this industry shift is creating new business and security challenges, as the IoT brings IT and OT together, and offered practical advice on how operational technology and traditional IT professionals are working together to tackle new security challenges presented by an onslaught of smart equipment, business automation and connected devices.

In October 2020, Damon Small also taught a cybersecurity workshop at Oilcomm on Protecting Assets in the Home Office and the Field, aimed at professionals across the oil and gas sector.

Throughout the year, Damon Small also discussed a range of cyberattacks on the energy industry in Energy Intelligence, The Energy Daily, and Oilman Magazine, and in June 2020 both Damon Small and Ollie Whitehouse discussed Trump's ban on foreign bulk power equipment and threats posed to the power grid in POWER Magazine, Utility Dive, and Greentech Media. We also discussed the difficulty of setting baselines when monitoring ICS security with Dark Reading.

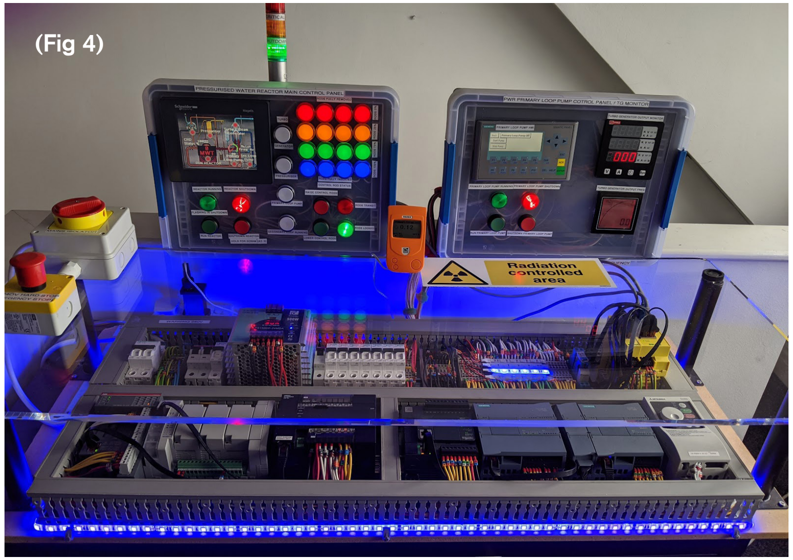

(Fig 4)

Reducing Vulnerabilities at Scale

We recognize that our industry requires scalable methods to prevent, detect, and mitigate vulnerabilities. Our research this year has included a multidisciplinary approach to this problem, including work into integrating security into devops workflows and creating tooling to support agile delivery of more-secure production software, alongside programming languages research and standards development to prevent specific bug classes and the study and experimentation of new tooling and program analysis techniques. We have also utilised considerable effort to create a more comprehensive threat model for the open source ecosystem, and identify scalable opportunities to offer indicators of security risk, help prioritize high-value pieces of internet infrastructure upon which there are a high number of transitive dependencies, and continued experimentation and development of tooling to identify vulnerabilities in source code and software binaries.

Securing the open source ecosystem

Launched in August 2020, NCC Group (together with GitHub, Google, IBM, JPMorgan Chase, Microsoft, OWASP Foundation and Red Hat) was a Founding Member of the Open Source Security Foundation (openssf.org), an initiative within the Linux Foundation which “brings together the industry’s most important open source security initiatives and the individuals and companies that support them.” At time of writing, the OpenSSF has six active working groups, including:

- **Identifying Security Threats:** “to enable stakeholders to have informed confidence in the security of open source projects… by collecting, curating, and communicating relevant metrics and metadata from open source projects and the ecosystems of which they are a part.”
- **Security Tooling:** “to provide the best security tools for open source developers and make them universally accessible... where members can collaborate together to improve upon existing security tooling and develop new ones to suit the needs of the broader open source community.”
- **Best Practices for Open Source Developers:** “dedicated to raising awareness and education of secure code best practices for open source developers.”
- **Vulnerability Disclosures:** “to help improve the overall security of the open source software ecosystem by helping mature and advocate well-managed vulnerability reporting and communication.”
- **Digital Identity Attestation:** “to enable open source maintainers, contributors and end-users to understand and make decisions on the provenance of the code they maintain, produce and use.”
- **Securing Critical Projects:** to help “allocate resources to secure the critical open source projects we all depend on.”

Jennifer Fernick currently serves on both the Governing Board and Technical Advisory Council of this initiative, and co-led the initial work within the Identifying Security Threats working group to create extensible security dashboards for open source projects to help illustrate the relative security posture of a given repository through the aggregation and interpretation of a number of key metrics and properties of a given codebase.

OpenSSF
OPEN SOURCE SECURITY FOUNDATION

We contributed to the OpenSSF whitepaper, "Threats, Risks, and Mitigations in the Open Source Ecosystem" led by Michael Scovetta of Microsoft. This paper seeks to offer an end-to-end analysis of the range of depth of security threats to open source projects, both as an educational resource to open-source maintainers, as well as a map of the territory to inspire further technical work in specific priority areas to more comprehensively help secure the open source ecosystem.

We also began a multi-stakeholder project on improving coordinated vulnerability disclosure with open source software projects and beyond, some of which was discussed in this blog post by our collaborators at the GitHub Security Lab. Further progress in this work will be discussed and released throughout 2021.

We were also excited to see our colleagues at OpenSSF announce a free, 3-course series on EdX on Secure Software Development Fundamentals. These courses teach requirements, design, and reuse, implementation, and more advanced topics including verification, threat modeling, and formal methods, and include them here as a resource for interested readers.

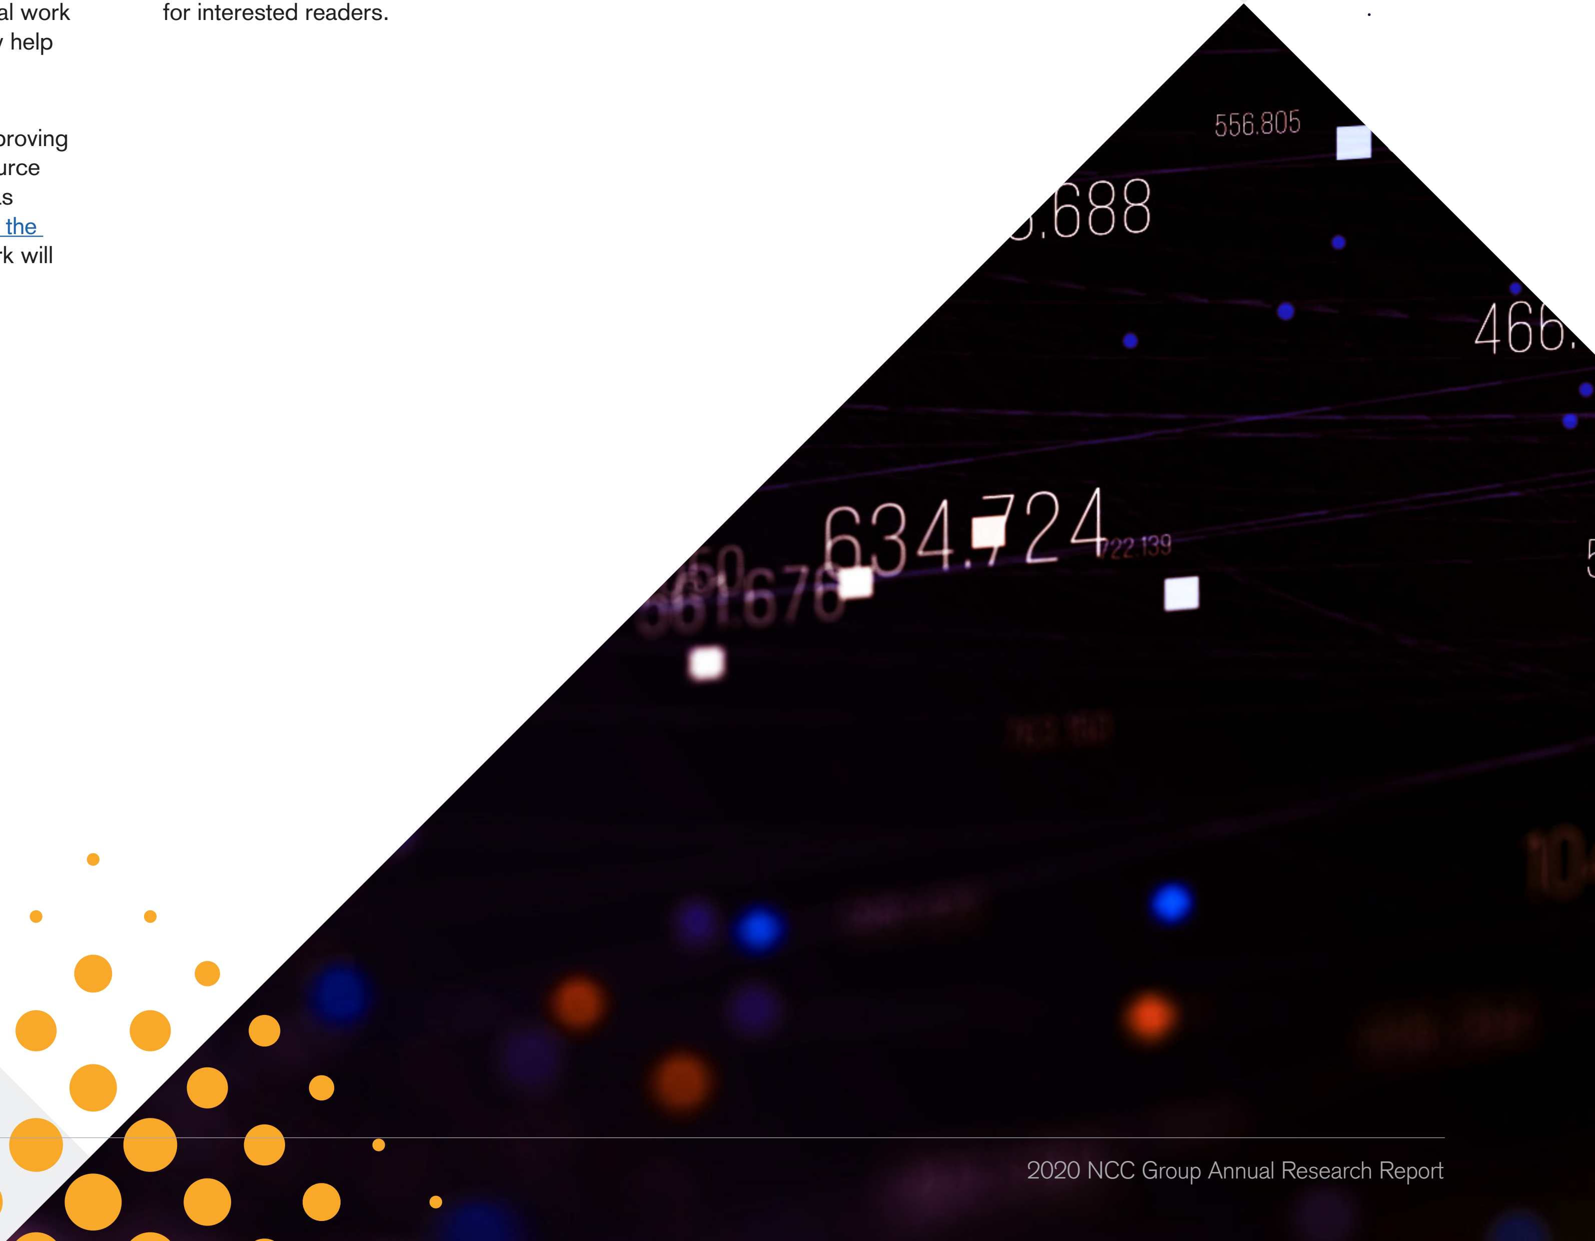

Programming languages and writing [more-]secure code

In our ongoing work on improving the ability of developers to write more secure code, Robert Seacord has continued his long-running work at the C Standards Committee, ISO/IEC JTC1[6]/SC22[7]/WG14[8] international standardization working group for the programming language C. He has also been asked as editor of "ISO/IEC TS 17961:2013 Information technology — Programming languages, their environments and system software interfaces — C secure coding rules" to prepare a version for publication as an International Standard (IS). He blogged for a more general audience about his efforts to improve software security through C language standards in February 2020.

He has developed several versions of a proposal for specific width length modifiers for C, which at time of writing have strong committee support and are undergoing minor final revision prior to standardization.

Robert Seacord has also developed a 57-page proposal titled "Defer Mechanism for C" with a group of collaborators from the C Standards Committee and presented it at the WG14 October 2020 meeting. The paper was well received and supported, and the committee agreed to create a new work item to develop an ISO/IEC Technical Specification for a Defer Mechanism which is in progress. This proposal aims to introduce a deferral of resource-freeing to the C programming language, inspired by the defer mechanism in Go.

He also looked at existing code in various open source systems and showed how they could be rewritten using various design options being considered for the Defer mechanism.

In addition, Robert Seacord co-authored an academic paper, C language mechanism for error handling and deferred cleanup[9], with Jens Gustedt from INRIA and the Université de Strasbourg that has been accepted at the 36th ACM/SIGAPP Symposium on Applied Computing (SAC '21), and has begun work on a base document for the ISO/IEC technical specification.

Robert Seacord also published a new book, titled, "Effective C: An Introduction to Professional C Programming" (No Starch Press, 2020). Coinciding with the book's release, he also published a blog post on the past, present, and future of effective C, in which he described the history of the C language and C language standardization, surveyed some of the exciting innovations being worked on today particularly with respect to security, and described some C language and library features that are on the radar for C23.

He also published an article about Effective C in IEEE Computer[10], which describes why the C programming language has succeeded, and what's next for the language from the perspective of a long-term C Standards Committee expert.

Beyond this, we have also undertaken deeper research related to potential weaknesses within memory-safe languages, and intend to publish our findings throughout 2021.

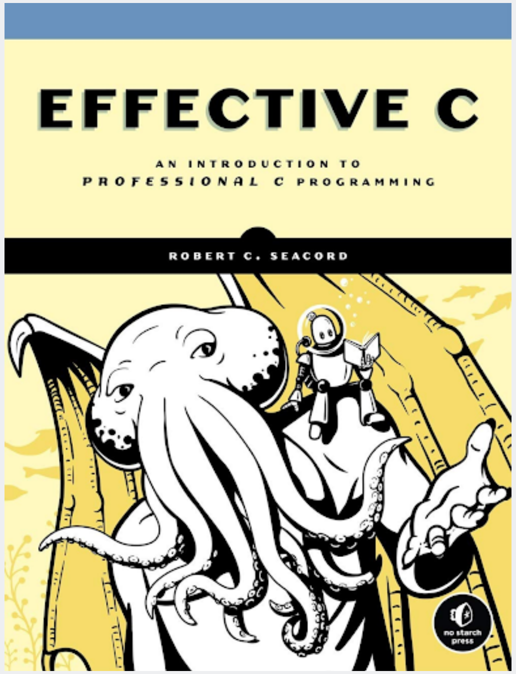

[6] https://www.iso.org/isoiec-jtc-1.html

[7] http://www.open-std.org/JTC1/SC22/

[8] http://www.open-std.org/jtc1/sc22/wg14/

[9] Jens Gustedt, Robert Seacord. C language mechanism for error handling and deferred cleanup. The 36th ACM/SIGAPP Symposium on Applied Computing (SAC'21), Mar 2021, virtual, South Korea.

[10] R. C. Seacord, "Effective C," in Computer, vol. 53, no. 11, pp. 79-82, Nov. 2020, doi: 10.1109/MC.2020.3016369.

Continuous delivery of security

In early 2020, Clint Gibler presented his DevSecOps State of the Union v2.0 at both RSA Conference and AppSec Cali. In this presentation, he offered attendees a distillation of the unique tips and tricks, lessons learned, and tools discussed in dozens of blog posts and more than 50 conference talks over the past few years, combined with knowledge gained from in-person discussions with security leaders at companies with mature security programs. Conclusions from this work formed 3 of TechBeacon's 5 key takeaways from RSA Conference 2020.

In February 2020, he presented his panel, Lessons Learned from the DevSecOps Trenches, at BSidesSF alongside external collaborators Zane Lackey (Signal Sciences), Astha Singhal (Netflix), Justine Osborne (Apple), and Doug DePerry (Datadog). This panel offered a frank discussion with security team leads at several forward-thinking companies on how they've built and scaled their security programs: what worked, what failed, and more.

Clint Gibler also presented How to 10X Your Company's Security (Without a Series D), where he summarized and distilled the insights, unique tips and tricks, and actionable lessons learned from a vast number of DevSecOps/modern AppSec talks and blog posts, showing where we've been, and where we are going. Security Boulevard published some of the top takeaways from this presentation, including automating as much of your security as possible, treating your company's developers as your customers, centralized vulnerability management, and the importance of continuous vulnerability scanning, and maintaining an inventory of IT assets.

In March 2020, Adam Rudderman of NCC Group's Bug Bounty Practice presented at Nullcon Goa, discussing where bug bounty fits into an organization's security program.

Research and Intelligence Fusion Team (RIFT)

NCC Group's Research and Intelligence Fusion Team (RIFT) is led by Christo Butcher. Below, we discuss a number of publications by RIFT, including on pandemic-related threat intelligence, our dissection and early reporting of the WastedLocker ransomware, the noted exploitation of Citrix's ADC vulnerabilities, and our in-the-wild observations about the remote code execution vulnerability in the F5 Big-IP administrative interface (CVE-2020-5902, disclosed July 2020), among other things including a profile of the threat actor commonly known as TA505.

In the third week of March 2020, we published a piece on our observation of global cyber-threat actors' early attempts to capitalize on current events and global anxiety related to the escalating pandemic. Ultimately, we found that there was not an increase of threat actors, the tactics of threat actors were changing, with many campaigns starting to leverage user interest in COVID-19 information as a means to distribute malware.

In April 2020, we blogged about our recent IETF draft (co-written with the UK National Cyber Security Center), titled "Indicators of Compromise (IoCs) and Their Role in Attack Defence". This article provides a draft that will help converge on the formalization as well as the definition, role, and value of indicators of compromise (IoCs), and "outlines the different types of IoC, their associated benefits and limitations, and discusses their effective use. It also contextualises the role of IoCs in defending against attacks through describing a recent case study."

In early June 2020, we wrote an in-depth analysis of the new Team9 malware family. Team9 is a new piece of malware that is being developed by the same group that created Trickbot. Researchers have already evaluated two main components of the malware thus far, the loader and the backdoor, and discuss details in the associated blog post.

In June 2020, Nikolaos Pantazopoulos, Stefano Antenucci, and Michael Sandee - in close collaboration with NCC's RIFT - published WastedLocker: A New Ransomware Variant Developed By The Evil Corp Group. This post discusses a new ransomware variant called "WastedLocker" which began popping up in May 2020. It is believed to have been created by the Evil Corp Group, perhaps best previously known due to their association with Dridex malware and BitPaymer ransomware. The ransomware is distributed via the SocGholish fake update framework, and it is used to distribute a custom Cobalt Strike loader. WastedLocker is a sophisticated piece of malware protected by a custom crypter referred to as CryptOne and will perform self-checks to ensure that the ransomware will run properly. In this post, we discuss the known history of Evil Corp, and offer technical analysis of WastedLocker, as well as some IoCs for research purposes, alongside an explanation about why IoCs for targeted ransomware have limited utility in defense, due to the bespoke nature of the build configurations per viction. This work was covered by Wired, Barron's, Dark Reading, SC Media, Security Boulevard, International Business Times, BankInfoSecurity, Yahoo! News, and Cyber Defense Magazine. Additional reporting by Dark Reading discussed how dozens of US newspaper websites were compromised and used to host malicious code used in the distribution of the associated ransomware.

In early July 2020, we published our whitepaper, Thematic for Success in Real-World Offensive Cyber Operations – How to make threat actors work harder and fail more often. In this work, we outlined the key Red Team techniques that continue to enable us to breach the largest and most sophisticated organisations on the planet, with each attack stage mapped against the MITRE ATT&CK framework. In publishing this, we hope to assist organisations in prioritising their security activities against contemporary offensive actors.

On July 10th 2020, NCC Group's RIFT published RIFT: Citrix ADC Vulnerabilities CVE-2020-8193, CVE-2020-8195 and CVE-2020-8196 Intelligence. This vulnerability went from disclosure to active, in-the-wild exploitation within less than 4 days, the timeline of which is enumerated in our post, which served as a summary of what the RIFT team knew as the situation unfolded. These consequences resulted from the July 7th disclosure by Citrix of a number of vulnerabilities in the application delivery controller (CVE-2020-8193, CVE-2020-8195, CVE-2020-8196). As of July 10th[11], RIFT has confirmed that the vulnerability reported and exploited by Donny Maasland on July 8th can be used to extract valid VPN sessions from a vulnerable instance.

On July 5th 2020, we published some of our initial analysis of the F5 Networks K52145254: TMUI RCE vulnerability (CVE-2020-5902). CVE-2020-5902 was disclosed on July 1st, 2020 by F5 Networks in K52145254 as a CVSS 10.0 remote code execution vulnerability in the Big-IP administrative interface, ultimately allowing arbitrary, active interception of any traffic traversing an internet-exposed, unpatched Big-IP node. By July 3rd, 2020 NCC Group observed active exploitation. NCC Group's RIFT observed that the REST API was being used to execute code. On July 14th, 2020 we noticed that the threat actors were starting to deploy more complex payloads and miners. As time went on threat actors started using staged exploitation and web shells, and were able to bypass mitigation attempts. The exploitation of this vulnerability provided the opportunity for threat actors to gain credentials, license keys, and private keys to any SSL/TLS certificates on these load balancing devices. Wired initially reported on this research in their July 6th article, Hackers Are Exploiting a 5-Alarm Bug in Networking Equipment. In this piece, they discussed the spike in RCE attempts we noticed against our honeypots exposing these vulnerabilities on the Sunday of the July 4th weekend, as well as into Monday. This further supports US CISA director Chris Krebs' July 5th tweet, "If you didn't patch by this morning, assume compromised"[12]. This research was also reported in Dark Reading, Threatpost, Decipher, BankInfoSecurity, and others. SC Media UK discussed this work in the context of U.S. Cyber Command's urge to organizations to quickly implement F5's patch for the vulnerability.

One week later, on July 12 2020, we wrote a guide to understanding the root cause of F5 Networks K52145254: TMUI RCE vulnerability (CVE-2020-5902).

Timeline from disclosure to exploitation for CVE-2020-8193, 8195 and 8196

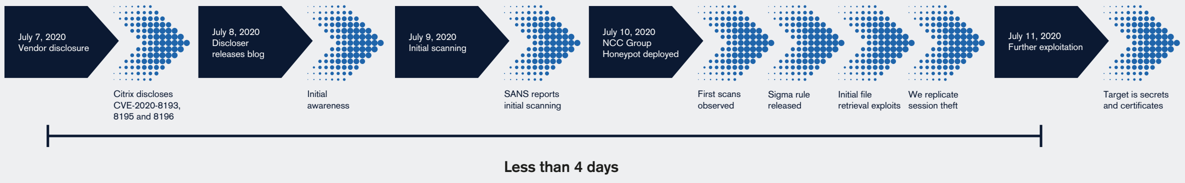

Timeline from disclosure to exploitation for CVE-2020-8193, CVE-2020-8195, and CVE-2020-8196

[11] https://dmaasland.github.io/posts/citrix.html

[12] https://twitter.com/CISAKrebs/status/1279851810094800902

In October 2020, RIFT released 3 months' worth of honeypot web traffic data showing attacks and exploitation of F5 Big-IP and a small amount of Citrix ADC exploitation. In this data, we observed that the threat actors exploited the vulnerable devices and did things like password dumps, the creation of multiple backdoors, the collection of various sensitive files, and deployment of bitcoin miners and botnet loaders, and emphasized the particular time-sensitivity of patching the F5 vulnerability to resist adversarial compromise of affected devices.

In November 2020, we published a profile on the cyber-threat actor known as TA505, a sophisticated, self-subsisting, threat actor who has engaged in multiple targets across multiple sectors for financial gain. TA505's simple but effective attack strategy is to "encrypt a corporate network with ransomware, more specifically the Clop ransomware strain, and demand a ransom in Bitcoin to obtain the decryption key." The group was able to succeed by utilizing a GetandGo-SDBbot campaign, consisting of a simple loader that gathered system information, beaconing, and command execution. We discuss how TA505 was able to stay under the radar for long periods of time, as well as qualities observed related to their working hours and other signatures of their attack methods.

TA505 Infection Chain

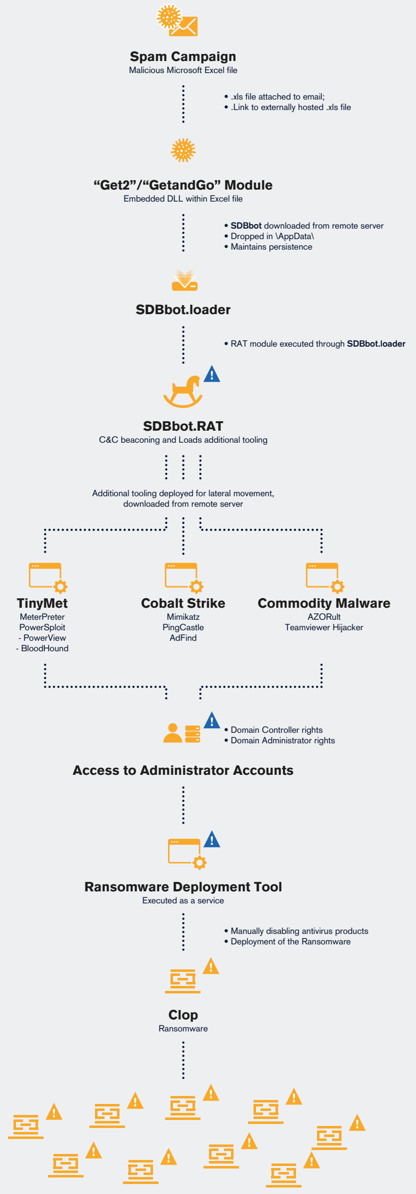
Spam Campaign
Malicious Microsoft Excel file
• .xls file attached to email;
• .Link to externally hosted .xls file
"Get2"/"GetandGo" Module
Embedded DLL within Excel file
• SDBbot downloaded from remote server
• Dropped in \AppData\
• Maintains persistence
SDBbot.loader
• RAT module executed through SDBbot.loader
SDBbot.RAT
C&C beaconing and Loads additional tooling
Additional tooling deployed for lateral movement, downloaded from remote server
TinyMet
MeterPreter
PowerSploit
- PowerView
- BloodHound
Cobalt Strike
Mimikatz
PingCastle
AdFind
Commodity Malware
AZORult
Teamviewer Hijacker
• Domain Controller rights
• Domain Administrator rights
Access to Administrator Accounts
Ransomware Deployment Tool
Executed as a service
• Manually disabling antivirus products
• Deployment of the Ransomware
Clop
Ransomware

TA505 Infection Chain

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 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 recent discussion of how in 2017, we unleashed our version of NotPetya on global commodities trading firm, Trafigura.

In June 2020, EDG published Striking Back at Retired Cobalt Strike: A look at a legacy vulnerability. This work looks at some of the communication and encryption internals of Cobalt Strike between Beacons and the Team Server in the 3.5 family, then exploring the subsequent exploitation of a vulnerability in Cobalt Strike 3.5 from 2016 to achieve remote unauthenticated code execution on the Team Server. We chose this type of archaeological bug hunting to illustrate to blue teams the encryption fundamentals underpinning Cobalt Strike, as well as to give some insight into detection engineering. For red teams, the exploitation of this vulnerability emphasizes the importance of hardening command and control infrastructure.

Throughout the year, EDG published a series of blog posts on CVE-2018-8611 Exploiting Windows KTM. This five-post series introduced CVE-2018-8611 and provided background information about the Windows Kernel Transaction Manager (KTM), performed patch analysis and basic triggering of the vulnerability with an assisted race condition, discussed triggering the race condition, and outlined some debugging tricks, and outlined how to then build an arbitrary read/write primitive that leads us to privilege escalation and presented a demo showing exploitation on an unpatched Windows 10 1809 x64 from approximately October 2018. The 5-post series concluded with their discussion of how the exploit they ended up writing differed from the exploit found in-the-wild in 2019, and discussed the future of KTM. This research was presented at POC (Power of Community) in 2019, s well as at OffensiveCon 2020 (slides, video).

EDG also released their open source tool, idahunt, which enables one to automate IDA Pro from the command line.

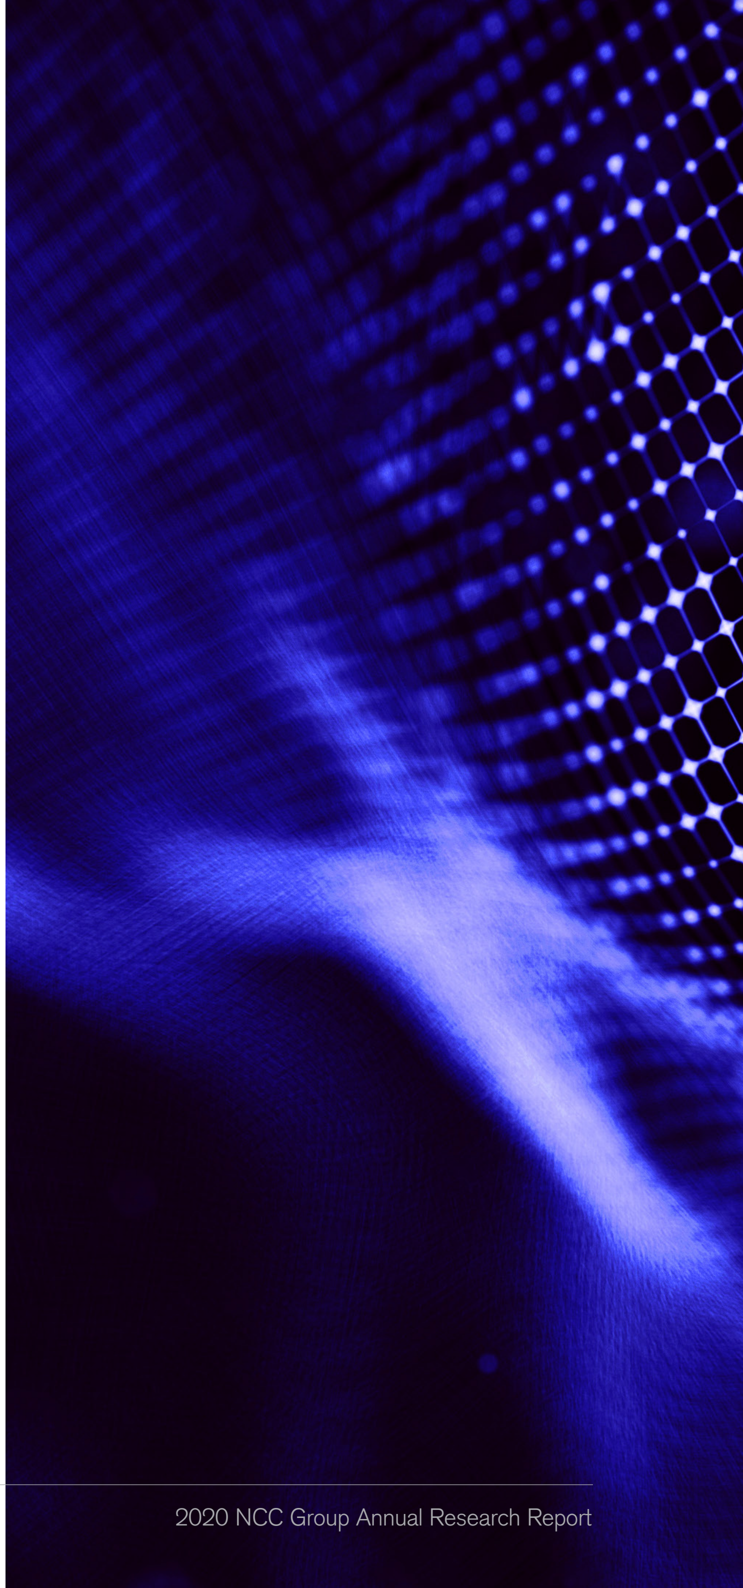

Other Research

Dirk-Jan Mollema presented Walking Your Dog in Multiple Forests – Breaking AD Trust Boundaries through Kerberos Vulnerabilities at Black Hat Asia and Hack in the Box Amsterdam. This research introduced a vulnerability in Kerberos and forest trusts that allows attackers to break the trust within Active Directory forest trusts in a new way. The presentation began with technical details on how Kerberos works over forest trusts and how the security boundary is normally enforced, then discussed a flaw in how AD forest trusts operate and how this 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), accompanied by a proof-of-concept and a demonstration of abusing the vulnerability.

Chris Nevin published Carnivore, his tool for assessing on-premises Microsoft servers such as ADFS, Skype, Exchange, RDWeb, and O365, and spoke about its' capabilities at both Black Hat USA 2020 and DEF CON Demo Labs.

Dhruv Verma, Michael Roberts, and Kuan Xiang Wen presented their Bad Active Directory (BAD) training at the DEF CON 28 Wall of Sheep/Packet Hacking Village.

Michaal Gough also gave a number of presentations about improving incident response.

Sanne Maasakkers presented on how to improve security awareness campaigns by applying phishing research at the Grace Hopper Celebration of Women in Computing.

Richard Appleby wrote an overview of cheating and anti-cheat methods in electronic games, offering a history and basic typology of game cheats and anti-cheat methods.

Sourya Biswas from NCC Group's Risk Management and Governance Practice gave a number of conceptual presentations on the psychology of phishing attacks, what can be learned from historical conflicts about the practice of cybersecurity, and a comparison of Wall Street banks vs Silicon Valley technology companies' approaches to enterprise security.

Roger Meyer and Gérald Doussot published a blog post on the impact of DNS over HTTPS (DoH) on DNS Rebinding Attacks. This work was built on top of their prior work on DNS rebinding attack and prevention techniques presented at BSidesLV and DEF CON in 2019. They had also released a tool, Singularity of Origin, to simplify the exploitation of vulnerable services. In this new work, they ultimately conclude that using DoH instead of plaintext DNS does not have an impact on DNS rebinding attacks - they were able to successfully perform all DNS rebinding attack strategies implemented in Singularity of Origin when targeting a vulnerable service, even when DNS-over-HTTPS is used.

Phillip Langlois and Edward Torkington published details about CVE-2019-1381 and CVE-2020-0859 in their blog post, How Misleading Documentation Led to a Broken Patch for a Windows Arbitrary File Disclosure Vulnerability. This post builds on top of their November 2019 blog post covering an elevation-of-privilege vulnerability they uncovered in Windows whilst conducting research into Windows Component Object Model (COM) services, both to discuss the intrinsically compelling root cause of this particular vulnerability, as well as to discuss how Microsoft's initial attempt to patch the vulnerability highlighted both the difficulty of successfully fixing this kind of bug and, perhaps of even more interest, some inaccuracies in Microsoft documentation that may have implications for other operating system components.

Matthew Pettit published on how order details screens for certain online shopping platforms may inadvertently leak personally identifiable information.

Simon Palmer published Crave the Data: Statistics from 1,300 Phishing Campaigns. In this post, he studied 1300 phishing campaigns targeting over 360,000 users, finding that across sectors, half of all victims who clicked a phishing link would supply credentials, but that attack success rates varied by sector - for example, victims from charitable organizations were 3x as likely to click a phishing link than victims from healthcare organizations.

Tanner Prynn published Code Patterns for API Authorization: Designing for Security, which describes some of the most common design patterns for authorization checking in web application code, and makes comparisons between the design patterns to help understand when each pattern makes sense as well as the drawbacks of the pattern, to help developers and architects understand what the different code patterns look like and how to choose between them, and to help security auditors to identify the most effective approaches to auditing authorization controls.

Balazs Bucsay wrote about Shell Arithmetic Expansion and Evaluation Abuse, discussing how arithmetic expansion and evaluation in Bash can be used to get a privileged shell on a Linux-based appliance that only presented a restricted shell.

Travis Knapp-Prasek presented a talk at BSidesSF titled Sans-Serif Rules Everything Around Me, which discussed how specific visual similarities in fonts can enable highly effective phishing attacks.

Lucas Rosevear disclosed a pre-authentication remote code execution vulnerability in playSMS (CVE-2020-8644).

David Cash discussed using SharePoint as a phishing platform, explaining how trusted Microsoft cloud infrastructure can be used to host a credential capture campaign.

Jesús Calderón Marín wrote a wildly-popular guide for security researchers about pentesting online casino roulette.

Richard Warren and David Cash published multiple vulnerabilities in Pulse Connect Secure including an arbitrary file read vulnerability in the pre/post logon message component, allowing an authenticated administrative user to read arbitrary files from the underlying OS (CVE-2020-8255), a Perl Template Injection vulnerability which can be exploited by an authenticated administrative user to execute arbitrary code as root (CVE-2020-8243), and an uncontrolled gzip extraction vulnerability which allows an attacker to overwrite arbitrary files, resulting in Remote Code Execution as root (CVE-2020-8260).

Joe Meyer of NCC Group's Risk Management and Governance practice presented on CCPA and the New United States Privacy Laws, to help navigate which states' privacy "laws" are actually passed, which ones are still in limbo, and what are the common criteria for compliance between them all.

Jeremy Boone of NCC Group's Hardware and Embedded Systems practice published a technical advisory on ARM MbedOS USB Mass Storage Driver Memory Corruption, outlining three memory safety vulnerabilities, allowing adversaries with physical access to corrupt kernel memory or disclose kernel memory contents.

Daniel López Jiménez presented Understanding and Hiding your Operations, a guide to opsec for red teamers, at No cON Name.

Erik Steringer added a considerable range of features to his open-source tool PMapper Principal Mapper), a script and library for identifying risks in the configuration of AWS Identity and Access Management (IAM) in an AWS account.

Jelle Vergeer published Decrypting OpenSSH sessions for fun and profit, which explores his research into OpenSSH that resulted from wondering if it was possible to decrypt the SSH session and gain knowledge of it by recovering key material from the memory snapshot, and also includes his release of a tool to dump OpenSSH session keys from memory and decrypt and parse sessions in combination with pcaps.

Andy Grant found and disclosed multiple vulnerabilities in macOS, including a local root privilege escalation bug in the macOS installer (CVE-2020-9817), which could enable a low-privileged user or process to gain arbitrary code execution with root privileges, effectively leading to a full system compromise, as well as a blog series demonstrating his finding that by using a carefully crafted calendar event, an attacker can retrieve semi-arbitrary files from a target victim's macOS system - all the victim has to do is click on an invite (CVE-2020-3882).

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, DEF CON, and countless others. In the past year alone our researchers 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, 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 us

Jennifer Fernick

SVP & Global Head of Research

jennifer.fernick@nccgroup.com

Matt Lewis

Director of Commercial Research

matt.lewis@nccgroup.com