BlueHat

BlueHat v17 General Audience Agenda

General Audience | November 8th, 2017

Track	Time	Speaker	Company	Talk Subject
Keynote	9:00 - 9:50 AM	Merike Kaeo	Farsight Security	Keynote
Track 1 – Encrypt All the Things	10:00 - 10:50 AM	Alban Diquet Thomas Sileo	Data Theorem	Where, how, and why is SSL traffic on mobile getting intercepted? A look at three million real-world SSL incidents
	11:00 - 11:50 AM	Joseph Salowey	Tableau Software	TLS 1.3 – Full speed ahead… mind the warnings – the great, the good and the bad
Track 1 – Battles in Silicon	1:00 - 1:50 PM	Alex Matrosov	Cylance	Betraying the BIOS: Where the Guardians of the BIOS are Failing
	2:00 - 2:50 PM	Niek Timmers Cristofaro Mune	Riscure B.V. & Independent Embedded Security Consultant	KERNELFAULT: R00ting the Unexploitable using Hardware Fault Injection
	3:00 - 3:50 PM	Rob Turner	Qualcomm Technologies	Raising the Bar: New Hardware Primitives for Exploit Mitigations
	4:00 - 4:50 PM	Gunter Ollmann	Microsoft	Extracting Secrets from Silicon – A New Generation of Bug Hunting
Track 2 – Hey Microsoft, you got it wrong!	10:00-10:50 AM	Casey Smith	Red Canary	You Are Making Application Whitelisting Difficult
	11:00-11:50 AM	Yong Chuan Koh	MWR Infosecurity	Corrupting Memory in Microsoft Office Protected-View Sandbox
Track 2 – Advancing products meet the new threats	1:00 - 1:50 PM	Saruhan Karademir David Weston	Microsoft	Securing Windows Defender Application Guard
	2:00 - 2:50 PM	Mark Wodrich Jasika Bawa	Microsoft	Mitigations for the Masses: From EMET to Windows Defender Exploit Guard
	3:00 - 3:25 PM	Dean Wells	Microsoft	Don't Let Your Virtualization Fabric Become the Attack Vector
	3:30 - 3:55 PM	Jonathan Birch	Microsoft	Dangerous Contents – Securing .Net Deserialization
	4:00 - 4:50 PM	Filippo Seracini Lee Holmes	Microsoft	Born Secure. How to Design A Brand New Cloud Platform With A Strong Security Posture

BlueHat

Talk Abstracts

Keynote | 9:00-9:50 AM | Merike Kaeo |Farsight Security

Track 1 – Encrypt All the Things | 10:00-10:50 AM | Alban Diquet and Thomas Sileo | Data Theorem
Where, how, and why is SSL traffic on mobile getting intercepted? A Look at three million real-world SSL incidents

Over the last two years, we've received and analyzed more than three million SSL validation failure reports from more than a thousand of iOS and Android apps available on the Stores, and used all around the world. From mobile banking to music apps, each report was triggered because an unknown or unexpected certificate was being served to the app, preventing it from establishing a secure connection to its server via SSL/TLS. We've analyzed each of these reports to understand what caused the SSL connection to fail, and then grouped similar failures into various classes of SSL incidents. Throughout this presentation, we will describe the analysis we've made and present our findings. First, we will provide a high-level overview of where, how, and why SSL incidents are occurring across the world for iOS and Android users, and describe the various classes of incidents we've detected. Some of these types of incidents, such as corporate devices performing traffic inspection, are well-known and understood, although we will provide new insights into how widespread they are. Then, we will take a closer look at a few notable incidents we detected, which have been caused by unexpected, or even suspicious actors. We will describe our investigations and what we found. Lastly, we will provide real-world solutions on how to protect apps against traffic interception and attacks, as a mobile developer.

Track 1 – Encrypt All the Things | 11:00-11:50 AM | Joseph Salowey | Tableau Software
TLS 1.3 – Full speed ahead… mind the warnings – the great, the good and the bad

Transport Layer Security (TLS) 1.3 is almost here. The protocol that protects most of the Internet secure connections is getting the biggest ever revamp, and is losing a round-trip. We will explore differences between TLS 1.3 and previous versions in detail, focusing on the performance and security improvements of the new protocol as well as some of the challenges we face around securely implementing new features such as 0-RTT resumption.

Track 1 – Battles in Silicon | 1:00-1:50 PM | Alex Matrosov | Cylance
Betraying the BIOS: Where the Guardians of the BIOS are Failing

This presentation is meant to serve as an alarum for hardware vendors; BIOS-level security researchers and defenders; and sophisticated stakeholders who want to know the current state of UEFI exposure and threats. The situation is serious but, with the right tools and knowledge, we can prevail. Hardware vendors such as Intel have introduced new protection technologies like Intel Boot Guard (since Haswell) and BIOS Guard (since Skylake). Boot Guard protects Secure Boot's "Root of Trust" from firmware-based attacks by verifying that a trusted UEFI firmware is booting the platform. When BIOS Guard is active, only guarded modules can modify SPI flash memory; this can protect from persistent implants. Both technologies run on a separate CPU known as the "Authenticated Code Module" (ACM), which isolates them from attackers and also protects from race condition attacks. Those "Guard" technologies are sometimes referred to as UEFI rootkit killers. Not many details are publicly available regarding these technologies. In this presentation, I will discuss particular implementations on hardware with the most recent Intel CPUs such as Skylake and Kaby Lake. Most of the information has been extracted from UEFI firmware modules by reverse engineering. This DXE and PEI modules cooperated with ACM-code for enabling, configuration and initialization. This talk will also cover some weaknesses of those guards. Where are the BIOS guardians failing? How difficult is it to bypass these protections and install a persistent rootkit from the operating system?

Track 1 – Battles in Silicon | 2:00-2:50 PM | Niek Timmers and Cristofaro Mune | Riscure B.V. & Independent Embedded Security Consultant
KERNELFAULT: R00ting the Unexploitable using Hardware Fault Injection

Fault injection attacks have been historically perceived as high-end attacks not available to most hackers. They used to require expensive tooling and a mysterious mix of skills which resulted them being out of reach for even the most skilled attackers. These days are over as low-cost fault injection tooling is changing the capabilities of the hacking masses at a rapid pace. Historically, fault injection attacks are used to break cryptographic implementation (e.g. Differential Fault Analysis) or bypassing security checks like performed by a pin verification function. However, nothing prevents them to be used on richer systems like embedded devices or IoT devices. Fault injection attacks can be used to change the intended behavior of hardware and software, due, among the others, to corrupted memory reads and instructions execution. In this talk we show that fault injection attacks and, more specifically, voltage fault injection, allow escalating privileges from an unprivileged context, in absence of logically exploitable software vulnerabilities. This is demonstrated using practical examples where the control flow of the Linux kernel is influenced in order to gain root privileges. The impacts of a new attack that allows for privileged code execution across security boundaries, bypassing all current software-based fault injection countermeasures, are analyzed, along with possible mitigations. All provided examples apply to a fully patched Linux operating system, executed by a fast and feature rich System-on-Chip.

Track 1 – Battles in Silicon | 3:00-3:50 PM | Rob Turner | Qualcomm Technologies
Raising the Bar: New Hardware Primitives for Exploit Mitigations

Almost three decades since the Morris worm and we're still plagued by memory corruption vulnerabilities in C and C++ software. Exploit mitigations aim to make the exploitation of these vulnerabilities impossible or prohibitively expensive. However, modern exploits demonstrate that currently deployed countermeasures are insufficient. In ARMv8.3, ARM introduces a new hardware security feature, pointer authentication. With ARM and ARM partners, including Microsoft, we helped to design this feature. Designing a processor extension is challenging. Among other requirements, changes should be transparent to developers (except poor compiler developers), support both system and application code, interoperate with legacy software, and provide binary backward compatibility. This talk discusses the processor extension and explores the design trade-offs, such as the decision to prefer authentication over encryption and the consequences of small tags. Also, this talk provides a security analysis, and examines how these new instructions can robustly and efficiently implement countermeasures.

Track 1 – Battles in Silicon | 4:00-4:50 PM | Gunter Ollmann | Microsoft
Extracting Secrets from Silicon – A New Generation of Bug Hunting

As reverse engineering tools and hacking techniques have improved over the years, software engineers have been forced to bury their secrets deeper down the stack – securing keys and intellectual property first in software, then drivers, on to custom firmware and microcode, and eventually as etchings on the very silicon itself. For the hackers involved, the skills and tooling needed to extract and monetize these secrets come with ever increasing hurdles and cost. Yet, seemingly as a corollary to Moore's Law, each year the cost of the tooling drops by half, while access (and desire) doubles. Today, with access to multi-million dollar semiconductor labs that can be rented for as little as $200 per hour, skilled adversaries can physically extract the most prized secrets from the integrated circuits (IC) directly. Understanding your adversary lies at the crux of every defensive strategy. This session reviews the current generation of tools and techniques used by professional hacking entities to extract the magic numbers, proprietary algorithms, and

WORN (Write Once, Read Never) secrets from the chips themselves. As a generation of bug hunters begin to use such tools to extract the microcode and etched algorithms from the IC's, we're about to face new classes of bug and vulnerabilities – lying in (possibly) ancient code – that probably can't be "patched". How will we secure secrets going forward?

Track 2 – Hey Microsoft, you got it wrong! | 10:00-10:50 AM | Casey Smith | Red Canary
You Are Making Application Whitelisting Difficult

Application whitelisting is a great defense. It really is. As difficult as it may be to implement, it gives organizations a strong defense to turn the tide against malicious binaries. The trouble is, administrators often trust all things that are signed by Microsoft. And... All binaries from Microsoft are signed in the same manner. This talk seeks to be a discussion of what Microsoft signs, how these binaries can be abused, and propose new strategies to move forward. How can we discover these binaries? What are capabilities that can be abused? What should Microsoft be signing? Should the same certificate be used for all binaries emitted from Microsoft? This talk will present a recent binary discovered to bypass Device Guard as a case study.

Track 2 – Hey Microsoft, you got it wrong! | 11:00-11:50 AM | Yong Chuan Koh | MWR Infosecurity
Corrupting Memory in Microsoft Office Protected-View Sandbox

The MS Office Protected-View is unlike any other sandboxes; it aims to provide only a text-view of the document contents and therefore does not have to provide full functionalities of the application. As a result, the broker-sandbox Inter-Process Communication (IPC) attack surface is greatly reduced. However this does not mean there are no vulnerabilities. This talk will discuss the methodology for fuzzing this IPC attack surface, from the test-case generation to the discovery and analysis of CVE-2017-8502 and CVE-2017-8692.

Track 2 – Advancing Products Meet the New Threats | 1:00-1:50 PM | Saruhan Karademir and David Weston | Microsoft
Securing Windows Defender Application Guard

Windows Defender Application Guard (WDAG) brings the next generation isolation into the browser space. It merges the best of Hyper-V virtualization and Microsoft Edge sandboxing technologies to bring hardware-enforced isolation of untrusted websites from the user's data and operating system. In this talk, we will walk through the WDAG security promise and architecture. We will explain how it was built from the ground up with security as the number one priority showcasing the architectural decisions that added layers of defense. Finally, we explore how Microsoft's internal security teams engaged from the very beginning of this feature's development, helping shape WDAG's design, finding and fixing critical vulnerabilities, and building additional defense-in-depth layers before the product reached a single customer.

Track 2 – Advancing Products Meet the New Threats | 2:00-2:50 PM | Mark Wodrich and Jasika Bawa | Microsoft
Mitigations for the Masses: From EMET to Windows Defender Exploit Guard

In the Windows 10 Fall Creators Update, we introduced Windows Defender Exploit Guard (WDEG) — a feature suite that enables you to reduce the attack surface of applications while allowing you to balance security with productivity in a realistic manner. With WDEG's smart attack surface reduction (ASR) rules and exploit protection, we are looking to provide security hardening for popularly used applications without losing sight of the complex environments being managed in most organizations. But what are these security hardening options? And how do we anticipate they will be put to work? In this talk, we will discuss why and how we embarked upon the WDEG journey, starting all the way from our passionate Enhanced Mitigation Experience Toolkit (EMET) customers, through the conception of the WDEG feature set, to the internal mechanics behind the rich set of protections it offers. We will also demonstrate how WDEG's smart ASR rules and exploit mitigation settings can be used to reduce the likelihood of exploitation of commonplace legacy applications, now directly from Windows 10.

Track 2 – Advancing Products Meet the New Threats | 3:00-3:25 PM | Dean Wells | Microsoft
Don't Let Your Virtualization Fabric Become the Attack Vector

Witness a whipper-snapper of an admin conduct a series of progressively more sneaky attacks against unsuspecting & ill-prepared virtualized workloads. Little did the whipper-snapper know, this was a guarded Hyper-V host--and guarded hosts come pre-loaded with anti-whipper-snapper technology. Stated another way: watch as Hyper-V defends itself against a series of fabric-level attacks by leveraging Windows Server 2016's remote attestation, key protection/release, hypervisor-enforced code integrity and shielded virtual machine technologies.

Track 2 – Advancing Products Meet the New Threats | 3:30-3:55 PM | Jonathan Birch | Microsoft
Dangerous Content - Securing .Net Deserialization

Serialization is a powerful tool in .Net, but if used incorrectly it can create vulnerabilities, including remote code execution. In this talk, I explain how .Net deserialization vulnerabilities occur, and why they can only be prevented by application developers. I explain four common forms of this vulnerability in detail, two using only .Net libraries and two using common vulnerable 3rd party libraries. For each of these I explain multiple ways to modify the vulnerable code to make it safe. I then use these as a basis to provide general guidelines for securing deserialization. Finally, I discuss coding best practices to prevent these vulnerabilities from being introduced, along with methods for detecting .Net deserialization vulnerabilities both through static and dynamic analysis. A handout will be provided listing potentially vulnerable API's and how to use them safely, along with useful notes on detecting this vulnerability.

Track 2 – Advancing Products Meet the New Threats | 4:00-4:50 PM | Filippo Seracini and Lee Holmes | Microsoft
Born Secure. How to Design A Brand New Cloud Platform With A Strong Security Posture

What if you could design a sealed, cloud infrastructure starting from a clean slate? What security posture would you adopt? This is the opportunity we had with Azure Stack! Starting from the assumption that the first "enemy" to protect from is the Administrator, we designed a tightly constrained management experience, protected by a military-grade OS security baseline, multiple levels of network ACLs and the latest encryption standards. In this talk, we discuss the security posture of Azure Stack and how we built the security principles of Assume Breach and Hardened by Default directly into the architecture of the cloud infrastructure. We will also describe the security assumptions we took, and how those heavily impacted the overall design of the on-prem cloud platform that analysts defined as the Microsoft' secret weapon in the cloud war.

BlueHat

General Audience | November 9th, 2017

Track	Time	Speaker	Company	Talk Subject
Track 1 – I Swear It Wasn't Me!	9:00 - 9:50 AM	Kymberlee Price Sam Vaughan	Microsoft	Down the Open Source Software Rabbit Hole
	10:00 - 10:50 AM	Sean Metcalf	Trimarc	Active Directory Security: The Journey
	11:00 - 11:50 PM	Alex Ionescu	Crowdstrike	Baby's First Bounty: Lessons from bypassing Arbitrary Code Guard
Track 1 – Cloud Chasing	1:00 - 1:50 PM	Nate Warfield Ben Ridgway	Microsoft	All Your Cloud Are Belong to Us; Hunting Compromise in Azure
	2:00 - 2:25 PM	Oran Brill Tomer Teller	Microsoft	Go Hunt: An Automated Approach for Security Alert Validation
	2:30 - 2:55 PM	Matt Swann	Microsoft	Scaling Incident Response – 5 Keys to Successful Defense at Scale
	3:00 - 3:50 PM	Greg Foss	LogRhythm	PIE – An Active Defense PowerShell Framework for Office365
	4:00 - 4:50 PM	Mathias Scherman Daniel Edwards Tomer Koren	Microsoft	Leveraging Honeypots to Train a Supervised Model for Brute-Force Detection
Track 2 – Phishing for Trust	9:00 - 9:50 AM	Billy Leonard	Google	10 Years of Targeted Credential Phishing
	10:00 - 10:50 AM	Alex Weinert Dana Kaufman	Microsoft	Account Compromise 2017: in the Trenches with the Microsoft Identity Security and Protection Team
	11:00 - 11:50 AM	Yacin Nadji	Georgia Institute of Technology	28 Registrations Later: Measuring the Exploitation of Residual Trust in Domains
Track 2 – Attacking Products	1:00 - 1:50 PM	Lei Shi Mei Wang	Qihoo 360 Inc.	Out of the Truman Show: VM Escape in VMWare Gracefully
	2:00 - 2:50 PM	Matt Nelson	SpecterOps	"____ Is Not a Security Boundary." Things I Have Learned and Things That Have Gotten Better from Researching Microsoft Software
	3:00 - 3:50 PM	Alexander Chistyakov	Kaspersky Lab	Detection Is Not a Classification: Reviewing Machine Learning Techniques for Cybersecurity Specifics
	4:00 - 4:50 PM	Andrea Lelli	Microsoft	Wannacrypt + Smbv1.0 Vulnerability = One of the Most Damaging Ransomware Attacks in History
Track 3 – Threat Intelligence	9:00 - 9:50 AM	Nick Anderson	Facebook	Detecting Compromise on Windows Endpoints with Osquery
	10:00 - 10:50 AM	Brian Hooper Jagadeesh Parameswaran	Microsoft	Tales from the SOC: Real-world Attacks Seen Through Defender ATP
	11:00 - 11:50 AM	Mark Parsons	Microsoft	Using TLS Certificates to Track Activity Groups
Track 3 – Threat Intelligence	1:00 - 1:50 PM	Chaz Lever	Georgia Institute of Technology	A Lustrum of Malware Network Communication: Evolution and Insights
	2:00 - 2:50 PM	Andrew Brandt	Symantec	Dyre to Trickbot: An inside Look at TLS-Encrypted Command-And-Control Traffic
	3:00 - 3:25 PM	Alexis Dorais-Joncas Thomas Dupuy	ESET	Sednit Reloaded: The Bears' Operations From Christmas to Halloween
	3:30 – 4:20 PM	Chuck McAuley	Ixia Communications	Disrupting the Mirai Botnet

Talk Abstracts

Track 1 – I Swear It Wasn't Me! | 9:00-9:50 AM | Kymberlee Price and Sam Vaughan |Microsoft
Down the Open Source Software Rabbit Hole

Many developers today are turning to well established third-party open source components and libraries to speed the development process and realize quality improvements over creating an in-house proprietary font parsing or image rendering library from the ground up. Efficiency comes at a cost though: a single OSS component may have multiple additional OSS subcomponents, and an application or service may have dozens of different third party libraries implemented. The result is that third-party and open source libraries have the ability to spread a single vulnerability across multiple products - exposing enterprises and requiring software vendors and IT organizations to patch the same vulnerability repeatedly. How big of a problem is this at Microsoft, and what can be done to minimize the risk? This presentation will detail a real-world example of how a single OSS component introduced over 100 vulnerabilities into over a dozen Microsoft products, and how that got resolved.

Track 1 – I Swear It Wasn't Me! | 10:00-10:50 AM | Sean Metcalf | Trimarc
Active Directory Security: The Journey

Active Directory is only the beginning… Attackers have set their sights squarely on Active Directory when targeting a company, though this typically isn't the primary objective. The motivation and end goals range from stealing data to impacting corporate operations. In this regard, gaining control of Active Directory is a means to an end; compromising Active Directory is an easy way to gain access to all critical corporate resources. Effectively protecting Active Directory has become critical in limiting the impact of a breach. This talk takes the audience on a journey covering the various security milestones and challenges with Active Directory. A variety of (fictionalized) companies and their AD security posture are highlighted along with the challenges they encounter with securing their systems. Key elements involve how enterprise "AD aware" applications can weaken Active Directory security and how leveraging cloud services complicate securing infrastructure. Also explored is what an attacker can do in an environment without having Domain Admin rights. The final section discusses the commonly heard excuses for not implementing security controls to protect Active Directory and the ways to counter these arguments. Join the author of ADSecurity.org as he covers the critical issues affecting organizations today, as well as the biggest challenges; current attack techniques; and the most effective defensive techniques to prevent and mitigate compromise (including limitations to these approaches).

Track 1 – I Swear It Wasn't Me! | 11:00-11:50 AM | Alex Ionescu | Crowdstrike
Baby's First Bounty: Lessons from bypassing Arbitrary Code Guard

Although this has now changed, the original Microsoft bug bounty program was targeted toward bypassing OS mitigations, and especially those that are mostly affecting remote attackers or local-remote attackers, such as CFG and ACG. Though I harbor a love for looking for and finding design-level vulnerabilities in the system, this work rarely takes me into mitigations that target browsers or other content renderers/parsers, as my interests mainly lie in obscure persistence, hooking, rootkit & stealth techniques. Every once in a while, however, such paths may cross, and this is exactly what happened around late October 2016, when I stumbled upon the new kernel-level Warbird runtime support, designed to assist the deployment of ACG in Edge, where DRM requirements called for the need to have dynamically generated memory. While I originally researched this as a stealth packing/unpacking technique, I eventually realized its potential to entirely bypass all of ACG, due to what I believed was a flaw in its design. By the you attend this presentation, the bug will have been fixed with the release of Windows Fall Creators Update (RS3), resulting in almost an entire year between the finding and the fix – and resulting in an entire feature update shipping with ACG vulnerable the whole time. Part of this delay can be immediately traced to an obvious culprit – I only submitted the vulnerability around April 2017, months after realizing its potential. What led to this delay on my part, and how did the bounty program terms (at the time) affect the decision? How can we best determine severity types for bugs such as these, and when should such adjudications be made -- at finding time, or at fixing time? How do we balance this process, or changes within it, with a researcher’s desire to receive a timely payment? What if a fix is never released, but the issue/finding still of great importance? The goal of this talk is to attempt to answer - or at least generate food for thought on - some of these issues, from the point of view of a researcher (and one new to such program), not the program's coordinators.

Track 1 – Cloud Chasing | 1:00-1:50 PM | Nate Warfield and Ben Ridgway | Microsoft
All Your Cloud Are Belong to Us; Hunting Compromise in Azure

MongoDB, Redis, Elastic, Hadoop, SMBv1, IIS6.0, Samba. What do they all have in common? Thousands of them were pwned. In Azure. In 2017. Attackers have shifted tactics, leveraged nation-state leaked tools and are leveraging ransomware to monetize their attacks. Cloud networks are prime targets; the DMZ is gone, the firewall doesn't exist and customers may not realize they've exposed insecure services to the internet until it's too late. In this talk we'll discuss hunting, finding and remediating compromised customer systems in Azure - a non-trivial task with 1.59million exposed hosts and counting. Remediating system compromise is only the first stage so we'll also cover how we applied the lessons learned to proactively secure Azure Marketplace.

Track 1 – Cloud Chasing | 2:00-2:25 PM | Oran Brill and Tomer Teller | Microsoft
Go Hunt: An Automated Approach for Security Alert Validation

How often did you find yourself analyzing a security alert only to find out you had already hunted similar alerts in the past? This Déjà vu happens quite often to cybersecurity analysts who work in a SOC. What if we told you that most security alerts can be assigned with a confidence score automatically, letting you, the analyst, focus on the most serious alerts? In this talk, we will present tools and techniques to automate human cybersecurity analyst by leveraging knowledge of past incidents, current security posture and a dash of crowdsourcing. Under the hood, we generate a ”tailor-made” hunting graph based on diverse data sources and security know-how which enables us to extract meaningful insights. By applying custom logic, aggregations and data science we will illustrate how to uncover patterns within the insights and assign a confidence score with appropriate reasoning to the alert, automatically.

Track 1 – Cloud Chasing | 2:30-2:55 PM | Matt Swann | Microsoft
Scaling Incident Response – 5 Keys to Successful Defense at Scale

As defenders, we watch our intrusion detection systems like a hawk so that we know when to jump into action. However, successfully evicting an adversary in a large-scale environment requires capabilities beyond detection. In this talk I describe 5 capabilities that network defenders must have in order to effectively respond to an intrusion in a large-scale service. I describe how we overcame these challenges in Office 365 with pointers to source code and reusable tooling.

Track 1 – Cloud Chasing | 3:00-3:50 PM | Greg Foss | LogRhythm
PIE – An Active Defense PowerShell Framework for Office365

Phishing is often the bane of a security analyst's existence. Even with all of the fancy tools in place, organizations still have to be prepared to handle targeted attacks, scams, generic spam, and more as they continue to reach end users. The toughest part is the fact that analyzing, tracking, and reporting on these attacks is a massive time sink - costing organizations valuable time and money. For this reason, we've been hard at work developing an open source toolset to help streamline and automate the entire process of tracking, analyzing, and responding to phishing emails, without the need for commercial software. The Phishing Intelligence Engine (PIE) -- a PowerShell Active Defense framework built around Office 365, that continuously evaluates Message Trace logs for malicious contents, and dynamically responds as threats are identified or emails are reported. All links, files, and other potential malware are dynamically sandboxed for analysis and results are streamlined to the SOC without the need for manual intervention. If the message is deemed 'dirty' PIE will hard-delete all mail from every recipient's inbox, and extract copies of each, along with refined metadata to a case folder for the SOC to analyze. The entire process is tracked, and a detailed report is generated and stored on a share, accessible by our Threat Research Team. The 'case' folders are organized by date and contain all emails, samples, and research results, allowing analysts to easily review pending cases, and capture metrics for reporting on events that transpired -- all without having to do anything aside from approving the quarantine action. This is just a small piece of what I'd like to dive into with this session. Email is a critical component of every organization and is one of the easiest ways in. I'd like to share my experiences and provide Open Source tools to help others take control of their email security, save loads of time, and proactively defend their networks.

Track 1 – Cloud Chasing | 4:00-4:50 PM | Mathias Scherman, Daniel Edwards and Tomer Koren | Microsoft
Leveraging Honeypots to Train a Supervised Model for Brute-Force Detection

In this talk we will discuss of the use of honeypots to find labels and build supervised learning models in the context of cloud security. We will demonstrate how we applied this method to build a model for detecting incoming brute-force attacks on SQL services. SQL services (Microsoft SQL Server, MySQL, …) are sensitive to breaches occurring by a brute-force attack. Breaches can lead for instance to ransomware attacks or data stealing. In a brute-force attack, the attacker tries to obtain the credentials information of the server in a trial-and-error manner, typically running automate software to generate many consecutive guesses. To mitigate this kind of risks, Microsoft offers an online service called Azure Security Center (ASC) which uses various techniques, including machine learning ones, to analyze logs from the VMs and raise security alerts. On IaaS machines, on the contrary of PaaS, the cloud provider is not maintaining the software, and thus has no access to operating system logs such as the success or failure of each connection attempt. 7% of Azure VMs have a SQL server installed. Thus, simple rule-based approaches using the frequency of connections are not applicable. Machine Learning techniques relying on external network data exist. However, as most of the data is unlabeled, they often rely on anomaly detection. Thus, such methods produce many false positives since many anomalies are not the result of a security incident. To overcome the lack of labeled data, we present an approach to leverage honeypot hits to tag malicious connections in Azure, which enables us to effectively train a supervised model. Due to the model’s accuracy and its customer impact, it has been implemented in Azure Security Center.

Track 2 – Phishing for Trust | 9:00-9:50 AM | Billy Leonard | Google
10 Years of Targeted Credential Phishing

While it's not kernel 0days or EMET bypasses, credential phishing has been a go-to in attackers toolboxes for many years, rising to prominence during the run up to the 2016 US Presidential Elections. Being able to access a target's email or files stored in the cloud without burning your prized 0day has proven to be too much for even the most advanced attackers to pass up. In this talk, we will look at how attackers have evolved and adapted their credential phishing operations over the past 10 years, from changes in delivery mechanisms to changes in persistence and exfiltration and how defenses have evolved during that same time.

Track 2 – Phishing for Trust | 10:00-10:50 AM | Alex Weinert and Dana Kaufman | Microsoft
Account Compromise 2017: in the Trenches with the Microsoft Identity Security and Protection Team

Join us for an action packed hour as we summarize the top incidents and attacker methodologies encountered by Microsoft's Identity Security and Protection team (ISP) in 2017. This talk will cover our wins and losses, breaking down attacks by type and volume, and be liberally decorated with stories of where it went sideways and what we learned about how account penetration is happening across our user base. Cyber criminals know that stealing a valid user's digital identity makes virtually all other defenses ineffective, so they bias towards account compromise to carry off their crimes. This puts the Identity Security and Protection team on the front lines of cybersecurity. Users on Office 365, Xbox, OneDrive, Outlook, Azure, and at millions of enterprise desks around the world rely on us to keep their identities – and all the services and assets they allow access to – safe from cybercriminals. The team is charged with preventing account hacking and fraud across all of Microsoft's Identity systems – consumer or enterprise, on any device or any service. We evaluate more than 10 billion logins a day for 400+ million unique daily user s (more than 4B accounts in total), and stop more than 30 million account attacks *every day*, giving us unparalleled insight into the state of affairs in cyber-crime. Join us for key insights from 2017,focused on issues impacting our organizational custom ers and key recommendations for protecting your organization in the face of evolving attacks.

Track 2 – Phishing for Trust | 11:00-11:50 AM | Yacin Nadji | Georgia Institute of Technology
28 Registrations Later: Measuring the Exploitation of Residual Trust in Domains

Any individual that re-registers an expired domain implicitly inherits the residual trust associated with the domain's prior use. We find that adversaries can, and do, use malicious re - registration to exploit domain ownership changes— undermining the security of both users and systems. In fact, we find that many seemingly disparate security problems share a root cause in residual domain trust abuse. With this study we shed light on the seemingly unnoticed problem of residual domain trust by measuring the scope and growth of this abuse over the past six years. During this time, we identified 27,758 domains from public blacklists and 238,279 domains resolved by malware that expired and then were maliciously re-registered. To help address this problem, we propose a technical remedy and discuss several policy remedies. For the former, we develop Alembic, a lightweight algorithm that uses only passive observations from the Domain Name System (DNS) to flag potential domain ownership changes. We identif y several instances of residual trust abuse using this algorithm, including an expired APT domain that could be used to revive existing infections.

Track 2 – Attacking Products | 1:00-1:50 PM | Lei Shi and Mei Wang | Qihoo 360 Inc.
Out of the Truman Show: VM Escape in VMWare Gracefully

Virtualization is one of the most complicated software in the world. The VMware workstation is very popular in many fields. The windows 10 has a lot of mitigation technology to get avoid of exploitation. It's a great challen ge to make a vm escape in VMware workstation under Win 10. Especially when the guest and host are both win 10 and the guest user are NO-ADMIN. This talk will present how to make a vm escape and execute arbitrary code in the host from a NO -ADMIN guest user under Win 10(both the guest and host are Win 10). They have developed three different exploitation. This talk will introduce them and show a very elegant exploitation technology of vm escape. Besides the vm escape technology, this talk will also show the exploitation technology in Win 10. It is quite attractive because there's a process continuation, saying that the guest can execute the exploitation without crashing/disturbing the host process(VMware workstation virtual machine process). The exploitation is very reliable, it reaches nearly 100% successful rate.

Track 2 – Attacking Products | 2:00-2:50 PM | Matt Nelson | SpecterOps
"____ Is Not a Security Boundary." Things I Have Learned and Things That Have Gotten Better from Researching Microsoft Software

A persistent "enlightened" attacker will invest the required resources to bypass any and all security features that might stand between them and their objective, regardless if these features are guaranteed to be serviced as security boundaries or not. This includes researching and developing attacks against Windows security features that may impose a hurdle in their attack chain. This talk will outline recent research into features such as User Account Control (UAC), the Antimalware Scan Interface (AMSI) and Device Guard and how these bypasses are useful to attackers in an operational context. Some examples include: UAC: If an attacker compromises a user that is running as a split-token administrator, bypassing UAC is required in order to perform any administrative actions; such as dumping credentials from memory. AMSI: With in-memory attacks becoming more prevalent via scripting languages, AMSI is the next logical step to facilitate detection. An attacker will need to bypass AMSI in order to safely operate in memory when using PowerShell, VBScript, or JScript. Device Guard: As organizations begin to consider whitelisting solutions, an attacker is required to adapt and develop a bypass to these technologies. One such solution is Device Guard, which can be used to heavily restrict what is allowed to execute on the system. In order to accomplish their objective, an attacker would need to bypass User Mode Code Integrity (UMCI). Such research can find novel ways to execute code in ways that are not likely to be detected. I will also cover some of the fixes that have been implemented in newer versions of the Windows Operating System. Fixing these bypasses will not only make Windows safer, but it will begin to disrupt attackers by raising the cost associated with successfully executing an attack.

Track 2 – Attacking Products | 3:00-3:50 PM | Alexander Chistyakov | Kaspersky Lab
Detection Is Not a Classification: Reviewing Machine Learning Techniques for Cybersecurity Specifics

While more and more security vendors are starting to use Machine Learning (ML) models for malware detection, the basic pipeline for the construction of these detectors usually looks the same: collect a dataset of benign and malicious samples, train a binary classifier to predict the correct label, use a positive prediction of the model to detect new malware. However, this approach does not take into account one important and natural property: no malicious code could become clean after the injection of any new functionality. As a result, an intruder can often avoid detection, simply by adding some obfuscated or clean-looking payload into the malware sample. In this talk we will show how to construct a ML detection model, that is provably secure against such attacks even, after the full reverse engineering. Using the real-time malicious activity detection problem as an example, we will review the classical step-by-step pipeline for designing, training and utilizing the ML classifier, and explain how to adapt it to the specifics of the malware detection problem. We will explain how to transform almost any applicable ML architecture (Deep NN, tree-based ensembles, kernel SVM, etc.) to make your static or dynamic malware detection model more secure; how to update the model's decision border without complete re-training; and how to explore the causes of the detection alert using the transformed architecture.

Track 2 – Attacking Products | 4:00-4:50 PM | Andrea Lelli | Microsoft
Wannacrypt + Smbv1.0 Vulnerability = One of the Most Damaging Ransomware Attacks in History

My presentation will trace the end-to-end WannaCrypt (also known as WannaCry) attack. I will start with an analysis of the underlying SMBv1 remote code execution kernel-mode exploit dubbed "Eternalblue", a powerful cyberweapon leaked by a hacker group known as "The Shadow Brokers". I will then describe how the Wannacrypt ransomware works, and show how the cybercriminals leveraged the EternalBlue exploit to spread the ransomw are and achieve a massive and unprecedented infection rate, leaving hundreds of thousands of machines affected. I will highlight the Windows 10 kernel mitigations that granted the OS immunity from the attack. I will also focus on some interesting characteristics that make WannaCrypt particularly sophisticated, like the file-wiping and space-consuming capabilities designed to make the recovery of the original files nearly impossible. I will conclude with a look into how much the perpetrators might have likely earned from the attack. An analysis of the Bitcoin transactions shows that the cybercriminals pooled around $140000 to date, which is a good amount of money, but doesn't seem to scale with the extent of infection. Not to mention, Bitcoin is a double-edged sword and the cybercriminals had a hard time trying to cash the money. In this section I will also mention some copycat malware that tried to spread using the same SMB vulnerability (e.g. NotPetya). I will end the presentation with advice on preventing, detecting, and responding to ransomware attacks.

Track 3 – Threat Intelligence | 9:00-9:50 AM | Nick Anderson | Facebook
Detecting Compromise on Windows Endpoints with Osquery

Just as Microsoft grows to embrace the open source community more and more, we must use open source tools to help us grow as a community. In this talk we'll explore the various advanced detection techniques we employ at Facebook using osquery for Windows. Specifically, we will examine instrumenting Windows Event Log data, inspecting detailed attack patterns on processes such as path hijacking, and mapping operating system state to detect deviations of a healthy system - all at Facebook scale. Building on these detection capabilities, we will then consider different response features currently available in osquery and how one can extend these capabilities to suit the needs of their own enterprise. By striving to make these advanced detection capabilities more approachable we hope to raise the bar of defenses employed by companies everywhere and encourage the security community to take a more proactive role in developing detection features used to catch advanced exploitation.

Track 3 – Threat Intelligence | 10:00-10:50 AM | Brian Hooper and Jagadeesh Parameswaran | Microsoft
Tales from the SOC: Real-world Attacks Seen Through Defender ATP

Windows Defender ATP gives defenders unparalleled visibility into the enterprise. Come spend an hour with us as we pull back the covers and go through detailed examples of real attacks that we saw as we defended the Microsoft corporate environment.

Track 3 – Threat Intelligence | 11:00-11:50 AM | Mark Parsons | Microsoft
Using TLS Certificates to Track Activity Groups

As the internet moves to encryption for standard protocols we have seen malware also following that trend by using TLS certificates for encrypting C2 communications. Using open source scanning data projects like Shodan, Censys and Rapid

7 sonar we will discuss ways to use this scanning data. This talk will go over examples of using TLS certificates for tracking multiple activity groups and their infrastructure , ways to find popular post exploitation frameworks and some examples of getting to know your own environment.

Track 3 – Threat Intelligence | 1:00-1:50 PM | Chaz Lever | Georgia Institute of Technology
A Lustrum of Malware Network Communication: Evolution and Insights

Both the operational and academic security communities have used dynamic analysis sandboxes to execute malware samples for roughly a decade. Network information derived from dynamic analysis is frequently used for threat detection, network policy, and incident response. Despite these common and important use cases, the efficacy of the network detection signal derived from such analysis has yet to be studied in depth. This paper seeks to address this gap by analyzing the network communications of 26.8 million samples that were collected over a period of five years. Using several malware and network datasets, our large scale study makes three core contributions. (1) We show that dynamic analysis traces should be carefully curated and provide a rigorous methodology that analysts can use to remove potential noise from such traces. (2) We show that Internet miscreants are increasingly using potentially unwanted programs (PUPs) that rely on a surprisingly stable DNS and IP infrastructure. This indicates that the security community is in need of better protections against such threats, and network policies may provide a solid foundation for such protections. (3) Finally, we see that, for the vast majority of malware samples, network traffic provides the earliest indicator of infection—several weeks and often months before the malware sample is discovered. Therefore, network defenders should rely on automated malware analysis to extract indicators of compromise and not to build early detection systems.

Track 3 – Threat Intelligence | 2:00-2:50 PM | Andrew Brandt | Symantec
Dyre to Trickbot: An Inside Look at TLS-Encrypted Command-And-Control Traffic

Back in 2014 and 2015, the Dyre (sometimes called Dyreza) Trojan was a distinctive crimeware tool for the simple reason that it appeared to employ, and experiment with, a whole range of sophisticated tactics, techniques and procedures: It was the first Trojan which exclusively employed HTTPS for its C2 traffic; It operated on a modular basis with a small cadre of other malware families, such as the Upatre downloader, which seemed to support it exclusively, as well as email address scraping tools and spam mail relayers; and it was at least as interested in profiling the environment it had infected as it was in exfiltrating any data it could find on the victim's machine. Then it disappeared suddenly, but re-emerged this year in the form of a Trojan now called Trickbot (aka Trickybot), completely rewritten but with many of the same features. In the lab, we permit Trickbot samples to persist on infected machines for days to weeks in order to perform man-in-the-middle SSL decryption on their C2 traffic. In this session, attendees will get a detailed forensic analysis of the content of some of this C2 traffic and the endpoint behavior of various machines (virtual and bare-metal) when left infected for an extended period of time. Finally, we will share what we know about the botnet's C2 infrastructure and its historical reputation. By understanding how Trickbot functions, and to where it communicates, we hope we can help identify infections more rapidly and, maybe, interpret the motives of whoever is operating this shadowy botnet to predict its next course of action.

Track 3 – Threat Intelligence | 3:00-3:25 PM | Alexis Dorais-Joncas and Thomas Dupuy | ESET
Sednit Reloaded: The Bears' Operations From Christmas to Halloween

The Sednit group, a.k.a Fancy Bear, Sofacy, or APT28, is one of the most prolific APT groups in existence. They have gained an increasing amount of attention from the media and researchers over the years. They have allegedly infiltrated strategic organizations during the past years, like the Democratic National Committee, the German Parliament, and the French media TV5 Monde. They are also known to pull out 0-day exploits in order to compromise their victims. Last year, we released an extensive analysis [1] of Sednit's toolkit, describing their arsenal as well as their operations. Since then, their ecosystem has kept evolving. In this presentation, we will talk about the current trends we've observed since December 2016, including components they stopped using and refinements to their existing toolkit. Here is what we will cover: - A few 0-day exploits we've found in the past ten months, we will quickly analyze how the exploits were used, how we tracked them. We will talk about our experience reporting these critical vulnerabilities, which were actively exploited in the wild. - An overview of the targeted campaigns we have seen recently, with the evolution of their toolkit and the disappearance of some of their components   - Recent discoveries on XAgent: Now at version 4, they keep working on their flagship backdoor. Now with new features, we will describe its evolution through the years. - And finally, we will talk about a new component we dubbed Tartine (a.k.a Zebrocy), which is a Delphi backdoor that was heavily used while attacking Eastern Europe institutions. This component went under the radar for quite some time and was recently linked to the Sednit group. [1] https://www.welivesecurity.com/wp-content/uploads/2016/10/eset-sednit-full.pdf

Track 3 – Threat Intelligence | 3:30-4:20 PM | Chuck McAuley | Ixia Communications
Disrupting the Mirai Botnet

The Mirai botnet has brought public awareness to the danger of poorly secured embedded devices. Its ability to propagate is fast and reliable. Its impact can be devastating and variants of it will be around for a long time. You need to identify it, stop it, and prevent its spread. I had the opportunity to become familiar with the structure, design, and weaknesses of Mirai and its variants. At this talk you'll learn how to detect members of the botnet, mess with them through various means and setup a safe live fire lab environment for your own amusement. I will demonstrate how to join a C2 server, how to collect new samples for study, and some changes that have occurred since release of the source code. By the end you'll be armed and ready to take the fight to these jerks. Unless you're a botnet operator. Then you'll learn about some of the mistakes you made.