



BlueHat 2024

Breakout Session Abstracts

Oct 29-30
2024
Redmond
WA, USA

Deprecating Azure AD Graph API is Easy and Other Lies We Tell Ourselves

Presented by Nestori Syynimaa from Microsoft

Azure AD Graph API and AzureAD PowerShell module were introduced in 201x to replace the MSOnline PowerShell module utilizing a “provisioning API”. In 2019, we announced that AzureAD and MSOnline modules and their back-end APIs will deprecate by June 2020. However, the deprecation date has been postponed multiple times ever since. Both PowerShell modules have been deprecated as of March 30, 2024, but the back-end APIs are still functional until June 30, 2025.

In April 2024, we announced the general availability of Microsoft Graph activity logs, which allows our customers to identify activities conducted by compromised accounts and other problematic behaviours of applications. However, these logs are not available for Azure AD Graph API. For this reason, adversaries likely keep using Azure AD Graph API with compromised user identities until its deprecation.

This talk covers the history of the Azure AD Graph API, sheds light on why its deprecation has been so difficult, how to migrate to MS Graph API, and finally, how to detect & prevent malicious Azure AD Graph API usage.

About the Presenter:

Dr Nestori Syynimaa is a Principal Identity Security Researcher at Microsoft Threat Intelligence Center. He has over a decade of experience with the security of Microsoft cloud services and is known as the creator of the AADInternals toolkit. Before joining Microsoft in early 2024, Dr Syynimaa worked as a researcher, CIO, consultant, trainer, and university lecturer for over 20 years.

Double Agent: Exploiting Pass-through Authentication Credential Validation in Azure AD

Presented by Avihai Ben Yossef from Cymulate

This research explores the vulnerabilities within Microsoft Azure Active Directory's Pass-Through Authentication (PTA) mechanism, focusing on the credential validation process. By analyzing PTA's architecture and operation, the study identifies specific weaknesses that could be exploited by attackers to bypass authentication and gain unauthorized access to cloud environments. The research provides a detailed examination of these vulnerabilities, supported by practical examples and potential attack scenarios, offering insights into improving the security of Azure AD deployments.

About the Presenter:

Avihai Ben Yossef from Cymulate

Echoes of Intrusion: Demystifying MS Graph API Attacks

Presented by Miriam Wiesner from Microsoft

The Microsoft Graph API can be a dangerous accomplice: combined with the right permissions, attackers can control everything in Microsoft 365. But it is still hard to detect and prevent Graph API related attacks. In my research I investigated and simulated MS Graph related attacks; parts of this work I will share in this presentation.

Take a peek behind the curtains of how adversaries (ab)use MS Graph API to breach corporate defences, implant persistence mechanisms, and exfiltrate sensitive data. Through a live demonstration, we will simulate the execution of potential attack scenarios, providing invaluable insights into the adversary's playbook.

Explore attack scenarios from the field and learn how to detect and prevent against such an attack. Are you ready to dive into real world threat actors' tactics?

About the Presenter:

Miriam C. Wiesner is a Sr. Security Research Program Manager at Microsoft, with more than 15 years of experience in IT and IT Security. She has held various positions, including Administrator/System Engineer, Software Developer, Premier Field Engineer, Program Manager, Security Consultant & Pentester, and Security Researcher.

She is also a renowned creator of open-source tools based in PowerShell, including EventList and JEAnalyzer. She was invited multiple times to present her research behind her tools at international conferences like Black Hat (USA, Europe & Asia), PSConf EU, MITRE ATT&CK workshop, and more.

Miriam is the author of the book "PowerShell Automation and Scripting for CyberSecurity: Hacking and Defense for Red and Blue Teamers", which was released in August 2023.

Creating a Transparent Cloud Industry

Presented by Justin Mourfield and Sesha Machiraju from Microsoft

Companies around the world leveraging cloud-based products inherently delegate much of the systems management to the Cloud Service Provider (CSP). This presents a unique and unprecedented challenge, ensuring the security of their environment and protection of data. CSPs entrusted with this responsibility must adhere to security and data protection policies but as technology evolves, so must the way we perform these duties. Reducing the threats posed by security incidents and the ever-changing threat landscape is more complex than ever before. Protecting customers is the responsibility of all CSPs but can only be achieved through enhancing security transparency, providing comprehensive vulnerability data, and cross-industry cooperation.

About the Presenters:

Justin Mourfield is a Senior Technical Program Manager in MSRC who is responsible for managing and coordinating all of Microsoft's vulnerability disclosures (CVEs), the "Patch Tuesday" process, and directly support MSRC's security and privacy incident response operations. I have 14 years of experience performing security operations and incident response across multiple domains, leading cross-org collaboration efforts, and interfacing with customers, account teams, and industry partners to build their trust and confidence in Microsoft. I am a fierce advocate for customers who leverages subject-matter expertise in security, privacy, and cloud engineering to provide the guidance and unparalleled support which enables them to perform critical functions or protect their environments.

Sesha Machiraju is a Technical Program Manager 2 within Azure CXP's Security Incident Response team who has 9 years of experience performing data analysis, incident response, and managing customer focused programs to improve their cloud security and reliability. I am directly responsible for defining Microsoft's holistic communication strategy during MSRC's security and privacy response process, manage the requisite publication of internal and external communications, and ensure compliance with Federal and International regulatory requirements. Additionally i lead multiple efforts that educate customers and their account teams on security best practices and incident response.

The two sides of UnOAuthorized

Presented by Eric Woodruff from Semperis and Cameron Vincent from Microsoft

For customers of Microsoft 365 and Azure, obtaining the role of Global Administrator (GA) is every attacker's dream - it is the Domain Administrator of the cloud. This makes Global Administrator every organizations nightmare of being owned by a threat group or hacker. Luckily, well-defined role-based access control and a strict application consent model can severely limit who gets their fingers on Global Administrator - or does it?

This session walks through the research background and the discovery of unexpected authorization in Entra, allowing an Application Administrator to abuse Microsoft service principals to take the role of Global Administrator, among some other discoveries.

Unique to this BlueHat session, we'll then pivot and dive into the MSRC response and resolution, directly from an MSRC researcher involved in the case. We'll explore how MSRC and the IDNA team worked to resolve a complex and layered case, provide insight into what goes on behind the scenes when researching a reported identity vulnerability, as well as additional background on the findings from the Microsoft side.

About the Presenters:

Throughout his 24-year career in the IT field, Eric Woodruff has sought out and held a diverse range of roles, including technical manager in the public sector, Sr. Premier Field Engineer at Microsoft, and Security and Identity Architect in the Microsoft Partner ecosystem. Currently he is a Senior Security Researcher working as part of the Security Research team at Semperis. Eric is a Microsoft MVP for security, recognized for his expertise in the Microsoft identity ecosystem. Outside of work, Eric supports the professional community, providing his insights and expertise at conferences, participating on the IDPro Body of Knowledge committee, and blogging about Entra and related cloud security topics.

Starting out as a full-time bug bounty hunter, Cameron Vincent now works on the vulnerabilities and mitigations team in the Microsoft Security Response Center. During his full-time bug bounty career, he was ranked as one of the top researchers for both Microsoft and Google's program numerous times. He now works on the V&M team within the MSRC side dealing with security issues internally to better protect Microsoft's ecosystem from within.

Tokens & Takeovers: Cloud-Powered Supply Chain Attacks

Presented by Nitesh Surana from Trend Micro and Gaurav Mathur from Microsoft

It takes one single misconfigured token to jeopardize cloud resources and their downstream dependent systems; a recent example being an overly permissive SAS token leading to the 2023 leak of 38 TB of Microsoft AI research data.

After Microsoft's account of the incident, we did our own part in hunting for overly permissive SAS tokens. We found two different ways of controlling a widely used official Microsoft tool called PC Manager. One could eventually execute a classic supply chain attack across multiple releases of MS PC Manager that were sprinkled across the web in multiple blogs, support forums using 'aka.ms' URL shorteners, an official Microsoft website and WinGet packages. Using the SAS tokens, we could takeover every release of MS PC Manager.

Furthermore, we will share our findings wherein one could inject malicious stored procedures on database backups for tutorials mentioned in official Azure docs, modify JavaScript resources being used on multiple websites using one single SAS token. To conclude, we will share what practitioners can do to proactively hunt for sensitive information in URL parameters such as SAS tokens in their environments.

About the Presenters:

Nitesh Surana from Trend Micro

Gaurav Mathur from Microsoft

DCOM Research for Everyone!

Presented by James Forshaw from Google

For almost 10 years I've been writing a tool, OleView.NET to aid in the analysis of Windows' COM attack surface for privilege escalation and remote code execution. What started as a merging of the SDK's OleView and Test Container GUIs it expanded out to an all-purpose security research tool adding support for Windows Runtime classes, security descriptor analysis, process parsing along with a PowerShell interface to augment the GUI.

The final important piece of the tool's functionality, the ability to call arbitrary COM interfaces at runtime without prior knowledge, was missing. This is crucial as it allows a researcher to not only inspect what COM services are registered but directly interact with them. While calling methods on a COM object was always supported for OLE Automation interfaces, for which the .NET Framework has good support, most of the interesting interfaces such as those exposed by privileged services do not just give you the information necessary to call them.

This presentation is about the research and implementation of this feature. It leverages the fact that almost all COM interfaces that can be called across process boundaries have registered proxies that can be analyzed to extract their marshaling information. That can then be used to reimplement a client. Of course, nothing is that simple, I'll describe many of the challenges that had to be overcome during the research and implementation, some of the false starts and some of the interesting security issues that came out of the work. I'll also demo the tooling so that attendees can have an understanding on how to use it for their own research.

About the Presenter:

James is a security researcher in Google's Project Zero. He has been involved with computer hardware and software security for over 10 years looking at a range of different platforms and applications. With a great interest in logical vulnerabilities, he's been listed as the #1 researcher for MSRC, as well as being a Pwn2Own and Microsoft Mitigation Bypass bounty winner. He has spoken at a number of security conferences including Black Hat USA, CanSecWest, Bluehat,

HITB, and Infiltrate. He's also the author of two security books "Windows Security Internals" and "Attacking Network Protocols", both available from NoStarch Press.

Pointer Problems - Why We're Refactoring the Windows Kernel

Presented by Joe Bialek from Microsoft

The Windows kernel ecosystem is facing security and correctness challenges in the face of modern compiler optimizations. These challenges are no longer possible to ignore, nor are they feasible to mitigate with additional compiler features. The only way forward is large-scale refactoring of over 10,000 unique code locations encompassing the kernel and many drivers.

This talk provides an overview of some optimizations the compiler performs that are often unexpected for developers. We will show real-life examples of these optimizations creating MSRC cases or correctness problems. Finally, we'll show what we're doing to fix the problem, the progress we've made in the latest Windows release, and what developers will need to do to ensure their code is also robust.

About the Presenter:

Joe Bialek, Microsoft, MORSE, Security Engineer. Joe has been working on low-level security at Microsoft for over 10 years in code bases such as the Visual Studio compiler, Windows Kernel, and Hyper-V stack. Joe currently leads or contributes to most exploit and vulnerability mitigations being built for Windows.

How Microsoft is Scaling DAST

Presented by Jason Geffner from Microsoft

Take a rare look inside how Microsoft is working to use automated reverse engineering techniques to secure its own web services. This talk discusses common challenges that enterprises face when scaling Dynamic Application Security Testing (DAST), including getting developers to provide service API specs, manual configuration of authentication and authorization, and scanning bandwidth limitations and latency. We present innovative automated solutions developed by Microsoft to address these challenges, including Automated Discovery of API Endpoints, Transparent Authentication & Authorization, and Accelerated DAST.

About the Presenter:

Jason Geffner, Principal Security Architect, Microsoft. Jason Geffner is an information security professional with an extensive history in application security, risk management, malware analysis, threat intelligence, incident response command, endpoint security, security automation, vendor security management, and security research & development.

Sweet QuaDreams or Nightmare Before Christmas? Dissecting an iOS 0-day

Presented by Christine Fossaceca from Microsoft and Bill Marczak from Citizen Lab

Not quite nation states but not quite independent corporations, “private sector offensive actors” (PSOAs) have become one of the latest sophisticated threats. These companies develop and sell surveillance and intrusion capabilities to governments around the world. While some governments responsibly use the tools to track criminals and terrorists, others instead opt to abuse the tools by spying on journalists, dissidents, or members of their political opposition.

The conversation about PSOAs often centers around NSO Group, and their infamous zero-click “Pegasus” spyware. However, an industry of competitors abounds. While the final payload of Pegasus has proved elusive for some time, Microsoft and Citizen Lab successfully obtained and analyzed the final payload associated with a separate zero-click mobile threat fielded by an NSO competitor, “QuaDream”. QuaDream’s spyware was used against targets around the world, including journalists, political opposition figures, and an NGO worker. This sample was deemed “KingsPawn” by Microsoft and the exploit named “ENDOFDAYS” by CitizenLab.

So what does it take to develop such a zero-click, zero-day attack? What does a modern, top-tier, iOS spyware implant look like? What is the state-of-the-art in mobile threats? And what is the likelihood of you or your employees being targeted by such an attack? In this talk, Bill Marczak and Christine Fossaceca discuss the discovery of QuaDream’s spyware, outline the zero-click exploit likely used to deliver it, and share their experience reversing engineering the attack surface from the ground up.

About the Presenter:

Christine Fossaceca, Microsoft, Senior Security Researcher. Christine Fossaceca is a Senior Mobile Security Researcher at Microsoft, specializing in iOS. She has a background in mobile exploit development, forensics techniques, red teaming, reverse engineering, and penetration testing. Christine’s current focus is on the Defender for Endpoint team analyzing iOS 0-days.

When the Levee Breaks: Exposing Critical Flaws in Wi-Fi Camera Ecosystems

Presented by Mark Mager and Eric Forte from Elastic

Wi-Fi security cameras are affordable, easy to use, and extremely convenient. They can be found in hundreds of millions of households across the world and are remotely accessible from anywhere via the Internet. An unfortunate byproduct of the industry-wide push to manufacture and distribute these IoT devices to the masses as quickly as possible is that they tend to suffer from inherently flawed security models. These cameras require constant connectivity to insecure cloud platforms and facilitate remote user access through vulnerable peer to peer protocols designed to circumvent secure network configurations.

Through extensive reverse engineering and vulnerability research over the past few months, we discovered several critical hardware and software vulnerabilities affecting millions of devices connected to a prominent IoT platform. Along with highlighting our process and key successes from this effort, we will provide details on how endemic these types of flaws are throughout the broader Wi-Fi camera industry along with countermeasures users can take to limit potential exposure and reduce their attack surface.

About the Presenters:

Mark Mager leads the Endpoint Protections Team at Elastic. He has served in prominent technical leadership roles in the research and development of advanced computer network operations tools and has provided malware analysis and reverse engineering subject matter expertise to government and commercial clients in the Washington, D.C. metropolitan area.

Eric Forte is a Security Research Engineer at Elastic with a background in embedded systems and streaming data analysis. He has worked in technical leadership roles in engineering Low Size Weight and Power (SWaP) capabilities and network security solutions. As part of this work, he managed an IoT research and reverse engineering lab to help in the development of these different capabilities for various organizations across the United States.

Outlook Unleashing RCE Chaos CVE-2024-30103 & CVE-2024-38021 & CVE-2024-38173

Presented by Michael Gorelik from Morphisec

Have you ever received an empty email and immediately thought it might be a reconnaissance attack? What if opening such an email in your Outlook client could trigger remote code execution through an invisible form? Yes, all Outlook message objects are based on MAPI Forms, and CVE-2024-21378 has flung open the gates to Outlook RCE chaos.

In our session, we'll dive into how this seemingly innocuous vulnerability can lead to mayhem. This vulnerability paved the way for us to discover a series of new remote code execution vulnerabilities in Outlook, including CVE-2024-30103, CVE-2024-38021 and CVE-2024-38173. But we're not stopping there.

Additionally, we'll uncover other vulnerabilities that can cause NTLM leaks from your domain-joined devices.

So, how did we get here? Join us as we construct an evolution timeline of this attack surface. From the origins of these exploits to their current incarnations, we'll cover it all. And because we believe in building a safer digital world, we'll conclude with specific, actionable recommendations on how to minimize these threats.

About the Presenter:

Michael Gorelik - Founder, Windows Reverse engineer, Red Teamer, MSC Computer Science, Vulnerability and Malware Researcher, Incident Responder, Speaker at Defcon, GovWare, BSides and more.

Scaling AppSec With an SDL for Citizen Development

Presented by Michael Bargury from Zenity/OWASP and Don Willits from Microsoft

Application security programs are difficult. Filled to the brim with vulnerabilities. Overloaded staff and inadequate budget. Challenging communication with developers. The common "solution" is to narrow scope and focus on crown jewel applications and their developers, playing on relative easy mode. What if instead we increase the scope to cover 100x developers and 1000x applications?

Surprisingly, it works. In the first 3 months of 2024, our program remediated >50K security vulnerabilities. 18K of them were remediated in a single night.

In this talk, we will share insights from two years in the making of a security program for applications built by business users using GenAI and low-code/no-code tools, a.k.a. Citizen

Development. We will share lessons learned and pitfalls not-avoided, and unique challenges for this kind of program. Applying SDLC to hundreds of thousands of citizen developers, with no security savvy. Working at 1,000x the AppSec scale relying on automation and guidance. Next, we will share the kind of vulnerabilities we see common in citizen development environments. Breaking access controls, allowing one user to impersonate another, leaking data to uncontrolled locations. We will demo exploits showing how they look like from the attacker's perspective.

We will finish off sharing our adoption of the SDL for citizen development, and showcase the OWASP Low-Code No-Code Top 10 as a framework to help you focus your program.

About the Presenters:

Michael Bargury, Zenity, Co-founder and CTO

Don Willits, Microsoft, Power Platform Security Architect

Embedding Sysmon Logs for Enhanced Threat Detection: A Practical Approach to Using RAG in Cybersecurity

Presented by Jose Rodriguez from George Mason University

As cybersecurity threats become more complex, AI-powered tools like language models (LM) can significantly improve the efficiency and accuracy of security investigations. One such tool, Retrieval Augmented Generation (RAG), enhances the relevance of analysis by retrieving and integrating key security information during the generation process. While RAG has been widely used to process unstructured data like threat intelligence reports, can it also be applied to structured and semi-structured data, such as security event logs?

In this presentation, I will introduce both traditional and LM-driven approaches to create embeddings from Sysmon logs that make it possible to implement RAG techniques on a case study on the APT29 adversary group, guided by the MITRE ATT&CK framework. This case study will demonstrate how these methods can enhance the detection and analysis of advanced cyber threats.

Attendees will learn how to preprocess event logs, fine-tune a pre-trained language model, and apply RAG techniques for log analysis. A Jupyter Notebook with all the steps will be provided, allowing participants to replicate the process in their own environments.

About the Presenter:

Jose is currently completing an MS in Data Analytics Engineering with a concentration in Cyber Analytics at George Mason University (GMU), where he also earned a BS in Statistics with a minor in systems engineering. As a former Cybersecurity Engineer for the MITRE-ATT&CK team, where he supported the development of data sources and components, he is passionate about combining cybersecurity and artificial intelligence to enhance autonomous decision-making in cyber defense. His experience includes developing machine learning models using Python, R, and Jupyter Notebooks, with a strong background in simulating attacker behaviors to generate incident data. He also actively contribute to open-source projects such as Threat Hunter Playbook, OSSEM, and Security Datasets. Jose is dedicated to creating machine learning solutions that help protect organizations and empower security teams to defend against cyber threats.

Patterns in the Shadows: Scaling Threat Hunting and Intelligence for Modern Adversaries

Presented by Mark Parsons and Colin Cowie from Sophos

In 2022 Sophos MDR launched a dedicated threat hunting and intelligence team to counter the evolving threats. This talk will share insights on implementing and scaling a successful threat

hunt & intel program, highlighting the symbiotic relationship between these disciplines to enhance security operations. By leveraging clustering techniques, our intelligence analysts provide context by profiling and attributing “patterns of life” comprised of TTPs, vulnerability abuse, infrastructure, tooling, malware and motivations over time. Threat hunters create, contribute, and monitor these clusters through various hunting methods.

We will present a case study on the STAC6451 cluster, when a ransomware actor automated Mimic ransomware deployment across multiple environments. We discuss the SPADE tool which identified STAC6451 through anomaly-based hunting; programmatically generating baselines to detect command and control activity across 25,000 environments. Then, we showcase practical challenges and successes of clustering when multiple threat actors may be operating within the same environment, and how hypothesis driven hunts from these clusters allowed for identification of newly compromised environments and expansion of the cluster.

We then discuss the nuances of clustering in scenarios involving multiple Nation State actors in the same environment, using distinct patterns of life to identify unique hunt opportunities.

Finally, we will emphasize the importance of tracking clusters over time. Covering how our CTI team utilizes statistical analysis and data viz tools to monitor the discrete aspects of clusters and initiate the creation of new clusters programmatically. Concluding with how tool specific hunts identified new intrusions, and how pivoting to infrastructure hunting, where we successfully predicted what threat actor C2 infrastructure would look like, led to the identification of vulnerabilities in security vendor programs, & tracking the actor over time.

About the Presenters:

Mark Parsons, Sophos, Senior Threat Hunter. Mark is a Threat Hunter for the Sophos Managed Detection and Response Team. He focuses on identifying new threats at scale, and has presented on his identifications at Black Hat USA, and PivotCon.

Colin Cowie, Sophos, Senior Threat Intelligence Analyst. Colin is a Threat Intelligence Analyst for the Sophos Managed Detection and Response team. He focuses on detecting emerging threats, threat actor identification, and incident response. In past roles, he has worked in the financial sector performing penetration testing as well as in mobile forensics for law enforcement.

Three Decades of Network Security Evolution

Presented by Vern Paxson from Corelight / UC Berkeley

30 years ago the Network Security landscape looked completely different, and subsequently at times it has evolved in quite striking ways. In the thick of the ongoing work of defending networks, it can be hard to step back to see the larger picture of how threats and defensive technologies change over time. How quickly do new threats arrive? How long does it take to actually deploy new security improvements? What do attackers probe for? How much encryption is actually in use, and how up-to-date is it?

We'll provide a look at this landscape by drawing upon years – and in some cases decades – of huge datasets from a dozen large operational environments, illustrating how some changes are lightning fast – and others, glacially slow.

About the Presenter:

Vern Paxson is co-founder and Chief Scientist of Corelight, a cybersecurity company based on network monitoring technology he and colleagues developed for many years. He is also a Professor of the Graduate School at the University of California, Berkeley, and formerly held staff and leadership positions at the International Computer Science Institute and the Lawrence Berkeley National Laboratory. At Corelight he leads a team of 24 researchers and developers in creating high performance algorithms for finding and contextualizing Internet attacks. He is a

Fellow of the ACM and has received ACM's SIGCOMM Award for lifetime achievement, the ACM Grace Murray Hopper Award, the IEEE Internet Award, and the Facebook Internet Defense Prize.

Minting Silver Bullets is Challenging

Presented by Josh Brown-White from Microsoft

With the advent of AI coding assistance such as GitHub Copilot there has been the obvious interest in using AI as a silver bullet to automatically correct source code to fix security vulnerabilities. Unfortunately minting this silver bullet is far more complex than simply calling the Azure AI APIs, and the organizations who have been rushing AI generated automatic remediations have seen results ranging from poor to disastrous. Modifying existing code behavior automatically is a very different exercise than a developer guiding Copilot to generate net new code, but the failures to understand those differences have resulted in quite poor outcomes. Many AI fix suggestions today catastrophically alter code, while only a minority of suggestions even technically correct the issues. Only a tiny percentage can be safely merged unmodified in a Pull Request. Most of the current attempts are causing more harm than utility.

This talk will explain the challenges that need to be accounted for when using generative AI to modify existing source code to correct security vulnerabilities, and detail how combining generative AI with existing deterministic analysis techniques can yield far better results. The future where the system not only automatically detects security vulnerabilities in source, but automatically remediates many (though not all) of them is attainable. It turns out that the reason why silver bullets aren't common is because they take a fair bit of work to make, but they can be made if you are willing to put in the work.

About the Presenter:

Josh is a Principal Security Lead at Microsoft on their Secure Development Lifecycle team, where he leads Microsoft's cross company usage of Static analysis, directs a team of security researchers and detection authors who leverage the current state of the art with Static Analysis and R&D new methods to further extend it, orchestrates cross company efforts around the security of Microsoft's products and services, and advises on security relevant public policy for Microsoft. Prior to this role Josh has been a Trustworthy Computing advisor for Windows, Azure, SQL Server, Windows Phone (RIP), and Windows Embedded, a product security architect at the payroll company ADP, a security analyst for FedEx, and in the distant past a developer on several now long dead technologies. Josh is also a contributor to SAFECode including contributing authorship to the SAFECode Security Fundamentals and SAFECode Threat Modeling guidance.

MSTIC Ghost Stories - A Threat Intelligence Year in Review

Presented by Rachel Giacobozzi from Microsoft

The Microsoft Threat Intelligence Center (MSTIC) discovers, tracks and disrupts the world's most sophisticated threats, from state-aligned actors to cyber criminals. MSTIC uses an adversary centric approach to understand risk to Microsoft and customers and apply that insight to detect, protect, and respond to the threats.

In this talk, we will look back over the last year of technical threats and share insights on the evolving threat landscape. The observations delivered in this talk ultimately seek to inform the future, empowering global people and organizations to achieve more, and do so securely. Examples covered in this presentation include:

- EURASIA/Russia Nation State Threats
- APAC
- MEA

About the Presenters:

Rachel Giacobozzi from Microsoft

PyRIT: From LLM Security Research to Practical Attacks

Presented by Richard Lundeen from Microsoft

The world of LLM security is wild. Papers are published at a rate where it's hard to keep up. And because of the probabilistic nature of LLMs, these attacks can often be combined and applied in unpredictable and interesting ways.

PyRIT is an open-source library developed by the Microsoft AI Red Team, and it aims to tackle this problem. PyRIT is built for researchers, pentesters, and defenders to help assess the robustness of LLM applications against different attacks and harm categories. The main scenario this talk focuses on is how PyRIT can be used to take research findings, combine them, and actually use them against a broad range of LLM endpoints. We want PyRIT to be the go-to tool researchers use to make proof-of-concept attacks and adapt them to real-world scenarios.

This talk walks through some of the most impactful public research we've incorporated into PyRIT - and how you can too! We'll show patterns of how we included external research like "tree of attack", "PAIR", "Crescendo", and "persuasive adversarial prompts". Using PyRIT for testing LLM security goes beyond downloading Hugging Face models and checking their security. We'll talk about how these proof-of-concepts can be included so researchers can test their entire pipelines against these attacks (including their derivations).

About the Presenter:

Richard Lundeen; Principal Software Engineering Manager at Microsoft; I lead an interdisciplinary group of red teamers, ML researchers, and developers to create software that emulates real-world attacks and finds failures against Microsoft's big bet AI systems. Our open source repository is here: <https://github.com/Azure/PyRIT>

SLIP: Securing LLMs IP Using Weights Decomposition

Presented by Adam Hakim and Lev Greenberg from Microsoft

Large language models (LLMs) have recently seen widespread adoption, in both academia and industry. As these models grow, they become valuable intellectual property (IP), reflecting enormous investments by their owners. Moreover, the high cost of cloud-based deployment has driven interest towards deployment to edge devices, yet this risks exposing valuable parameters to theft and unauthorized use. Current methods to protect models' IP on the edge have limitations in terms of practicality, loss in accuracy, or suitability to requirements. In this paper, we introduce a novel hybrid inference algorithm, named SLIP, designed to protect edge-deployed models from theft. SLIP is the first hybrid protocol that is both practical for real-world applications and provably secure, while having zero accuracy degradation and minimal impact on latency. It involves partitioning the model between two computing resources, one secure but

expensive, and another cost-effective but vulnerable. This is achieved through matrix decomposition, ensuring that the secure resource retains a maximally sensitive portion of the model's IP while performing a minimal amount of computations, and vice versa for the vulnerable resource. Importantly, the protocol includes security guarantees that prevent attackers from exploiting the partition to infer the secured information. Finally, we present experimental results that show the robustness and effectiveness of our method, positioning it as a compelling solution for protecting LLMs.

About the Presenter:

Adam Hakim, Microsoft, Data Science Team Lead. Adam holds a PhD in Neuroscience from Tel Aviv University and has over a decade of experience working and consulting as a Data Scientist and AI Researcher. He initially worked with various startups and later joined Microsoft, where he is leading AI security projects for large language models (LLMs).

Isolation or Hallucination? Hacking AI Infrastructure Providers for Fun and Weights

Presented by Hillai Ben-Sasson and Sagi Tzadik from Wiz

More and more companies are adopting AI-as-a-Service solutions to collaborate, train and run their artificial intelligence applications. From emerging AI startups like Hugging Face and Replicate, to mature cloud companies like Microsoft Azure and SAP – thousands of customers trust these services with their proprietary models and datasets, making these platforms attractive targets for attackers.

Over the past year, we've been researching leading AI service providers with a key question in mind: How susceptible are these services to attacks that could compromise their security and expose sensitive customer data?

In this session, we will present our novel attack technique, successfully demonstrated on several prominent AI service providers – including Hugging Face and Replicate. On each platform, we utilized malicious models to break security boundaries and move laterally within the underlying infrastructure of the service. As a result, we were able to achieve cross-tenant access to customers' private data, including private models, weights, datasets, and even user prompts. Furthermore, by achieving global write privileges on these services, we could backdoor popular models and launch supply-chain attacks, affecting AI researchers and end-users alike.

Join us to explore the unique attack surface we discovered in AI-as-a-Service providers, and learn how to mitigate and detect the kind of vulnerabilities we were able to exploit.

About the Presenter:

Hillai Ben-Sasson is a security researcher based in Israel. As part of the Wiz Research Team, Hillai specializes in research and exploitation of web applications, application security, and finding vulnerabilities in complex high-level systems. Hillai is a frequent speaker in security conferences, and has been recognized in MSRC's Most Valuable Researchers leaderboard.

Breaking LLM Applications - Advances in Prompt Injection Exploitation

Presented by Johann Rehberger from embracethered.org

Prompt Injection is a novel security threat that impacts large language model (LLM) applications. Confidentiality, Integrity, and Availability can all be impacted by a successful prompt injection exploit. This talk showcases many real-world exploit examples in well-known LLM applications and chatbots.

Specifically we will do deep dives on the following threats:

- * Misinformation, Scams and Phishing (including advanced attacks such as conditional prompt injection payloads)
- * Automatic Tool Invocation without human in the loop
- * Data Exfiltration Techniques
- * Attacks on LLM memory and how an adversary can achieve persistence
- * ASCII Smuggling, learn how LLMs can craft invisible text and decode hidden secrets and understand what that means for the security of your LLM application

Additionally, for each threat the details on how vendors mitigated vulnerabilities, including in systems such as OpenAI ChatGPT, Microsoft Copilot, GitHub Copilot Chat, Anthropic Claude, Google Bard/Gemini, Google NotebookLM and many others.

About the Presenter:

Johann Rehberger, Red Team Director. Johann Rehberger has over twenty years of experience in threat modeling, risk management, penetration testing, and red teaming. As part of his many years at Microsoft, Johann established a Red Team in Azure Data and led the program as Principal Security Engineering Manager. He also built out a Red Team at Uber, and currently is Red Team Director at Electronic Arts. He enjoys providing training and was an instructor for ethical hacking at the University of Washington. Johann is the author of the book "Cybersecurity Attacks - Red Team Strategies", and holds a master's in computer security from the University of Liverpool. He regularly blogs about his research at <https://embracethered.com>

Lessons Learned from Red Teaming 100 Generative AI Applications

Presented by Blake Bullwinkel from Microsoft

This talk covers the big lessons learned by the Microsoft AI Red Team in identifying safety and security vulnerabilities in flagship AI systems like Bing Copilot, Security Copilot, M365 Copilot, and models such as GPT-4, DALL-E, and the Phi series:

1. Prompt Injection gets all the attention, but traditional security failures are still top billing (example case study: credentials in Copilot source code, code execution via jailbreak in Code Interpreter)
2. As models get better, risk evolves (case study: GPT-4o which supported audio, video modalities had to be assessed for its ability to have romantic relationship with user)
3. LLM Guided Red Teaming can help us cover more of the risk landscape but is still finicky. Here we walk through an example of how our OSS automation tool PyRIT helped with saving close to 160 hours of manual probing, but how the scorer we used in evaluating frequently broke when we did RAI red teaming.
4. No free lunch in making AI systems safe: Tradeoffs that we have observed (example: in a facial recognition model, the more attempts were made to suppress the model from observing the face, the more the model focused on clothing. In another example, we found that smaller models are more immune to jailbreaks compared to larger counterparts since they
5. The difficulty in making AI systems safe: simple attacks have large impact (we show how a simple jailbreak could lead to dropping tables in production database that had Copilot turned on) and the inability to distinguish inadvertent failures and intentional failures

About the Presenters:

Blake Bullwinkel is an Offensive Security Engineer on the AI Red Team (AIRT) at Microsoft and previously worked as a Data Scientist in Cloud+AI. On the AIRT, he leads red teaming of the Phi-3 series of language models, conducts research into novel adversarial attacks, and tests a variety of generative AI products for harmful content and security vulnerabilities. He graduated with a masters in Data Science from Harvard University.

Automate AI Red Teaming in your existing tool chain with PyRIT

Presented by Joris de Gruyter and Shiven Chawla from Microsoft

Microsoft's open source tool PyRIT provides a set of extensible APIs and scripts that can be used for red teaming AI applications that use LLMs. The strength of the tool lies in the continuous addition of new techniques to red team LLMs as research is published. To overcome the challenge of learning a new tool set or potentially a new language (Python), we present PyRITShip – a layer on top of PyRIT to enable integration easy into existing tools, which can be run locally in Python directly, or using a container.

In this session we will discuss the basic uses and features of PyRIT. We then move onto a live demo of our BURP Suite Extension to enable LLM red teaming using BURP's repeater module with PyRIT as the driver. Finally, we will outline some future ideas for extensibility in other tool sets, and a call to action on making PyRIT work for you and how you can contribute.

About the Presenters:

Joris de Gruyter, Senior Offensive Security Engineer at Microsoft. Joris has a strong background in engineering and product management in developer tools for business applications. He rolled into AI Red Teaming at the product level as a Responsible AI Champion inside Microsoft's Power Platform group, and is now an offensive security engineer in the AI Red Team.

Shiven Chawla, Senior Offensive Security Engineer at Microsoft. Shiven has worked in security engineering at several Fortune 500 companies. His engineering skills led him to working on machine learning related projects at Amazon and Microsoft, ultimately becoming an AI Red Teamer before it was cool.



BlueHat 2024

Lightning Talk (15 min) Abstracts

Oct 29-30
2024
Redmond
WA, USA

World of Scams - A systematic analysis of online scams using the Scam Tactics and Techniques Framework

Presented by Amit Tambe from F-Secure

Online scams are a multi-billion-dollar industry that targets consumers on the Internet. According to Global Anti-Scam Alliance, over \$1 trillion is lost to scams worldwide.

Problem - The threat landscape news is fraught with online scammers targeting consumers every day with ingenuity. While several cybersecurity players have attempted to analyse the scam landscape, these efforts are at best, ad-hoc. There is no single systematic approach that can describe to the reader, in detail, all the techniques and methods used by scammers. Whereas a scam taxonomy exists, it serves to provide examples of rather than defining systemically the tactics and techniques (T&T) adopted by scammers. Defining such a systematic framework can go a long way in enabling defenders to protect consumers against scams.

Our paper showcases two unique contributions:

- Scam kill chain - sequence of events involved in a scam attack on a consumer
- Scam T&T framework - covers a comprehensive set of tactics and techniques used by scammers in the wild. The framework has been created based on our wide experience in the security field, including scams. Inspired from the MITRE® ATT&CK Matrix, our framework gives a thorough overview of steps taken by scammers to achieve their goals. A top-level tactic maps to a more detailed technique in the scam kill chain to characterise a mini-goal achieved by the scammer. Techniques represent the methods used by a scammer to achieve that specific mini-goal (i.e. tactic). This framework consists of 8 tactics and a total of 31 techniques.

In the presentation, we will see how the kill chain and “scam tactics and techniques framework” can be applied to conveniently visualize the progress of any scam, starting from how a user is targeted and ending with how money is cheated out of victims. Additionally, we will also see how the framework can be utilized by organisations to assess and thus improve their security posture against scams.

F-Secure

About the Presenter: Amit Tambe is a researcher at F-Secure Corporation, and primarily focuses on research about prevalent threats and scams targeting consumers. Additionally he also analyses Android malware

My Best Frenemy: A Synergy Between Red Team and Blue Team in Oracle's SaaS Security

Presented by Svetlana Gaivoronski and David B. Cross from Oracle

At Oracle, we maintain a robust security posture through the dynamic synergy of our Red Team and Blue Team, functioning as best frenemies. While the Red Team simulates attackers by seeking out vulnerabilities and testing our defenses, the Blue Team serves as vigilant defenders, monitoring systems, detecting threats, and responding to incidents. However, the true challenge lies in managing the human aspects of this relationship.

The success of a Red Team exercise can often feel like a failure for the Blue Team, leading to potential friction and demoralization. To navigate these human challenges, we have fostered a culture of communication and empathy. After each exercise, both teams come together for debriefs, where they share insights and discuss what worked and what didn't. This approach ensures that every Red Team success is viewed as an opportunity for Blue Team improvement, fostering a sense of shared accomplishment.

When the Blue Team detects Red Team activities, they are encouraged to monitor and analyze rather than immediately shutting them down. This strategy allows the Blue Team to learn and enhance their defensive measures while understanding the tactics used by potential adversaries. It's a controlled environment where both teams can test and improve their skills without fear of failure.

By maintaining a clear separation yet promoting collaboration, we create a seamless integration that strengthens our security posture. This careful balance helps manage the emotional aspects of the Red and Blue Team relationship, ensuring continuous improvement and a robust security framework.

In this talk, we'll cover all these challenges and the practical steps we're taking to overcome them, as well as share some anecdotes from day-to-day Red/Blue collaboration.

About the Presenters:

David B. Cross, Senior Vice President, Chief Information Security Officer (CISO) of Oracle SaaS Cloud | Gartner Peer Community Ambassador. David first started his work in security with his five years' active-duty service with the aviation electronic warfare community of the United States Navy. David was awarded with numerous honors including a Navy Achievement Medal, Southwest Asia Service Medal, Armed Forces Service Medal and NATO medal for his combat-based tours. David is now the SVP and CISO for the Oracle SaaS Cloud Security organization. Previously, David was a Director and built the Google Cloud Security Engineering organization for 3 years with his preceding 18 years spent with Microsoft in numerous security platform, cloud, product and engineering leadership roles.

Svetlana Gaivoronski, Director of Security Engineering, Oracle SaaS Cloud. As the Director of Security Engineering at Oracle, Svetlana leads the Detection and Response Team (DART) Automation project, focusing on detection engineering, compromise investigation, and automation. With a PhD in Cyber Security, her background includes founding and architecting the Automatic Forensics Triage Service at Microsoft, as well as contributing to Azure Sentinel. Svetlana holds multiple papers and patents in intrusion detection, forensics, and SDN networks and is currently pursuing the LEAD certificate at Stanford University. Her mission is to solve complex security challenges while mentoring and empowering others in the security community.

Lessons Learned: Scaling Out Securing Open Source

Presented by Zach Steindler from Microsoft

Programming language package repositories are juicy targets for attackers as they serve billions of requests per day. For the same reason, it's a great place for defenders to see high impact from security capabilities. And yet, each package repository ecosystem has unique community values and architecture - so how do you support developing security capabilities in ecosystems you aren't familiar with? Over the past year, the OpenSSF Securing Software Repositories Working Group has done so by providing roadmaps (like the "Principles for Package Repository Security" co-published with CISA), implementation guidance (like our biggest success to date

"Trusted Publishers for All Package Repositories"), and partnering with other organizations to fund people in these ecosystems to implement these capabilities. As a result, we've seen tremendous progress in package repository security capabilities in the past year, and a healthy roadmap for what's ahead. These new security capabilities, as well as the support framework we used to facilitate their implementation in open source ecosystems, can inform your organization's security roadmap.

About the Presenter: Zach Steindler, GitHub, Principal Engineer. Zach is a member of the OpenSSF's Technical Advisory Council and co-chair of the Securing Repositories Working Group which helps coordinate security improvements in programming language package repositories like PyPI and Ruby Gems. In early 2024 he co-published "Principles for Package Repository Security" with US CISA. He works at GitHub on securing software development for open source and enterprises. Away from computers he enjoys gardening and welding.

A Security Engineer's Journey: Creating a Developer-Friendly Security Tool

Presented by Susan Krkasharian from Microsoft

In this session, I will share my journey as a junior security engineer developing "AntiSSRF," a secure software library designed to protect Microsoft services against Server-Side Request Forgery (SSRF) attacks. SSRF vulnerabilities, now present in both the OWASP Top 10 and OWASP API Top 10, are critical because they allow attackers to make unauthorized requests from a server, potentially exposing sensitive data and internal services. In Azure, this can lead to severe breaches, including unauthorized access to cloud resources and data exfiltration.

Targeted at a broad audience, including security engineers and software engineering managers who aim to enhance application security through code, this session will provide insights into the building of AntiSSRF, highlighting the importance of providing developers with robust tools to safeguard applications, as developers may not always be aware of the intricate nuances required to protect against such vulnerabilities.

I will share the journey of developing AntiSSRF, from initial ideation and prototyping to production-ization. I will cover challenges faced, such as securing DevOps pipelines and maintaining compliance with Microsoft SFI security standards. I will also discuss the strategies used to drive adoption, including evangelism, addressing performance concerns and code defects, and leveraging static code analysis to surface additional services that could benefit from AntiSSRF.

This lightning talk aims to provide valuable insights into the overall effort required to create and maintain a secure code library, highlighting that creating a secure library involves much more than just writing code; it requires a holistic approach to security, compliance, and user engagement.

About the Presenter:

Susan Krkasharian, Microsoft, Security Software Engineer 2. Susan Krkasharian began her journey at Microsoft as an intern and transitioned to a full-time role after earning her degree in Computer Science from UCLA. Susan has been on the DevSec team at Microsoft for a little over four years now, during which she developed comprehensive strategies for SSRF vulnerability prevention, detection, and mitigation.

Firmware Security: The Middle Child of Security

Presented by Nithin Sade from Google

Firmware Security can often be de-prioritized in favor of other security focus areas. This lightning talk aims to bring attention to some of the recent firmware threats and challenges around managing firmware for enterprises. For e.g. BlackLotus, managing BIOS settings, Firmware Updates, Computrace, etc. At the end, this talk attempts to urge more standardization across OEM's to help improve Firmware Security space.

About the Presenter:

Nithin Sade - Google LLC - Security Engineer. Nithin Sade is a Security Engineer at Google, dedicated to building secure and user-friendly systems. His passion lies in crafting scalable security solutions that minimize risk without compromising user experience. With expertise in endpoint security, Nithin strives to create a safer digital world.

Getting "In Tune" with an Enterprise: Detecting Microsoft Intune Lateral Movement

Presented by Brett Hawkins from IBM

Organizations continue to implement cloud-based services, a shift which has led to the wider adoption of hybrid identity environments that connect on-premises Active Directory with Microsoft Entra ID (formerly Azure AD). To manage devices in these hybrid identity environments, Microsoft Intune has emerged as one of the most popular device management solutions. Since this trusted enterprise platform can easily be integrated with on-premises Active Directory devices and services, it is a prime target for attackers to abuse for conducting lateral movement and code execution.

This research will give a background on Microsoft Intune (Intune) and how it is being used within organizations. It will then show how to use this cloud-based platform to deploy custom Windows applications to achieve code execution on user devices. Additionally, this research includes the public release of new Microsoft Sentinel rules to help defenders detect the usage of Intune for lateral movement and defensive hardening guidance for the Intune platform.

About the Presenter:

Brett Hawkins, Capability Lead, Adversary Services, IBM X-Force Red. Brett has been in Information Security for several years working for multiple Fortune 500 companies across different industries. He has focused on both offensive and defensive disciplines, and is currently on the Adversary Services team at IBM X-Force Red. He holds several industry recognized certifications, and has spoken at several conferences including Black Hat, DerbyCon, Wild West Hackin' Fest, BSides, and Hackers Teaching Hackers. Brett is also a member of the open-source community, as he has contributed to or authored various public tools, such as SharPersist, DueDLLigence, SCMKit, ADOKit and InvisibilityCloak.

Ransomware Resilience: Turning the Tide Against Cyber Extortion

Presented by Tom Williams from True Zero Technologies

Ransomware attacks have become a pervasive threat to global organizations, causing substantial financial losses over the past decade. The increase in successful ransomware attacks is driven by multiple factors and the issue has proven challenging to track by the industry and researchers. A major obstacle to accurately quantifying the impact lies in the fragmented nature of available data and the lack of centralized oversight. This call to action takes a multidisciplinary perspective to analyzing cybersecurity risk and ransomware threats across various industries and provides insights into the vulnerability profiles of sectors targeted by ransomware operators. Furthermore, this study examines the challenges associated with mitigating these risks, including the technological evolution of attack variants, sophistication of tooling in the AI/ML era, and the importance of readiness.

About the Presenter:

Dr. Tom Williams III, True Zero Technologies, Principal Systems Architect. Dr. Tom Williams is a security practitioner with over 15 years of experience in protecting organizations from cyber threats. Throughout his career, Tom has led highly technical teams and advised clients on risk management, cyber resilience, and security strategy for Fortune 500 companies and DOD agencies. His passion for the technology sector is reflected in his ongoing research and thought leadership on adjacent business and technology topics.

SafeChatAI: Enhancing Cybersecurity Awareness Using Artificial Intelligence.

Presented by Ayobami Olatunji from Microsoft

The proliferation of digital devices and increased online communication have led to numerous cyber threats, particularly social engineering and phishing attacks, resulting in identity theft and impersonation. A significant factor contributing to the rise in online fraud is the lack of adequate cybersecurity awareness. Although various platforms and frameworks exist to enhance cybersecurity, they often fall short in providing real-time protection. This paper proposes SafeChatAI, a system designed to deliver real-time alerts when sensitive information, such as a Social Security Number (SSN), is requested via chat or email by an unknown individual. Additionally, SafeChatAI allows users to report suspicious activities directly to the Federal Bureau of Investigation (FBI), verify the identity of potential scammers, flag accounts as scammers, and notify other users of flagged individuals. By combining AI-driven monitoring with user participation, SafeChatAI aims to create a safer online environment and reduce online fraud.

About the Presenter:

Ayobami Olatunji from Microsoft

AI's got Muffins- the RAG-a-muffins!!!

Presented by Vivek Vinod Sharma from Microsoft

Retrieval Augmented Generation [RAG] is the heavily used for building GenAI apps. But what happens when RAG gets poisoned! We uncover importance of RAG and how best to secure it from misuse/abuse in GenAI -LLM Application development

About the Presenter:

Vivek Vinod Sharma from Microsoft

Entitlements on macOS X and why they matter

Presented by Yves Younan from Cisco Talos

This short talk goes over what entitlements are on macOS, how they work and why they are important. They are different from how traditional permissions have worked in many operating systems. It is important for developers to understand how to use them. To prevent vulnerabilities like:

<https://blog.talosintelligence.com/how-multiple-vulnerabilities-in-microsoft-apps-for-macos-pave-the-way-to-stealing-permissions/>

The talk will also discuss the issues that Talos discovered and why they're important to the security of macOS users.

About the Presenter:

Yves Younan; Cisco Talos, Sr. Manager Vulnerability Discovery & Research. Yves Younan leads the Vulnerability Discovery & Research team within the Talos Security Intelligence and Research Group at Cisco. Cisco purchased Sourcefire in 2013.

Prior to joining Sourcefire's Vulnerability Research Team, he worked as a Security Researcher with BlackBerry Security. Before joining BlackBerry, he was an academic, founding the Native Code Security group within the DistriNet research group at the Katholieke Universiteit Leuven in Belgium.

He received a PhD in Computer Science from the Katholieke Universiteit Leuven. His PhD focused on efficient mitigations against code execution attacks.