10 Hot Talks From Black Hat US 2022

These 10 sessions are a good predictor of important issues to watch in the next year.

Brought to you by informa tech

FEATURE

10 Hot Talks From Black Hat US 2022

These 10 sessions are a good predictor of important issues to watch in the next year.

By Ericka Chickowski, Contributing Writer, Dark Reading

Security researchers, red teamers, and defenders alike came to Black Hat USA 2022 — either in-person in Las Vegas or via the virtual platform — to hear about the latest zero-day vulnerabilities, the newest offensive security tools, and the most sophisticated methods of defense driving the bleeding edge of cybersecurity today.

There were more than 100 talks over the course of two days that spanned a range of topics and themes, including identity and access management threats, container and cloud vulnerabilities, and, of course, the software supply chain. Additionally, after months of geopolitical conflict in Ukraine, cyberwar and infrastructure defense implications were front and center. These themes frequently cropped up in the 10 hottest talks from this year's Black Hat conference, described below.

Identity and Access Management Threats

Hot Talk #1: Elevating Kerberos to the Next Level

There was identity and access management (IAM) research aplenty at this year's show, with numerous talks taking on authentication weaknesses in endpoints, design flaws in identity platforms, and more. One key talk in this mix was a rundown of research done by James Forshaw, security researcher for Google Project Zero, and Nick Landers, head of adversarial R&D for NetSPI. The research demonstrated how Kerberos can be abused to carry out a wide range of local privilege escalation (LPE) attacks.

Kerberos is the main authentication protocol for Windows enterprise networks, but, to date, most security research has focused on attacking Kerberos to carry out remote exploits or aid in lateral movement across the network. Forshaw and Landers took a fresh approach this year and spent months digging into the protocol.

"Kerberos is a complicated protocol, and there's been a lot of work in deconstructing how it operates between hosts and servers but a little bit less around the nuances of the protocol, especially in how a local machine handles it," Landers said.

One of the main bugs the pair examined was Privileged Attribute Certificate (PAC) validation bypasses using user-to-user authentication, but they explained that the issues go deeper than that.

"We found issues with UAC — James found issues in code execution with remote credential guard," said Landers. "We looked at service account elevation, and at the end of our slides we have a huge array of assorted issues and bugs across LSASS (Local Service Authority Server Service) and the Kerberos protocol, all of which have been fixed in the last three to four months."

Defenders need to get a fundamental understanding of Kerberos, in spite of its complications.

Landers said one of the big takeaways from his discoveries is that defenders need to get a fundamental understanding of Kerberos, in spite of its complications.

"It's got a lot of confusing cryptography, and the better defenders start to understand how that cryptography comes into play — and the full scope of the specification for Kerberos and its implementation — the better off they'll be."

Hot Talk #2: IAM the One Who Knocks

On the cloud-based identity and access management (IAM) front, a marquee presentation was a compilation of research conducted by Igal Gofman, head of research for Ermetic, tackling the security gaps that exist in the connections between multicloud and hybrid cloud infrastructure components.

The more that organizations with multicloud infrastructure environments patch together different identity and access management methods to work with each cloud provider's IAM tools, the more room there is for misconfigurations, improper storage of security tokens and other credentials, and other toxic combinations of permissions. This pea soup of different IAM controls can make it easier for attackers to escalate privileges and move laterally across different infrastructure.

"At one time, those [security tokens] didn't give much to the attacker beyond what was on a local machine," Gofman said. "But now, those security tokens have much more access because everyone in the last few years moved to the cloud and has more access to cloud resources."

Gofman discussed how organizations can use native IAM services provided by cloud providers to improve visibility and control over a multicloud identity portfolio, but noted that security teams have to be mindful of the differences between each tool's capabilities and defaults to ensure blind spots don't persist. He also demonstrated Access Undenied, an open source tool developed by his firm as a third-party option to bridge some of those gaps.

He recommended that organizations start out of the gate by using cloud providers' logging features to start building a list of which users and machine identities are actually operating across their environments.

"These tools are not actually used extensively, but they're good options to better understand what's going on in your environment," he explained. "You can use logging to reduce the attack surface, too, because you can see exactly what users are using and what permissions they have."

Container Escapes

Hot Talk #3: The COW Container on Windows Who Escaped the Silo

This year's slate of sessions at Black Hat reflected the growing importance of container security in the enterprise, with several talks and tools geared toward container risks, particularly container escapes.

"You can use logging to reduce the attack surface, too, because you can see exactly what users are using." —Igal Gofman, Ermetic

Case in point was one presentation given by SafeBreach researcher Eran Segal. He dug into the inherent architectural design problems in how Windows containers are isolated from real host settings. The issues make it easy for attackers to break down the barriers of container isolation, Segal said.

"Windows containers isolated as 'process isolation' are not isolated well, and it is possible to impact the host from inside," explained Segal, who added that Linux kernel architecture is built to handle containers better than Windows' architecture. "It [container isolation] is harder to implement in Windows. Microsoft has a workaround for that named Hyper-V containers, but they are not really containers — they are more similar to VMs than containers."

Segal demonstrated to the audience how this design issue makes it possible to gain a permissions system inside a container, cause a denial of service to the host, and access the entire kernel memory — from where it's possible to start moving laterally because kernel memory contains passwords. He shared details on how process-isolated Windows containers work and how researchers can take his vulnerabilities and hunt for more.

Software Supply Chain

Hot Talk #4: RCE-as-a-Service: Lessons Learned From Five Years of Real-World CI/CD Pipeline Compromise

Software supply chain security was understandably a big undercurrent running throughout the show. One talk by Iain Smart and Viktor Gazdag of NCC Group addressed head-on one of the elephants in the room when it comes to modern DevOps: the insecurity of many continuous integration/continuous delivery (CI/CD) pipelines. For today's most efficient development teams, CI/CD pipelines and the tools that comprise them are akin to automated factory lines. They allow organizations to build and refactor software more quickly than ever, but abuses of the functionality in these pipelines could essentially turn a development team's workbench into what Smart and Gazdag call "remote code execution as a service."

"Attack outcomes can include stealing source code or intellectual data, backdooring an application that is deployed to thousands of customers, [and] gaining access to multiple environments such as development and production — both on-prem or in the cloud or both," Gazdag explained.

Some of the most common weaknesses he has uncovered in his research include how attackers can potentially embed their own SolarWinds-type flaw into the CI/CD pipeline, hardcoded credentials in source code management or version control systems, over-permissive roles, and a lack of audit, monitoring, or alerting to track or control what is occurring daily.

Gazdag recommends that organizations start allocating more time to create threat models for their pipeline and its tools, exploring how different environments connect, where boundaries exist, and how and where secrets are stored.

Hot Talk #5: Scaling the Security Researcher to Eliminate OSS Vulnerabilities Once and for All

Open source software (OSS) lies at the foundation of so much modern enterprise software. It's been a boon for faster, more efficient software builds, but it also poses one of the most existential threats to application security today. Many security advocates believe that the fate of application security and the

Gazdag recommends that organizations start allocating more time to create threat models for their pipeline and its tools.

software supply chain rests on how effectively the open source community and the security world can work together to shore up vulnerable code within open source software projects.

A presentation by Jonathan Leitschuh, the inaugural Dan Kaminsky Fellow at Human Security, offered a ray of hope on this front by providing a scalable method for the security research community to collaborate with open source maintainers to make it easier to address larger numbers of high-risk flaws at scale. His talk discussed the practical application of using automated bulk pull-request generation to smooth the process of fixing vulnerabilities that span numerous open source projects. This helps address the particularly thorny problem of how a single vulnerability can ripple across hundreds of different pieces of OSS due to open source dependencies.

Leitschuh has spent his time as the first Dan Kaminsky Fellow to refine methods using tools like CodeQL, GitHub's code query language, to scan for flaws across hundreds of thousands of OSS projects. He has also used tools like OpenRewrite, a style-preserving refactoring tool developed by Moderne, to help scale the triaging, reporting, and fixing of flaws. In one example, he demonstrated how a CodeQL query gave him a list of 900 repositories potentially vulnerable to Zip Slip, a flaw in unzipping zip files that can lead to remote code execution.

"Of the 900, we have made 86 Zip Slip fix pull requests to date, which means 86 critical security vulnerabilities that now have possible fixes," he said. Beyond that one instance, his work has generated almost 600 different pull requests to fix old and new vulnerabilities he has discovered in open source software, he added.

While he's only one researcher working on the problem, Leitschuh hopes that the security community armed with his tooling and method can start making a big, automated dent in the OSS vulnerability pool.

"It is possible to fix widespread and common security vulnerabilities at scale. We have the technology. All we need to do is leverage it," he says. "Fixing these vulnerabilities is not an interactive problem. We can solve it with math, science, technology, and security."

Systemic Security Labor Problems

Hot Talk #6: Bug Bounty Evolution: Not Your Grandson's Bug Bounty

For the past decade, bug bounties have been hyped as a way to increase the scalability of an application security program by putting crowdsourced researchers to work in finding previously unknown vulnerabilities. It sounds rosy on paper, but, in practice, many a security team has run into problems leaning on bug bounties to supplant the fundamentals of application security — the boring stuff like asset management, vulnerability management, and developer training — rather than using it as a valuable supplement.

This was the drive behind a talk by Katie Moussouris, founder and CEO of Luta Security, who took the stage to air the dirty laundry of bug bounties as they're run today. This includes problems in execution that the industry has struggled with from the start.

"I think that there's room for a ton of improvement, not just in how bug bounties are designed and executed, but also in the holistic picture of the ecosystem in which a bug bounty operates," Moussouris said. A bug bounty problem shouldn't be around to highlight the low-hanging fruit that can be discovered from traditional application security practices, added Moussouris, but instead be used for surfacing the complex, hard-to-find and harder-to-exploit flaws.

"We want organizations to be not just prepared to fix the

It sounds rosy on paper, but, in practice, many a security team has run into problems leaning on bug bounties to supplant the fundamentals of application security — the boring stuff like asset management, vulnerability management, and developer training — rather than using it as a valuable supplement.

bugs thrown over the fence in a vulnerability disclosure program or bug bounty program, but actually looking at their core security investments and using bug bounty programs as an indicator of health of their overall security program," she said. "Because, if you think about it, every bug is a symptom of an underlying disorder in their security system."

Moussouris also tackled the fact that bug bounties aren't necessarily a great gig for the bounty hunters, either.

"It's like the worst gig economy job you could possibly get — worse than an Uber or Lyft job because you get paid with every gig that you take with Uber and Lyft. You do not get paid for every single bug you find if you are a bug bounty hunter," she said.

In addition to offering tips for what individual organizations need to do to get the most out of bug bounties, she used her talk as a platform to discuss how the industry can expand the marketplace for security labor — for example, by using apprenticeship models and building a pipeline for developing talent and education around vulnerability remediation and application security resilience.

Threat Hunting

Hot Talk #7: The Open Threat Hunting Framework: Enabling Organizations to Build, Operationalize, and Scale Threat Hunting

While offensive security discoveries and demonstrations of new bugs will always hog the limelight at Black Hat, this year's show didn't forget about the defenders seeking out new techniques and frameworks to combat emerging threats.

One of the big highlights came from a team at IBM Security X-Force, which released the new Open Threat Hunt Framework (OTHF) to help security teams interested in starting a threat hunting program or taking their existing program to the next level. According to John Dwyer, head of research for IBM Security X-Force, the framework is the product of two years of work interviewing threat hunting teams to figure out the factors that differentiate high-performing threat hunting programs. This includes factors like goal setting and scoping of hunting activities, communication of value to company leadership, training and role development for hunters, and, most importantly, the development of repeatable processes.

The framework offers some practical guidance to the threat hunting community, which primarily leans on tooling to help them with the nitty-gritty technical details of conducting a single hunt or a series of hunts. These tools, however, don't widen the scope of that lens to offer ad-

vice in running a sustainable program.

"The industry has focused on the technical bits. But, at the end of the day, what we really want to do is drive impact," Dwyer said. "So, the things we are talking about is offering examples of how to build a process, how to use it, and how to achieve continuous improvement within your program. When organizations do that, they crush it."

AI/ML and Cybersecurity

Hot Talk #8: Malware Classification with Machine Learning Enhanced by Windows Kernel Emulation

Artificial intelligence (AI), machine learning (ML), and data science were a particularly hot topic category, with a whole track dedicated this year to both adversarial AI themes — attacking AI models and platforms for nefarious purposes — and themes focused on the use of advanced AI techniques to scale up security detection and defense.

Some of the adversarial AI research topics presented included exploration of potential attacks against graph neural networks (GNNs) that can steal proprietary model training data to steal the model's functionality itself, as well as a deep dive into how an AI programming tool trained on GitHub repository data could potentially sully the software supply chain with vulnerabilities it has learned from the insecure code it's trained on.

On the flip side, several researchers presented talks on how to improve the use of AI in cybersecurity defense. One examined weaknesses in AI detection mechanisms for picking up on deepfake audio and synthesized voices. Another cutting-edge talk investigated how the Generative Pre-trained Transformer 3 (GPT-3) natural language model can be used to solve tough security problems, such as querying SIEM data based on a natural language question and parsing the results into human-readable answers.

Arguably, one of the most immediately practical applications of ML was presented by Dmitrijs Trizna, a Microsoft researcher. He presented a way of using ML to combine traditional static analysis of malware with more dynamic analysis to classify new malware, improve detection rates, and decrease false positives beyond what any AI/ML classifiers currently can do.

In concert with the presentation, Trizna released the Quo Vadis ML model as an open source research prototype.

The heart of the method employed by the tool is the use of a Windows kernel emulator developed by Mandiant and running the emulation reports it generates through a neural network.

"The approach takes a portable executable file and processes it through a machine learning algorithm to get better results than we have right now in the public space," he said. The combination of static and dynamic analysis provides greater detection fidelity, and the machine learning makes it scalable, he added.

"[When you examine] statically it means you are analyzing without actually launching malware logic in any environment, whereas dynamic analysis detonates malware in a controlled environment [and] collects telemetry from what it does on a system — like network connections and file system manipulations," he said. "We show how this type of data can be used alongside the machine learning algorithm."

That added component of ML can help security teams scale up their analysis efforts, but ML is meant to be supplementary to the human thinking power of an analyst team, he explained. He believes that there's currently an alignment problem in ML-backed security as researchers try to figure out how to effectively apply ML to security problems. One of the big challenges is getting the right combination of data science skills and security fluency to come together in a single team or a single researcher. He hopes that more projects like his can bridge the gap so security researchers can take models and information and use them without needing deep knowledge of ML to apply them.

Cyberwar in Ukraine

Hot Talk #9: Industroyer2: Sandworm's Cyberwarfare Targets Ukraine's Power Grid Again

The cyberwar waged in concert with physical warfare in Ukraine was a weighty topic that was addressed with two big presentations at the show. The first was given by Robert Lipovsky, senior malware researcher at ESET, who unveiled details of his research into the activities of the Russian-

Insurers now require third-party attestation or hire experts to assess a client's cybersecurity infrastructure and incident response plans.

operated threat group Sandworm. Lipovsky has been tracking Sandworm's return to the proverbial well with more attacks against the Ukrainian power grid to support Russian aggression against the Ukraine.

"We found that Sandworm attackers for the third time tried to disrupt the flow of electricity in Ukraine," said Lipovsky, whose team worked closely with the Ukrainian CERT team in the spring fending off the attack.

In his talk, he covered the technical details around his team's reverse engineering of Industroyer2, the latest version of industrial malware that was previously best known for causing a major blackout in the Ukrainian power infrastructure back in 2016.

"This was a very interesting discovery because after the Sandworm attack in 2016 a lot of the industry was asking these questions, like what's next? [We've] heard from Sandworm since then, but we haven't seen power grid attacks for five years, until now, as there's a full raging war," Lipovsky said. He added that the good news was that the threat group was unsuccessful in its attempt, causing only a blackout that lasted about an hour.

This was puzzling to Lipovsky's team, as the internals of Industroyer2 showed higher ambitions than that, with sophisticated functionality that could have caused much worse damage.

Another part of the analysis he presented at the show was a dive into CaddyWiper, a destructive wiper that was deployed alongside Industroyer2.

"A wiper is a pure act of sabotage — it's like ransomware without the financial motivation," he said, explaining that this disruptive element targeted industrial workstations to disrupt operations at targeted facilities. "Sandworm has a thing with wipers, and we've been seeing an evolution of wipers since the beginning."

Hot Talk #10: Real Cyberwar: Espionage, DDoS Leaks, and Wipers in the Russian Invasion of Ukraine

Lipovsky's CaddyWiper analysis dovetailed with findings presented in a Ukraine cyberwar talk given by researchers Juan Andres Guerrero-Saade and Tom Hegel of SentinelOne.

The researchers examined a broad range of wiper malware strains employed in the aggression against Ukraine by Russian threat actors.

The duo also detailed cyberwarfare tactics like DDoS, espionage, and other disruptive cyber operations. One of the themes they explored was how messy cyberwar can look on the ground when physical warfare is ongoing.

"I think what we're learning is how these particular groups adapt to that situation, what the targeting actually looks like, and how little control [attackers] actually have on that environment," said Guerrero-Saade. "People think about cyberwar as the attacker being in complete control, and they know everything and they've tested everything. But it doesn't work that way. It turns out that when there's bombs falling and the power's going out and people aren't at their desks to click on a phishing email, things look very different."

He explained that the research into these messy activities is still nascent. As the security community and government work to understand how these cyber operators have conducted themselves in the Ukrainian conflict, it will give them food for thought on how to prioritize protection of critical infrastructure and commercial interests in preparation for potential future conflicts. For example, there's a world of difference between what happens when a wiper wreaks havoc at a news organization versus when it causes damage at an industrial organization or a financial system.

"At some point, you have to understand what the prioritization is, and with threat actors all of this comes down to telemetry. Do you have visibility?" he said. This is why it has been interesting to observe the rampant use of wiper malware, he said, adding that his team tracked at least seven strains in play within this conflict.

"What wipers are showing us is actually a failure of visibility. The wiper is burning the house down after you're done," he said. "The fact we are seeing all these operations at the wiper level means we've missed the initial intrusion, the lateral movement, data theft, or whatever they did. We just watched them burn the house down."

"We need to accept the fact that cyber operations are part of warfare, and these things are accompanying," he added. "The operations and we, as an industry, need to be paying attention and play our part in defending organizations, private and public, as well."

Find Out More

Black Hat USA 2022 adopted a hybrid format to make it possible for those who couldn't travel or who might have missed a few sessions to watch and rewatch sessions like these to maximize the lessons learned. Check out replays on Black Hat's YouTube page.

About the Author: *Ericka Chickowski specializes in coverage of information technology and business innovation. She has focused on information security for the better part of a decade and regularly writes about the security industry as a contributor to Dark Reading.*