SANS | Research Program

White Paper

SANS 2024 Top Attacks and Threats Report

Written by **Lee Crognale**

July 2024

AUDITBOARD

Introduction

Each year, the RSA Conference brings together like-minded practitioners, leaders, innovators, and vendors. The topics always center around the biggest threats encountered each year and how to protect your enterprise in a fluid landscape. This year, there was no shortage of capabilities, solutions, best practices, and lessons learned shared by some of the most influential minds in the business. Attention-grabbing topic areas included *security-focused software development and solutions for multi-cloud security*, *SOC automation*, *securing generative AI*, and *policy and compliance*. The dominant underlying theme in almost every area this year was the pros, cons, and use cases for artificial intelligence (AI) and machine learning (ML). If you made it to RSA this year, you've likely taken many of these key takeaways back to your place of work by now. If you weren't able to attend, we will break down some of the most formidable issues on the horizon, help you choose what solution may be best for your situation, and also highlight some of the discussions presented by the SANS panel of experts, who continually share their industry knowledge with RSA attendees.

Understanding the Current Threat Landscape

Let's first take a look at data harvested from the previous year. We are again using information collected from the Identity Theft Resource Center (ITRC), a nonprofit organization aimed at reducing the risks associated with identity compromise. This data consists of publicly reported breaches and illustrates the severity of the problem. In 2023, a new record high 3,205 compromises were reported, affecting more than 300 million individuals. Although the total victim count was less than previous years, the number of compromises increased over 70% from the previous year.[1] Some of the most noticeable increases were reported in the attack vectors associated with malware, zero-day attacks and other/unspecified attacks falling under the categories of cyberattacks and system and human error (see Figure 1). Also notable was the increase in supply chain attacks, at more than 2,700 last year, which is an alarming 2,600% increase since 2018. In their report, Verizon highlighted the "Exploitation of Vulnerabilities" as the leading cause of supply chain attacks.[2] See Figure 2.

How does this differ from previous years? Last year, we noted that although zero-day attacks are the attention grabbers, most exploited vulnerabilities are known and fixable. While this is still mostly true, we are seeing a shift in how AI is being used to accelerate the exploitation timeline, leaving organizations scrambling to identify issues, roll out patches, and perform necessary remediations. On average, just slightly more than 50% of identified vulnerabilities were remediated by the 60-day mark, painting a clear picture of how attackers are leveraging current technology to scan systems and exploit vulnerabilities faster than ever before. See Figure 3.

	2023	2022	2021	2020	2019	2018
Cyberattacks	2,365	1,584	1,611	877	928	754
Phishing/Smishing/IRC	438	467	537	383	490	379
Ransomware	246	293	353	159	83	53
Malware	118	73	141	103	112	102
Non-Secured Cloud Environment	14	10	24	51	15	15
Credential Stuffing	29	18	14	17	3	10
Unpatched Software Flaw	-	-	4	3	3	-
Zero Date Attack	110	8	4	1	-	-
Other	30	17	424	159	223	195
Not Specified	1,380	698	110	1	-	-
System & Human Error	729	163	179	153	231	261
Failure to Configure Cloud Security	25	18	54	58	56	49
Correspondence (Email/Letter)	380	55	66	55	89	121
Misconfigured Firewall	19	30	13	4	4	29
Lost Device/Document	53	7	12	5	19	17
Other	220	36	34	31	63	45
Not Specified	32	17	-	-	-	-
Physical Attacks	53	46	51	78	118	152
Document Theft	6	7	9	15	19	25
Device Theft	23	21	17	30	57	57
Improper Disposal	5	5	5	11	14	21
Skimming Device	9	6	1	5	4	12
Other	5	6	19	17	24	37
Not Specified	5	1	-	-	-	-
Data Leak	2	-	7	-	-	-
Unknown	56	8	12	-	2	8

Figure 1. Security Compromises by Attack Vector – ITRC 2023 YE Data

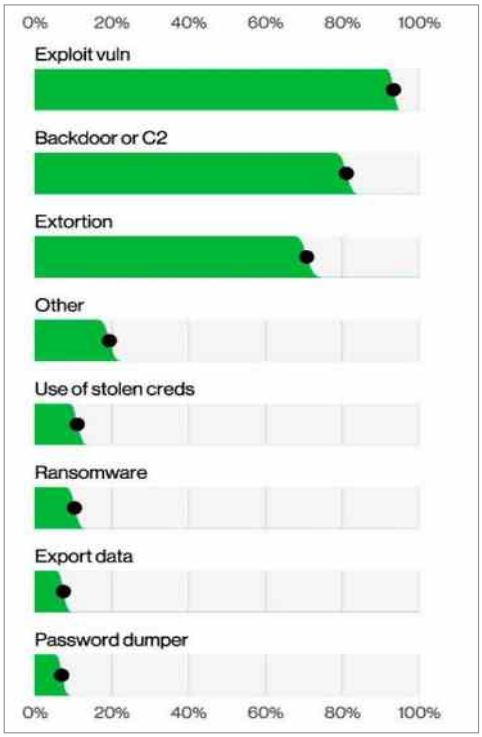

Figure 2. Supply Chain Attack Vectors – Verizon 2024 DBIR

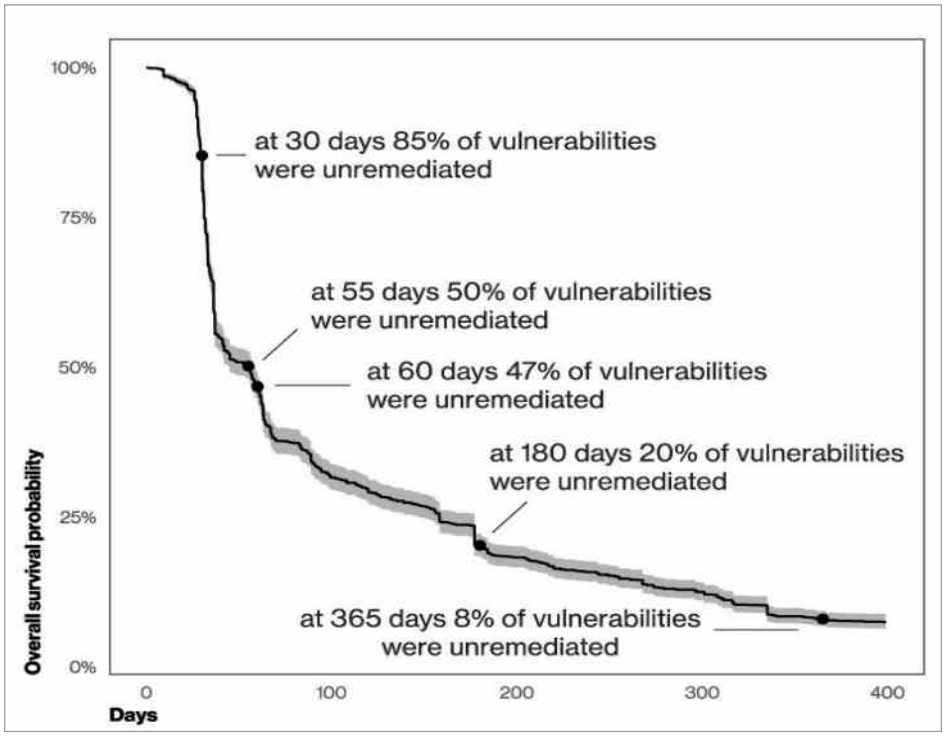

Figure 3. Timeliness of Remediation – Verizon 2024 DBIR

[1] "ITRC Annual Data Breach Report," 2023, www.idtheftcenter.org/publication/2023-data-breach-report

[2] "2024 Data Breach Investigations Report," www.verizon.com/business/resources/reports/dbir

Security-Focused Software Development and Solutions for Multicloud Security

SANS continued its 15-year tradition of highlighting some of the year's most concerning emerging threats. Dr. Johannes Ullrich, dean of research at the SANS Technology Institute College and founder of the SANS Internet Storm Center, kicked off the panel discussions by highlighting the security cost associated with technical debt. This is a problem affecting many organizations, and it's not just the enterprise applications but the entire security stack that is at risk, as Ullrich pointed out. He lamented over the issues that exist with legacy code hanging around, falling under your umbrella after an acquisition or reorganization. This can cause issues when the original development language isn't in practice anymore, the vulnerability isn't clearly understood, the code is outdated and not understood by the current workforce, and sadly nothing is well documented.

These small software issues compound into larger problems compromising the security stack entirely. To get to the root of this problem, some important questions you should be asking your organization include:

- Do we have legacy systems and source code?
- What is our current patch/update process?
- How do we deal with third-party integration?
- How good is our documentation regarding all the above?

Ullrich offered up his own advice regarding patching, developing, and maintaining software: "Apply patches incrementally as they're being released." Sometimes this meets with resistance because the library or source code being updated doesn't seem to add to or enhance your own needs. Over time, however, it is the collective, unpatched code that becomes ripe for a vulnerability, requiring a massive rewrite and a lot of developer hours and insight into the original, old code's purpose.

It may also come as a surprise that it's not just old technology that suffers from this problem. Companies sometimes rebrand old products, and code reuse may result in a product with slightly new features and only slightly fewer vulnerabilities.

His topic posed an interesting question for those suffering from similar circumstances: *Where and when is the best place to inject SECURITY into your safe coding practices?*

Dr. Johannes Ullrich
SANS Technology Institute Dean of Research, Internet Storm Center Founder

DevSecOps vs. SecDevOps: What Is the Difference and What Is Right for Your Organization?

To help organizations make informed decisions, it's best to elaborate on some of the following security approaches. The difference between DevSecOps and SecDevOps may seem inconsequential or an arbitrary placement of words, but the impacts on your organization are often directly affected by the implementation approach you choose. In fact, the arrangement of words is not arbitrary at all.

SecDevOps highlights security as being paramount. Some benefits of a security first approach include coding practices that implement security at each step. This includes both development and testing, so those vulnerabilities that are often exploited by attackers have been eradicated before the code ever hits the production environment. This is achieved through automation, which requires some initial overhead that may not necessarily stifle the latter approach. At first glance, this may seem like the best option before weighing the pros and cons, but the challenge of finding developers with a true background in best security practices or training your existing employees on security-first basics can be a heavy lift for some organizations, as it requires time and resources that sometimes are not available.

A **DevSecOps** approach implies that security is part of the overall software development life cycle (SDLC) but it's often not a bottleneck or a deterrent to code moving onto the next phase or even making its way to the production environment. Although this approach still requires a clear vision of security best practices, often the budget constraints, tooling, or delivery deadlines trump the introduction of security measures. And tying this back to Ullrich's point about legacy code, these are the areas that are often glossed over to get an update or feature-rich version of software out the door to keep pace with an organization's goals. If a developer is unfamiliar with extremely old code or development languages, they may address a smaller, but easier-to-tackle vulnerability that leaves legacy code wide open for a cunning attacker.

This discussion leads into some interesting topics that may be worth introducing in your own organizations including:

- Can your detection and response tools assist with patch management and/or can they provide additional safeguards against the improper or unauthorized use of AI in your organization?
- How do your current policies deal with the use of AI in day-to-day practices and/or in security tools, automation, and detection?
- What are some effective measures for identifying and defending against third-party risks

SOC Automation

With more threat actors, more volume, and more devices in our enterprises than ever before, it is not shocking that our SOCs are overworked, understaffed, and constantly looking for ways to automate processes. Effective SOC automation was at the forefront of the conference this year. How can we arm employees with tools that prevent burnout, reduce false positives and identify and mitigate against emerging threats as they continue to evolve? To pick the most powerful solution, you must first traverse the myriad terms and tools that aim to provide detection and response.

Understanding the Technology: Which Detection and Response Tool Is Right for My Organization?

Endpoint Detection and Response (EDR)

Most of us are familiar with the traditional EDR tools and how they operate. But each new year brings new solutions aimed to ease burnout and identify incidents with far better accuracy. And to do it well before they actually reach your endpoints: the ultimate goal of your detection and response platforms. Traditional EDRs require that each endpoint (laptop, server, virtual machine, or mobile device) is being monitored in realtime by a locally installed "agent." They sit and monitor these endpoints as they are the attackers' point of entry into your organization. Although many organizations no longer use EDRs as a standalone utility, they are the backbone for many of the emerging platforms used to detect threats.

Network Detection and Response (NDR)

NDR tools use a physical or virtual appliance or sensor to detect network threats and anomalies. Threats could come in the form of unauthorized protocols, anomalous ingress or egress transmission sizes, atypical ports, or other unusual behavior based on normal observed activity. This is ideal for networks that are unable to set up an agent on every endpoint due to the increasing "remote or work from home" landscape. NDRs can be used to trigger alerts, drop suspicious traffic, quarantine untrustworthy devices, and even create forensic images for a more detailed analysis if needed.

Extended Detection and Response (XDR)

An XDR is a platform designed to ingest data from your existing EDR, NDR, and other products and make them visible and controllable from one standard interface. For large organizations, having one unified platform designed to integrate and provide visibility across all of their existing products can be a game-changer. Although this may seem like the best approach, keep in mind that sometimes this technology limits you to specific vendors. Also, it is only as effective as your underlying EDR and/or NDR.

Managed Detection and Response (MDR)

MDR solutions are a bit different than the above, as they are not a product but a software-as-a-service (SaaS) approach that many organizations rely upon for real-time monitoring, threat-hunting, and remediation when they don't have the knowledge, personnel, or resources to do all of these tasks and more in-house. Although not a replacement for an EDR security solution—as many MDR providers will leverage the existing EDR as part of their process—they provide the SOC-level expertise and threat-hunting skills many smaller organizations lack. MDR providers rely upon their vast knowledge and skill set in this arena and can provide the type of 24/7 support that only large corporations could effectively manage. Depending on your company's size and the depth and breadth of the IT/security department, MDRs are often a cost-effective approach for securing your environment and preventing incidents before they occur due to their incorporation of active threat hunting. Human-enacted and automated remediation efforts include threat isolation and vulnerability fixes. This can prove to be valuable for organizations that have adopted the DevSecOps approach described above and those who may be operating with older, exploitable source code in their environments they have not yet remediated through their SDLC.

Steve Sims
SANS Offensive Cyber Operations Curriculum Lead and Fellow

Automation and AI as a Force Multiplier for Offensive Operations

One of the final panel topics this year was presented by SANS Institute fellow and the current curriculum lead for Offensive Operations, Steve Sims. His session focused on the use of LLMs, like *Shell GPT*, to hyper-accelerate the exploitation life cycle.

He referenced some newer terms that are getting a lot of buzz lately, like adversarial AI, which means using AI to assist with attacking existing LLM algorithms. This disrupts those original AI-based solutions put into play by the defenders, allowing for faster, stealthier, and more targeted attacks that circumvent traditional security mechanisms than those we have seen in the past. Although the terminology may be new, using machines to identify vulnerabilities in hardware and software and attack other machines based on their findings is not a new concept. In his presentation, he referenced a 2016 DARPA challenge that asked researchers and universities to do just that. The challenge was to identify vulnerabilities and automate patching and/or weaponization of the vulnerabilities based on the findings. Although this "challenge" was exciting for its time, it still required skilled humans vetting and understanding the nuances of the unique bugs that were uncovered.

But the biggest threat in this scenario, as Sims pointed out, was the fact that "the speed at which we can now discover vulnerabilities and weaponize them is extremely fast and it's getting faster." If enterprising attackers can locate a vulnerability and weaponize it faster than the patch cycle allows, they've met their objective. So, for defenders, this is a real threat that will prove to be difficult to defend against because the efficacy of adversarial AI will only continue to improve. See Figure 4. If you were thinking that all of this automation will put you out of your job, Sims offered a bit of reassurance stating, "We're still needed—meaning researchers—someone still needs to be there to groom the vulnerability and the exploit and get the whole thing to work."

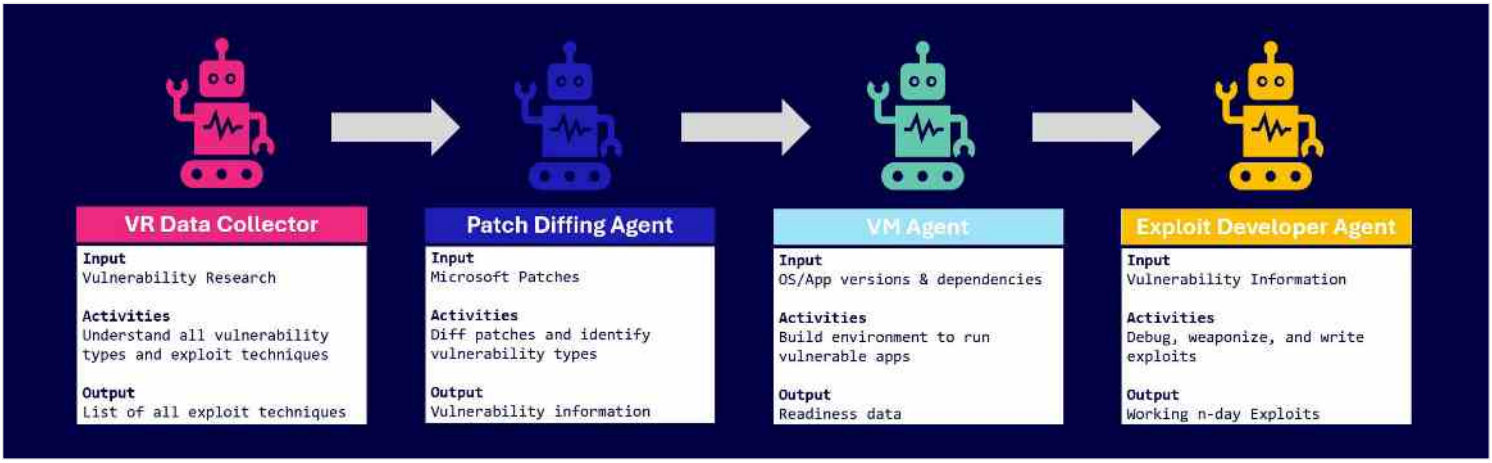

Figure 4. AI-Expidited Attck Cycle

This presents some real-world scenarios, though, that we need to consider, like the concept of an AI arms race. Will defensive AI be able to defeat the number of AI attacks now flooding the environment thanks to the advances in machine learning? Right now, it seems like adversarial AI has the upper hand. What's alarming is that attackers know organizations are unable to defend against an AI-based attack due to lack of proper tooling and skill sets required to do the job, and they will seek to exploit these weaknesses to their benefit. It has catapulted the rise in cybercriminal groups as well as nation-states that are leveraging AI to poke holes in our current network defenses, most worryingly those around critical infrastructure such as power grids. And this problem won't be slowing down anytime soon. While US-based companies struggle to find skilled employees with extensive knowledge in AI and ML, nation-states and criminal organizations are actively recruiting folks with the same skills to do work for the dark side.

In their latest year end Global Threat Intelligence Report, BlackBerry reported a soaring *70% uptick in unique malware files discovered, which equates to a staggering 5.2 novel malware samples per second*.[3] This illustrates the previous points that our current defensive controls are not able to withstand the onslaught of new AI-assisted attacks being orchestrated. Of the attack samples analyzed, threat actors were observed targeting the finance and healthcare industries most heavily during this cycle.

[3] "Global Threat Intelligence Report," BlackBerry, June 2024, www.blackberry.com/us/en/solutions/threat-intelligence/threat-report

Organizational Risks to Generative AI

After so much discussion on how AI is being used by attackers to infiltrate our existing defenses, the logical solution is to also use AI to make those defenses stronger. Many organizations are already focused on how to better leverage the advancements in this technology to not only beef up security but to augment and automate processes in their organizations to provide better speed, efficiency, cost savings, and other business and customer benefits. The advantages of AI are numerous and can help the decision-making process, realize reduction in human errors, create higher quality output, increase innovation, foster service and product improvements, and provide the type of speed and efficiency that can really positively impact productivity and profitability.[4]

But innovative technology is never without risk, and where AI is concerned, the risks run the gamut from subjects like accuracy and accountability, intellectual property, and even legal risks.[5]

Understanding how AI is used in your environment helps to proactively identify some of the potential risks involved. Most employees are well-versed in the practice of protecting sensitive data, but what they may not realize, is that utilizing AI to help solve problems often opens them up to potential data leaks. Fortunately, some of the same products and services aimed at monitoring and detecting incoming threats to your network can also be used to detect the inappropriate or unauthorized use of AI technologies, which results in possible sensitive data leaving your organization.

Although AI is not new, most organizations are merely on the threshold of realizing all the ways they can take advantage of the technology's benefits without introducing unnecessary risk to their organizations. At the very least, this will likely be the basis for many important and sometimes polarizing debates in the years to come, so organizations should begin preparing now if they haven't already. This was foreshadowed at this year's event by the increased focus on policy and compliance.

[4] "12 key benefits of AI for business," TechTarget, June 2023, www.techtarget.com/searchenterpriseai/feature/6-key-benefits-of-AI-for-business

[5] "The Top Five Real Risks Of AI to Your Business," Forbes, June 2023, www.forbes.com/sites/rscottraynovich/2023/06/22/the-top-five-real-risks-of-ai-to-your-business

Policy and Compliance

It is more important than ever to make sure your policies align with your business goals and that security is discussed early on to create an effective plan before a crisis hits. There were a plethora of vendors showcasing all-encompassing auditing, automation, and reporting platforms to make sure key stakeholders have visibility at each level.

Third-Party Risks and Mitigation Efforts

Continuing with threats related to software are those specifically attributed to relationships with third parties. Your company's products and services can be directly affected by the entities you partner with and one misstep in their security practices could, in turn, shutter your business' reputation or livelihood. Ultimately, those vulnerabilities are passed onto your organization, which makes solutions like SecDevOps and MDRs a viable option for businesses seeking to lower their risk. In addition to cybersecurity risks, third parties carry financial-, operational-, and even compliance-based risks, which should all be assessed prior to entering a relationship. Fortunately, there are frameworks and services available to walk organizations through the intricacies of assessing these risks and how to best leverage third-party products and services effectively.

IT Risk and Compliance Challenges

When it comes to compliance risks, it's often a topic that proves difficult for a lot of organizations. Depending on the type of business you conduct and the information you process, you are typically subject to different regulatory, legal, and industry standards, and evaluating your compliance is often done through different auditing measures. Organizations that don't adhere to the set of rules set forth for their industry can face fines, legal issues, suspensions, and even forced closures, so it's certainly an area that shouldn't be overlooked. Some of the most common risks that are unearthed during these assessments include out-of-date or improperly coded applications and lack of proper authentication mechanisms for accessing data.

Zero Trust in Verifying Identities in the Age of AI

Another important and upcoming threat that was highlighted on this year's panel was an AI-related topic that is becoming a bigger concern across the board: zero trust in verifying identities in the age of AI. As Ullrich pointed out, thanks to advances in ML, bots are more than 15% more likely than humans to correctly pass captchas, and they can do it faster to boot.[6]

[6] "An Empirical Study & Evaluation of Modern CAPTCHAs," July 2023, https://arxiv.org/pdf/2307.12108

But just how you identify someone as a real-life human in today's digital environment is proving more difficult and he sees it as a two-tiered problem: 1) How do you establish an identity to an online entity—the part of the process that takes the most research and investment cost upfront? and 2) How do you do it in such a way that it isn't so intrusive that you alienate your customer base?

One dilemma lies in our need for identity verification to access sensitive information, but another presents itself when users push back when they feel that our chosen verification methodologies overstep what people perceive to be acceptable. One such failed attempt was exhibited by the IRS. Their process left people unhappy because it required too much information to prove customers were, in fact, human and the owners of their own information. So, then, what constitutes a normal and acceptable verification method and how do companies implement mechanisms that don't ostracize their customer and/or employee base?

Ullrich believes this problem may be alleviated with good information security practices that strive to establish mutual trust between the user and the services upfront. As more entities get it right, they can leverage that trust and apply it to additional products and services.

Deepfakes and Digital Trust: Developing Detection Mechanisms

All of these verification methodologies don't necessarily address the emerging threat that comes with the rise in the easy availability of technology to almost perfectly imitate images, videos, or audio of real people. There also has been a huge reduction in cost of those technologies, which has caused a surge in synthetic media known as *deepfakes*. So, how do companies holding your information confirm you are a real person with a legitimate purpose attempting to authenticate yourself, and how do you truly know to whom you're talking online?

The term deepfakes typically has a negative connotation when the techology is used to impersonate humans for profit, propaganda, or other nefarious purposes. But just like most innovations in the technological realm, the ability to "synthesize" or digitally recreate real people has been around for many years, and there are many legitimate businesses harnessing that innovation for good, including digital artists, gaming companies, and film studios. There is no shortage of applications designed to digitally touch-up, stylize, or even totally re-create pictures, video, and audio, and today, the results are so lifelike, they are being incorporated into many services to efficiently assist with repeatable or scriptable tasks.

One such company, Synthesia, says they have eliminated almost all the glitches that made previous AI-generated avatars "creepy," and are now able to trick even the most discerning viewers into believing these avatars are real likenesses. They credit swift advancements in generative AI and a myriad of companies feeding massive datasets of human speech, actions, expressions, emotions, and gestures from social media, YouTube, and the internet at large into their AI models.

This makes threat detection at scale much more difficult. Currently, there are free and open-source tools that claim to be able to spot deepfakes, but there are even more resources designed for creating them. This is one area that is sure to set vendor products and services apart from their competitors in the next few years.

RSA's mission is to bring together cybersecurity professionals to discuss current and future concerns and to share ideas and solutions that will create safer environments. SANS shares those goals, and for the 15th year, took the stage to showcase what they perceived to be the top cyber threats facing many organizations. AI was the recurring thread woven into many of the guided discussions, vendor solutions, expert panels, and networking events this year. Although the risks may seem daunting, we are thankful to have a place for agencies, institutions, and businesses large and small to come together, share their obstacles and successes, and find solutions for better securing their information in a landscape that continues to change and challenge us each year.

Best Practices and Key Takeaways

It's never too late to implement better security practices. This could come in the form of:

- Implementing safer coding practices and utilizing better patch management strategies
- Identifying your organization's weak areas and arming yourself with better automation to detect and mitigate risks before they lead to a breach
- Embracing AI but putting the necessary safeguards in place to ensure you don't open your organization up to unnecessary data exposure
- Educating employees
- Continually evaluating relationships with third parties
- Developing trust and open lines of communication with your customer-base so they aren't overwhelmed by increased protections imposed by your organization
- Adapting to change

With innovation comes the need to evaluate our current practices, look for ways to improve our processes, and share our success and failures with the community.

Conclusion

RSA's mission is to bring together cybersecurity professionals to discuss current and future concerns and to share ideas and solutions that will create safer environments. SANS shares those goals, and for the 15th year, took the stage to showcase what they perceived to be the top cyber threats facing many organizations. Supply chain attacks continue to increase, placing even greater importance on the relationships and trust you build with third-party entities. The rise in unique malware and zero-day attacks targeting organizations is also growing, attributable to the surge in use of AI by attackers to identify weaknesses and exploit them at lightning speeds. Although the risks may seem daunting, we are thankful to have a place for agencies, institutions, and businesses large and small to come together, share their obstacles and successes, and find solutions for better securing their information in a landscape that continues to change and challenge us each year.

Sponsor

SANS would like to thank this paper's sponsor: