ARCTIC WOLF LABS

Principle networks

ARCTIC WOLF LABS

2025 PREDICTIONS

Introduction

Adaptability and Opportunism in 2025

Looking ahead to 2025, we believe two forces will exert a significant influence on tomorrow's threat landscape: adversaries' ability and willingness to adapt, and their propensity for financial gain. Unencumbered by laws, certain ethical standards, or institutional inertia, attackers of all types — from nation-state agencies, to ransomware groups, to hacktivists, to individuals — benefit from an opportunistic advantage and nimbleness of execution that defenders may lack. Closely related, threat actors perceive the world through a lens of opportunity. With this outlook, practically any new technology, or emerging crisis, or change in IT architectures — to list just a few circumstances — can be leveraged.

Against this backdrop, our specific observations lead to our five core predictions for 2025:

01 **Many organizations will see a continued breakdown of their perimeter defense as threat actors target VPN gateways and other edge devices.**

02 **Malicious actors will continue to refine their social engineering methods, creating opportunities for large-scale campaigns.**

03 **Ransomware attacks will increasingly exploit weaknesses in identity and access management (IAM) configurations.**

04 **Critical infrastructure will continue to be targeted, both for extortion and to prepare the digital battlefield for potential hybrid conflicts.**

05 **The widespread availability of advanced AI reasoning will allow threat actors to rapidly uncover novel initial access techniques.**

These predictions are the work of several of our brightest minds who aim to prepare security teams for the challenges of the year ahead to mitigate the risks posed by threat actor activity. It is also important to note that these listed predictions highlight areas of concern but are not presented in a ranked or hierarchal format. We suggest determining the priority of each topic based on the specifics of your environment.

Dan Schiappa
Chief Product & Services Officer

About Arctic Wolf® Labs

Arctic Wolf Labs is a group of elite security researchers, data scientists, and security development engineers who explore security topics to deliver cutting-edge threat research on new and emerging adversaries, develop and refine advanced threat detection models with artificial intelligence, including machine learning, and drive continuous improvement in the speed, scale, and detection efficacy of Arctic Wolf's solution offerings.

With their deep domain knowledge, Arctic Wolf Labs brings world-class security innovations to not only Arctic Wolf's customer base, but the security community-at-large.

01 Organizations Will See a Continued Breakdown of Their Perimeter Defense as Threat Actors Target VPN Gateways and Other Edge Devices

Once the cornerstone of a strong cybersecurity posture, the perimeter has undergone unprecedented transformation in recent years — resulting in both a larger attack surface and a plethora of remote access tools that threat actors are eager to exploit.

Defense in depth – with multiple layers contributing unique capabilities and adding redundancy – has always been a stalwart of enterprise security.

The Only Constant Is Change

For both technical and psychological reasons, perimeter defenses would likely top past surveys about the relative importance of different layers. However, the past few years have brought considerable change as multiple trends converged.

First, enterprise perimeters themselves have transformed. From IoT devices and cloud infrastructure to customer-facing applications and many other interface points, today's perimeters bear little resemblance to those of even the recent past — and one result of this transformation is a vastly expanded attack surface.

Second, the adoption of remote work has imposed new requirements and introduced new challenges that influence perimeter architectures and tooling. Even as return-to-office (RTO) initiatives draw some workers back, extended and partially remote workforces remain common. To enable and accommodate this new operational reality, organizations introduced and continue to refine their tooling, including VPN gateways and zero trust network access (ZTNA) tools.

Third, threat actors have adjusted their tactics, techniques, and procedures (TTPs) to — as is always the case — take advantage of these shifts.

More edge devices and services mean more configurations to harden and vulnerabilities to patch, with any mistake or delay creating an opportunity for an attacker to gain initial access. Of course, more edge devices and services also mean a higher exposure to zero-day vulnerabilities.

Turning the Tables

Threat actors have also had considerable success attacking the very tools organizations rely upon for enabling remote work.

For instance, the rapid introduction of VPN gateways, and the rearchitecting that often accompanied their rollout, led to misconfigurations that adversaries are quick to discover.

But more significantly, even solid implementations can give a false sense of security. For example, 2024 saw widespread exploitation of software vulnerabilities within VPN gateways (e.g., Palo Alto Networks, SonicWall), necessitating a coordinated response between security service providers, the appliance vendors, and end customers to remediate. We must be cautious to not overlook that all networked devices, including security appliances, run software and all software has the potential for flaws or misconfiguration.

Additionally, adversaries frequently leverage stolen credentials to turn an otherwise secure gateway into a convenient means of surreptitiously accessing the enterprise environment.

In organizations that permit working from home (or elsewhere), a large percentage of the workforce may have remote access — increasing the chances that a set of compromised credentials can be used to access protected IT environment.

When threat actors successfully penetrate VPN devices, there can be a very quick turnaround from initial access to acting on objectives — as seen with Akira affiliates taking less than two hours to exploit a SonicWall vulnerability and detonate ransomware.

Yet VPN gateways are not the only devices at risk. Other on-prem perimeter devices such as IP cameras, VoIP devices, industrial IoT gateways, medical devices, and others have the same potential to be leveraged by malicious actors to conduct attacks. This highlights the need for a detailed inventory of your environment and continuous monitoring to identify such threats.

Manufacturers Beware...

Although this threat is not limited to any specific industry, we have observed a significant impact in manufacturing.

For example, a recent ransomware threat event involving SonicWall SSL VPN access affected a wide variety of industries — but manufacturing represented 32% of all intrusions investigated. Similar disproportionality was observed in 2024 Arctic Wolf Incident Response case data, with the manufacturing industry representing 44% of all cases investigated.

Every Organization Is at Risk

Because every connected, digital organization has a perimeter — even as the exact composition of edge devices and services, and remote connectivity gateways varies — these attacks aren't limited to any specific industry.

Moreover, they're not necessarily limited to any particular attack type — although ransomware deployment, cryptocurrency mining, industrial espionage, and data theft are common attacker objectives.

That said, comparatively smaller organizations may have allocated fewer resources into network segmentation and other safeguards, which can make them seem like easier prey. In 2024 we witnessed an increase in small- to medium-sized organizations being exploited in opportunistic ransomware attacks (e.g., Fog and Akira affiliates).

Recommendations

- Ensure network and endpoint logs are available for examination and correlation — visibility into initial exploitation can be limited, making such telemetry especially important for identifying potential intrusions as quickly as possible.
- Train the workforce on credential hygiene best practices and consider subscribing to a **threat intelligence** service that includes monitoring of credential dumps.
- Reduce the blast radius of compromised accounts and devices by using **network segmentation** to help prevent unauthorized users from accessing specific network-connected resources and to create micro-perimeters around critical assets and network components.
- Implement a **vulnerability management program** that prioritizes continuous vulnerability remediation and assessment, with other components of the program complementing and assisting overall remediation and mitigation.
- Subscribe to **security bulletins** and incorporate their recommendations into your regular security operations.
- Pay special attention to **securing Internet of Things (IoT) devices** — these devices often come with little or no baked-in security, and can often be quickly forgotten once installed in your environment.

02 Adversaries Will Continue To Create and Leverage Opportunities for Large-Scale Social Engineering Campaigns, While Incorporating New TTPs

Social engineering offers a cheap and effective way for threat actors to bypass technological defenses, and new tools — particularly generative AI — make it even easier to execute even more effective attacks.

Even as organizations invest heavily in technological defenses, one extremely fallible element remains: the humans who make up the employees, contractors, vendors, and other third parties comprising the modern extended workforce.

For a threat actor launching a multi-phase attack, it can be easier and more efficient to use social engineering to bypass defenses than to employ technological means. After all, it only takes one mistake from one person to unwittingly open the door.

> **“Never let a good crisis go to waste.”**
> *- Winston Churchill*

We regularly see threat actors immediately spring into action when disaster strikes, hoping to exploit the desperation and chaos that follows. The past year provided attackers with two prime opportunities:

- CrowdStrike’s infamous update-gone-wrong, which led to widespread IT outages
- CDK’s global downtime following a ransomware attack

In both cases, we observed phishing campaigns targeting impacted parties — often dangling service restoration as a lure.

Unfortunately, there’s no reason to think attackers will abandon such effective approaches.

If something big happens in 2025, it’s a safe bet that social engineering campaigns will be crafted and launched within hours.

We may even see two-stage attacks in which a threat actor first disrupts a major player, and then immediately launches campaigns to take advantage of the resulting chaos.

Don’t Believe Everything You See and Hear

Compounding the threat, generative AI is lowering the bar to entry for crafting convincing messaging and creating deepfakes that increase the effectiveness of phishing attacks.

Voice phishing (or vishing), in particular, **is growing as a threat**, with adversaries masquerading as employees and targeting call centers, help desks, and other departments that interact remotely and can grant access (e.g., via password recovery/reset flows). Plus, even live video feeds can be manipulated to make attackers look and sound like legitimate employees. Today’s deepfake tools require only a few still photographs — which are easily sourced from LinkedIn or a team member’s public social media presence.

Without a reliable mechanism for remote identity verification — security questions don’t count, but offline hardware keys do — anyone providing remote assistance will remain an attractive target.

Looking beyond phishing, in 2024 we witnessed **a historic campaign of high sophistication targeting the XZ Utils project**. This episode was likely conducted by a nation state-affiliated group, and is an example of social engineering playing out over a longer term.

MFA Implementation Details Matter

In response to heightened awareness of phishing, many organizations have deployed multi-factor authentication (MFA). However, it's important to note that MFA — while immensely valuable — isn't a panacea.

Threat actors are executing adversary-in-the-middle (AiTM) attacks and employing **MFA fatigue** against organizations with MFA enabled. As an example, Arctic Wolf observed widespread exploitation of the **Axios phishing campaign** which leveraged AiTM techniques.

Correct configurations should render these types of attacks impossible, but reaching that state often requires specialized expertise and modern MFA techniques like those employing the **WebAuthn**/**FIDO2** phishing-resistant standards.

Recommendations

- Educate users about phishing and conduct phishing attack simulation — create a culture of **security awareness** that forgoes assigning blame to encourage individual accountability.
- Implement email controls as a defensive layer: restrict external email inbound, add headers to emails to warn users when external inbound emails arrive, and use email security products from vendors like Mimecast or Barracuda.
- Implement and enforce modern, phishing-resistant MFA.
- Reduce the blast radius of compromised accounts and devices through **network segmentation** and **least-privilege access controls**.
- In addition to network (for phishing lures) and endpoint (for post-compromise activity) telemetry, ensure SaaS logs (e.g., Microsoft 365, Microsoft Entra, Okta, Duo) are available for examination and correlation.
- Where possible and practical, consider limiting your exposure to outages and other disruptions by employing a multi-vendor strategy.
- Incorporate third-party outages into your disaster recovery planning.
- As part of the due diligence process when evaluating potential third-party vendors and service providers, pay particular attention to their continuity and recovery plans, and compliance certifications.

03 Ransomware and Other Attacks Will Increasingly Exploit Weaknesses in Identity and Access Management (IAM) Configurations

Identity has rapidly risen to prominence as one of the most important and complicated cybersecurity domains — unfortunately, misconfigurations and permissive policies play right into the hands of ransomware affiliates and other threat actors.

Identity and access management (IAM) systems are essential elements of the modern enterprise technology stack. Among other functions, IAM infrastructure provides:

- Authentication, to establish with confidence that entities — primarily employees, systems, and devices —are who they say they are
- Authorization, to enable the appropriate level of access to privileges, resources, applications (etc.)
- Identity management, to enable users, administrators, and systems to make updates and changes to identity data and related information

Of Details and Perceived Trade-Offs...

However, identity is a specialized and challenging domain, and even the use of an IAM system is no guarantee against errors. The results of these errors — including overprivileged access, orphaned accounts, and shadow directories — can be exploited to gain unauthorized access to systems and resources.

Plus, security objectives often run up against productivity needs. Safeguards are intended to prevent unauthorized access, but too much friction for users can impede their ability to do their jobs. In many cases, such perceived trade-offs force IT to soften security measures or maintain workarounds — even when the risks of doing so are known.

As if the situation wasn't complicated enough, Active Directory (AD) configurations are not specifically designed to be "secure by default." This could lead to multiple opportunities for detrimental misconfigurations that require updating or modifying. Unfortunately, making changes to authentication infrastructure can be disruptive to end users, causing needed upgrades/migrations to fall behind. For example, on-premises Active Directory infrastructure often lags behind other high priority work and remains operational longer than it should, even with versions of Windows that have reached end-of-life (EoL) status.

Infostealers and Credential Abuse

Naturally, threat actors are all too willing and able to take advantage of any weakness, vulnerability, or misconfiguration in identity infrastructure — including in MFA implementations (as noted previously). In particular, **the use of infostealers** to acquire credentials or active session cookies and the subsequent reuse of those credentials and cookies are major threats.

To put the risk in perspective, **Verizon's 2024 Data Breach Investigations Report (DBIR)** indicates that over 80% of breaches involve compromised identity. In practice, this can mean **gaining initial access** or performing intrusion actions like reconnaissance, privilege escalation, and establishing persistence.

We often observe that VPN credentials serve as the root point of compromise in our Arctic Wolf Incident Response cases. In 2024 we saw affiliates with Akira and Fog ransomware opportunistically exploiting SSL VPN accounts without MFA enabled.

While the wider identity industry is taking steps to help strengthen and standardize defenses — including the establishment in September 2024 of the **Interoperability Profiling for Secure Identity in the Enterprise (IPSIE) Working Group** — we expect threat actors to have success exploiting IAM configurations well into the foreseeable future.

Why Break in When You Have the Keys?

With good reason, many security leaders are especially concerned about zero-day exploits. However, while potentially devastating and definitely headline-grabbing, such attacks represent a small fraction of overall incidents.

In contrast, the use of stolen credentials and general exploitation of poor credential hygiene are much more common — despite contributing to intrusions that can be considerably more difficult to detect.

Technologists are primed to pay particular attention to sophisticated attack vectors, but threat actors are all too willing to take advantage of misdirected vigilance by employing more mundane TTPs.

Recommendations

- Work with your IAM and application providers to strengthen your defenses against account takeovers (ATOs) and — more generally — to enforce strong credential controls and phishing-resistant MFA.
- Block authentication attempts from hosting-based traffic — this may be extended to include proxies and anonymization services, which criminals also use to hide their origins.
- Set automated blocking on authentication attempts to hinder password-spraying activities, and implement geolocation-based blocking (e.g., restricted countries, impossible travel scenarios).
- Configure syslog to forward your organization's VPN and firewall logs to your security operations provider.
- Implement **network segmentation** to limit the ability of threat actors to move laterally.
- Ensure telemetry from on-premises (e.g., Active Directory) and SaaS authentication providers (e.g., Microsoft 365, Microsoft Entra, Okta, Duo, etc.), and endpoints (for post-compromise activity) is available for examination and correlation.

04 Critical Infrastructure Will Continue To Be Targeted, Both for Financial Gain and in Preparation for Potential Hybrid Conflicts

Continuing a troubling trend, key sectors will be subjected to disruptive attacks and stealthy intrusions — as adversaries look for financial payouts and aim to prepare the digital battlefield for potential conflict.

Critical Infrastructure Is Under Attack

Issued in February 2013 by the Obama administration, **Presidential Policy Directive 21 (PPD-21)** designates 16 sectors as "critical infrastructure."

As the Cybersecurity & Infrastructure Security Agency (CISA) **explains**, these are sectors "whose assets, systems, and networks, whether physical or virtual, are considered so vital to the United States that their incapacitation or destruction would have a debilitating effect on security, national economic security, national public health or safety, or any combination thereof."

Although companies and other organizations within these sectors always face cyber threats, 2024 saw an increase in threat activity against a subset. For example:

- Among a string of incidents in the Water and Wastewater Systems sector, the largest wastewater company in the US **was attacked in October**
- As reported by Check Point, the Energy sector endured **a 70% year-over-year increase** in attacks against utilities
- Sophos noted **a four-year high** in ransomware attacks against Healthcare and Public Health organizations, with surveys suggesting two-thirds of such organizations were impacted
- In February, **a joint statement issued by CISA, the FBI, and the NSA** warned that the People's Republic of China was seeking to pre-position itself on U.S. communications networks in case of a conflict with the U.S.

Many of the attacks followed the typical ransomware playbook of disrupting operations and exfiltrating data to extort a payment. However, there's suspicion among the cybersecurity community that some of these incidents may have also been intended to distract from a strategic objective of establishing stealthy persistence within these environments.

From Disruption to Destruction

The coming year could represent a geopolitical tipping point. Adversaries often attempt to take advantage of change, and a number of Western nations have recently experienced leadership transitions, while elections are scheduled or anticipated in several others.

At the same time, there's no shortage of ongoing and potential regional conflicts — any one of which could rapidly emerge as a flashpoint triggering a larger-scale war.

Should such a situation arise, nation-state affiliated actors who've been refining their skills for economic benefit may soon be called upon to escalate from disruption to outright destruction. Taking a water treatment plant offline for a few hours is one thing, but disabling a plant for days or weeks during a hot/kinetic war will force a reallocation of attention and resources.

In fact, we don't even need to speculate about what a modern hybrid war may look like, as Russia's ongoing invasion of Ukraine has seen countless attacks against **energy assets**, **water systems**, and other critical infrastructure.

The Best-Case Scenario: More of the Same

Unfortunately, aside from a possible regression to the mean following a 'busy' year, there's no reason to think that attacks against critical infrastructure will decrease.

In the best-case scenario, threat actors will continue to attack for financial gain, to perform reconnaissance, and to establish footholds that can be leveraged as needed in the future.

In a worst-case scenario, facilities throughout the Western world will endure attacks unlike any experienced before.

Recommendations

- Implement a vulnerability management program that prioritizes continuous vulnerability remediation and assessment, with other components of the program complementing and assisting overall remediation and mitigation.
- Subscribe to security bulletins and incorporate their recommendations into your regular security operations.
- Ensure network and endpoint logs are available for examination and correlation — visibility into initial exploitation can be limited, making such telemetry especially important for identifying potential intrusions as quickly as possible.
- Reduce the blast radius of compromised accounts and devices through network segmentation and least-privilege access controls.
- Ensure a comprehensive, realistic, and up-to-date disaster recovery plan is in place for your organization. Arctic Wolf customers may discuss this with their concierge security team for a better understanding of best practices for their environment.
- Maintain proper backup practices — while backups don't address the issues around data exfiltration, being able to restore operations can buy your organization time and limit the ripple effects of the attack.
- Understand and account for the shared responsibility model of cloud services — the cloud/SaaS provider and the SaaS customer (i.e., you) each assume ownership of particular responsibilities with respect to data security.
- Follow the 3-2-1 principle of backup — 3 copies of data (1 primary and 2 backup), 2 copies stored (at separate locations), 1 off-site storage (ideally in a secure private cloud).
- Test your recovery processes and capabilities — a real-world incident is not the time to discover that your backups don't work or that they are incomplete.

05 The Widespread Availability of Advanced AI Reasoning Will Allow Threat Actors to Rapidly Uncover Novel Initial Access Techniques

Thus far, even the most advanced AI models have failed to replicate human reasoning capabilities, but that may soon change — and once it does, threat actors will undoubtedly harness this newfound power to uncover new ways to break into protected environments.

Imagine a near-term future where capabilities wielded by today's most experienced penetration testers are embedded within advanced AI models. For good or bad — or, more likely, for good and bad — this is where things are heading, as advanced reasoning capabilities are further refined and become a widely available core element of AI.

While frontier large language models (LLMs) already possess some decent programming capabilities — albeit with key limitations — they haven't yet led to a significant increase in new initial access techniques emerging from these technologies.

However, the bleeding edge of AI is advancing at a breakneck pace. In just a few years we've seen extraordinary — and arguably unparalleled — progress in text, audio, image, and video generation capabilities.

In that time, one lingering question has remained: when will AI lead to a step-function change in how threat actors execute attacks?

Lowering Barriers to Entry

So far, existing AI technology has opened up wider access to programming and malware authoring in general, and has helped to craft phishing lures and automate campaign workflows, but limitations in AI's reasoning capabilities have prevented meaningful advancements in TTPs.

More specifically, logical reasoning is a key part of the ability to write code that functions as expected, especially for complex codebases — and today's LLMs have known gaps in their ability to reason.

However, based on the trajectory of AI development, it's safe to presume that such reasoning may only be one or two iterations away.

The State of the Art

Already, recent developments in LLMs have brought about considerable improvements in this space. For example, while the capabilities of OpenAI's o1-preview model are still being studied, preliminary results show improvements in benchmarks measuring capabilities in math, physics, chemistry, and formal logic.

Once LLMs are able to competently reason about the flow of data through an application, they are expected to facilitate discovery of novel vulnerabilities or to chain together vulnerabilities in a manner that is more difficult for humans to achieve.

In fact, we've already seen the development of an open-source tool that uses Anthropic's Claude AI model to find zero-day vulnerabilities in Python codebases — even completing the entire call chain from user input to server output.

An Ongoing Arms Race

Very soon, penetration testers and hackers may regularly employ advanced LLMs in their efforts. Of course, this is a double-edged sword: while it may allow programmers to write more secure code and organizations to bolster their own defenses, it can just as easily be employed for malicious purposes.

The warning signs are already here: in 2024, **OpenAI identified a cluster of ChatGPT accounts using the platform for scripting and vulnerability research tasks**. OpenAI identified three threat actors — all with nation-state ties — performing these actions.

Almost certainly, this is just a glimpse into how state-aligned threat actors are exploring the use of LLMs in their vulnerability research programs.

Trickle-Down TTPs

For its part, **Microsoft has introduced a set of TTPs to describe LLM-based attacks**. Within that list, "LLM-assisted vulnerability research" stands to become a more common practice as reasoning capabilities improve further.

Nevertheless, it stands to reason that nation state-affiliated adversaries are in the best position (i.e., due to their ample resources) to operate at the bleeding edge. Of course, new techniques don't stay secret for long, and successful approaches will quickly trickle down to the wider cybercrime ecosystem.

Practically, this means that the shorter-term risks of novel LLM-enabled initial access techniques are likely concentrated among targets considered high-value to nation states — particularly those with intellectual property in key domains, or whose disruption would assist in a large-scale cyber conflict (e.g., critical infrastructure).

However, as these TTPs become more common and affordable, they will be incorporated into more attacker toolsets and will be used with very little discretion.

Recommendations

- Implement **network segmentation** and the **principle of least privilege** to help limit an adversary's ability to perform intrusion actions, should they successfully achieve initial access.
- As part of your **vulnerability management program**, conduct audits and penetration tests to identify areas of weakness and low visibility, and identify and prioritize remediation of significant vulnerabilities.
- Employ **managed detection and response (MDR)** to provide continuous monitoring capable of identifying common post-compromise activities.

Conclusion

Anticipating the Future Helps Us Prepare for It

It's interesting to note that four of our five predictions center on the ways in which attackers seek to gain access to protected environments. Such a result was not planned in advance, but is instead a reflection of what drives our team. By examining the evolution of attacker TTPs — especially those focused on initial access — we're able to propose reasonable answers to the question, "What will adversaries try next?"

These answers, in turn, help to inform preventative measures that organizations can prioritize based upon their own operating context and risk appetite.

For what it's worth, many of these measures are about getting the details right, like:

- Create a full inventory of systems
- Diligently apply patches, and prioritizing perimeter-facing devices
- Check and recheck configurations
- Practice sound credential hygiene
- Monitor for stolen credentials
- Implementing a security awareness program to minimize end user risks

These things aren't flashy, but they're fundamental — and they can make the difference between withstanding an attack and becoming a victim.

However, no defense is impenetrable, and plenty of today's threat actors are motivated, patient, and well-resourced.

Consequently, just as important as preventative measures are the abilities to detect and respond — quickly and effectively — to cyber attacks that do occur.

Detection capabilities largely boil down to two things:

01 **Collecting as much telemetry from as much of your IT environment and ancillary systems as is possible**

02 **Making sense of it in real time**

Easier said than done, of course — but entirely feasible with the right tools.

Response is as much about people and processes as it is about technology, because an effective and timely response depends upon people doing the right thing while under immense pressure.

Accordingly, perhaps the most important elements of response are preparation and practice. Overconfidence is a killer, and you'd much rather find a flaw in your disaster recovery process — whether procedural or technical — during a controlled exercise than in the chaos of an intrusion.

No one knows exactly what the future will hold, but getting the basics right will put your organization in a position to withstand and recover from whatever adversaries throw your way.