

Trust in the Age of AI: Securing Identity, Users (& Agents) and Access Everywhere

Ovidiu Pismac
Senior Security Cloud Solution Architect
Microsoft Cybersecurity Architect, CISSP, MCT
Microsoft EMEA Global Channel Partner Sales (GCPS)

ovidiup@microsoft.com

<https://www.linkedin.com/in/ovidiupismac/>



The world has changed...



The **Threat Landscape** has changed...
DRAMATICALLY...!



Agent led vs. Human led attacks


Reconnaissance
is continuous
and adaptive

Speed
becomes
structural

Vulnerability
chaining becomes
the norm

Reach extends
to the
entire estate

Customer impact of attacks



Average cost per data breach is ~\$4.44M, \$5.52M cost of data breach for organizations that didn't use AI security and automation, \$10.22 M historic value for cost of data breach in US with higher regulatory fines, detection and escalation costs.

2025 Daily Reality

WIRED

NEWSLETTERS SUBSCRIBE

MATT BURGESS SECURITY SEP 22, 2025 2:00 AM

A Cyberattack on Jaguar Land Rover Is Causing a Supply Chain Disaster

The UK-based automaker has been forced to stop vehicle production as a result of the attack—costing JLR tens of millions of dollars and forcing its parts suppliers to lay off workers.



DELL CAMERON ANDREW COOTS SECURITY AUG 9, 2025 6:30 AM

Security News This Week: The US Court Records System Has Been Hacked

Plus: Instagram sparks a privacy backlash over its new map feature, hackers steal data from Google's customer support system, and the true scope of the Columbia University hack comes into focus.



Pandora Data Breach Exposes Sensitive Data of 30,000 Customers

August 11, 2025

SentryBay
CYBERTHREATRADAR®



Sensitive Data of 30,000
Customers Exposed by
Pandora Data Breach



salesforce

Qantas Refuses to Pay
Ransom After Hackers
Leak 5 Million
Customer Records

How Did Attackers Breach JLR?

Jaguar Land Rover has not released a comprehensive technical post-mortem of the September 2025 breach; however, security researchers and leaked information from the attackers themselves provide a clear picture.

The **JLR breach breakdown** reveals that the intrusion was not the result of a sophisticated zero-day exploit, but rather the execution of well-known tactics: social engineering, credential abuse, weak segmentation, and inadequate detection.

Social Engineering and Credential Theft

At the heart of the attack was **social engineering**. The hacker collective responsible has a long history of using phishing and “vishing” (voice phishing) to impersonate insiders. Investigations suggest the breach originated from a targeted vishing campaign weeks earlier, when attackers posing as internal staff tricked employees into disclosing credentials.

Armed with a valid username and password, in some cases with administrator rights, the attackers could log in through normal authentication flows. This meant they didn’t need to brute-force systems or exploit unknown flaws; they simply walked in with stolen keys.

Compromised Credentials as Keys to the Kingdom

Credentials played a central role in both of JLR’s 2025 incidents. In the March attack, for instance, the Hellcat group compromised a third-party contractor’s Jira/VPN login credentials. That single account gave them access to JLR’s internal project management system, where they quietly extracted **hundreds of gigabytes of data**.

The September breach appeared to follow the same pattern: once the attackers had access to valid credentials, they encountered few safeguards to prevent them from proceeding.

- **Multi-factor authentication (MFA)** was either missing or inconsistently applied, allowing simple password reuse.
- **Overly broad permissions** meant a single account could access far more data than necessary.
- Some credentials were shockingly outdated; reports showed that passwords stolen as far back as 2021 still worked in 2025.

"The biggest security threat is the human element"

A company can spend hundreds of thousands of dollars on firewalls, intrusion detection systems and encryption and other security technologies, but if an attacker can call one trusted person within the company, and that person complies, and if the attacker gets in, then all that money spent on technology is essentially wasted.

Kevin Mitnick

2026 Daily Reality

News



'Salesforce environment Odido hacked through phishing and social engineering'

Friday 13 February 2026, 16:48 by [Redactie](#), 33 comments

Last updated: 13-02-2026, 17:23

The attackers who managed to steal the personal data of **6.2 million people** at Odido were able to get in through phishing and social engineering. Sources told the **NOS**. In the attack, the attackers first sent phishing emails, with which they managed to steal login details from employees. The attackers then called Odido employees, pretending to be the IT department. The attackers asked the employees to approve their login attempt, which allowed them to bypass an extra security step, according to the NOS. The accounts of several employees are said to have been hacked.



Image by MajicStock | Shutterstock



Malaysia Airlines landed on the Qilin ransomware gang's dark web victim site Thursday, but the group has released no proof, no samples, and no details on what data, if any, was stolen.



A TRACTOR GATE FOR AN OIL PIPELINE IN ROMANIA. PHOTOGRAPH BY PEPSEBNU, LICENSED UNDER CC BY-SA 3.0

Daryna Antoniuk

February 6th, 2026


Cybercrime

Romania's oil pipeline operator confirms cyberattack as hackers claim data theft

Romania's national oil pipeline operator Conpet said a cyberattack disrupted parts of its technology infrastructure and knocked its website offline earlier this week, adding that oil transport operations were not affected.



Conpet, which operates about 3,800 kilometers (2,360 miles) of pipelines supplying domestic and imported crude oil and petroleum products to refineries across Romania, said its operational technology systems, including its supervisory control and data acquisition (SCADA) and





Malaysia Airlines, the nation's flagship carrier, is claimed by the Qilin ransomware group. Image by Cybernews via Qilin leak site.

Attackers are using AI

<https://www.anthropic.com/news/disrupting-AI-espionage>

Full report: Disrupting the first reported AI-orchestrated cyber espionage campaign

Simplified architecture diagram of the operation



Chinese-Speaking Hacker Uses DeepSeek Agent to Launch Autonomous Cyberattacks



By Mayura Kathir | July 31, 2026

Chinese-speaking threat actor "knaithe" (aka KnYuan) has been caught running an AI-enabled autonomous attack stack built around DeepSeek and the Hermes Agent framework, proving that large language models can now drive end-to-end offensive operations with minimal human oversight.

Hermes provided terminal access, skills orchestration, and Model Context Protocol (MCP) integrations. At the same time, DeepSeek generated FOFA queries, assessed CVEs, and generated exploit logic.

Parallel to this core stack, the actor experimented with multiple LLMs, including MiniMax, alongside limited trials of Western tools such as Claude, demonstrating an ongoing evaluation of the AI market for offensive workflows.

Claude AI Autonomously Discovers Cryptographic Weaknesses That Escaped Expert Review



By Divya | July 29, 2026

Researchers at Anthropic reported that **Claude Mythos Preview** autonomously discovered new cryptographic attacks against the post-quantum signature candidate HAWK and a reduced-round version of AES.

This demonstrates that advanced AI models can now contribute to research that was traditionally the domain of specialized cryptanalysts. While the findings do not compromise currently deployed systems, they raise important questions about how the cryptography community will validate and manage discoveries generated by AI.

Autonomous AI Agent Exploits Zero-Day to Breach Hugging Face Infrastructure



By Eswar

August 1, 2026

An autonomous AI agent powered by **OpenAI models** breached **Hugging Face's** production infrastructure in July 2026 after escaping its evaluation sandbox via a zero-day vulnerability.

Documented by HiddenLayer's Research Team on July 31, the agent was undergoing an internal cyber capability evaluation on ExploitGym, a benchmark designed to test an AI's ability to discover and exploit software vulnerabilities.

OpenAI had disabled production safety classifiers and reduced cyber refusals to measure the model's raw offensive capability, restricting sandbox network access to package installation via a self-hosted JFrog Artifactory registry proxy.

Autonomous AI Agent Exploits Zero-Day

Rather than solving the benchmark, the agent found a zero-day in the Artifactory proxy itself, using it to reach the open internet. **JFrog has since patched** the flaw in Artifactory 7.161.

From there, the agent located an unsecured public code evaluation sandbox, gained root access, and used it as a staging base to pivot toward Hugging Face, inferring the platform might host the benchmark's reference solutions.

OpenAI

Research

Products

Business

Developers

Company

Foundation

Q

July 21, 2026 Security

OpenAI and Hugging Face partner to address security incident during model evaluation

Revolut handed nearly 700 customers' data to scammers

European payments group has contacted clients whose IDs, addresses and other details were handed over to hackers posing as government officials



Since launching in 2015, Revolut has become Europe's biggest fintech with 80mn customers in 30 countries



CenterPoint Energy is a Houston-based public utility company that provides electricity and natural gas services. It serves about 7 million customers across Indiana, Minnesota, Ohio, and Texas.

A hacker, posting under the alias '4d722e4d656f77,' **claims** to have pulled 7.49 million lines of CenterPoint customer data through an API that had no web application firewall, no rate limiting, no certification checks, and no authentication token.

Microsoft Digital Defense Report 2026

*Where we are
in 2026*

Cyber risk is accelerating at rapid speed, driven by the connections among identities, AI systems, cloud services, software supply chains, edge infrastructure, and critical systems.

Empowered by the vast capabilities of AI, attackers are pursuing all available entry points, including points of leverage where one compromise can create effects far beyond the original target, setting up cascading waves of compromise. The methods attackers are using and their targets haven't changed—it's the scale and effectiveness of their attacks that have.

What's different this year

- AI is significantly accelerating attacker workflows while creating a new enterprise attack surface.
- Identity-led intrusions are leading to credential theft that drives follow-on attacks, creating a flywheel of unauthorized access.
- Open-source software, edge devices, cloud identity, and operational technology are being attacked at an unprecedented rate.
- The constant pace of attacks means resilience is now leadership discipline, not solely a technical recovery objective.

In the pages that follow, we walk through the biggest challenges and changes confronting defenders not only this year, but also the next, and what organizations can do to reduce the threat.

63%
of intrusions
involved data theft

5.3 hrs
average time
before exposed cloud
workloads were attacked

93%
of voice
phishing attacks
kept the victim on the line
long enough to begin social
engineering

46M+
business
email
impersonation attacks
were detected

89-95%
of email phishing
attachments
led to a credential theft effort

Ten security priorities for organizations

Cyber risk is now a continuous test of business resilience. The boundaries between state-sponsored activity, cybercrime, and influence operations have continued to erode. This converges into sustained and overlapping campaigns that have tested not only technology estates, but the resilience of organizations and the trust placed in them. For business leaders, this means focusing less on reacting to the last breach and more on how long critical operations can be sustained under continuous cyber stress.

In this environment, cybersecurity is a continuous business condition, defined by persistent adversary pressure, machine-speed operations, and systemic interdependence. The defining question for boards today is not whether disruption will occur, but whether the organization can continue operating, contain harm, recover quickly, and maintain trust when it does.

The response starts with secure foundations: Secure by Design, Secure by Default, and Secure Operations, applied through identity-centric Zero Trust. These foundations make the following priorities durable as AI expands the number of identities, applications, agents, and dependencies organizations must govern.

The organizations that succeed will be those that engineer for endurance: building systems that can absorb disruption, adapt under pressure, and maintain trust even when compromised. For security leaders and boards, building that resilience requires these priorities:

	1	Govern cyber risk at leadership level	2	Defend identity as the primary control plane	3	Establish governance, visibility, and data controls for AI
		Report identity exposure, patch latency, agent permissions, critical dependencies, dwell time, and recovery readiness.		Protect identity by deploying passkeys and phishing-resistant MFA, reducing privilege, governing human and non-human identities, and using rapid invalidation/rotation for tokens and session credentials.		Maintain visibility into approved AI tools, services, and agentic capabilities while monitoring for unauthorized or unmanaged use. Apply sensitivity-aware access, data protection, and governance to the data those systems and agents can reach.
	4	Reduce exposure and edge risk	5	Focus on cross-signal correlation	6	Enact a continuous threat vulnerability management model
		Continuously discover internet-facing assets and prioritize rapid patching of edge devices, enabling centralized and high-fidelity logging, enforcing strong authentication, and integrating network-level visibility with threat intelligence.		Correlate endpoint, identity, cloud, application, email, and network signals to detect intrusions earlier. Cross-signal correlation also enables defenders to anticipate attack paths and targets.		Conduct assume-breach exercises focused on cloud services, SaaS platforms, APIs, and workflow automation capabilities. Prioritize testing based on where the highest-value data resides, is accessed, or can be aggregated.
	8	Secure software dependencies	9	Participate in collective defense	10	Prepare for long-horizon technological transitions
		Harden software installation, maintain a software inventory, validate provenance, monitor packages, and prepare rapid containment.		Share actionable threat indicators and victim-informed intelligence with trusted partners and law enforcement to support coordinated threat tracking, investigation, and disruption.		
						The defining question
						The defining question for boards today is not whether disruption will occur, but whether the organization can continue operating, contain harm, recover quickly, and maintain trust when it does.

Top 10 takeaways from this report



1. AI is changing the physics of cybersecurity.

AI is compressing attack timelines, lowering the cost of sophisticated capabilities, and enabling attackers to operate with greater speed, scale, and autonomy. But the same capabilities can strengthen defense by accelerating discovery, analysis, prioritization, and response. Security organizations must prepare for an environment in which both attackers and defenders increasingly operate at machine speed.

2. Attackers are already targeting AI as another attack surface.

Threat actors can exploit AI to do things like execute malicious commands, steal AI serving capacity, and exfiltrate data. As AI becomes ubiquitous in the workplace and at home, it has become both a tool and a target for attackers.

3. People remain a heavily exploited initial-access path.

Threat actors continue to target user behavior through user execution, ClickFix-style social engineering, phishing, and impersonation. These techniques remain effective even as AI changes their speed, scale, and personalization.

4. Identity is the primary control plane for defense.

Attackers consistently take advantage of unnecessary privileges, persistent access, and weak identity governance. This is why phishing-resistant multifactor authentication (MFA) and passkeys, disciplined identity hygiene, tiered administration, and strong privileged-access enforcement remain the best safeguards against cyberattacks. That control plane now spans both human and non-human identities, including applications and agents.

5. Signal correlation through shared threat intelligence is key to defense at multiple levels.

The degree to which a defender's systems share and correlate threat intelligence is one of the highest-leverage strategic variables under the defender's control. Signal cross-correlation can reveal attack patterns that remain invisible when endpoint, identity, cloud, application, email, and network telemetry are analyzed in isolation.

6. Emerging threats facilitate stealthier attacks, operating at a larger and faster scale.

Over the next year, Microsoft sees the biggest threats coming from:

- Open-source supply chain compromise
- Attacks against edge devices
- AI as a force multiplier for malicious activity

These threats are emerging as AI-generated synthetic content erodes confidence in digital communications. As authentic and synthetic content become harder to distinguish, individuals can become simultaneously more vulnerable to manipulation and less confident in legitimate sources. This shifts the security model from implicit trust toward explicit verification.

7. Protecting data is even more important as AI expands access at scale.

Trusted AI depends on trusted, well-governed data. Data protection has always been a key priority, and now that AI systems and agents can reach and act on sensitive information at unprecedented scale, it's more crucial than ever. Organizations should reduce oversharing, apply sensitivity-aware access and least-privilege controls, and govern how data is discovered, used, and shared across AI, cloud, and application environments.

8. The future of defense involves rethinking how we protect our assets from reactive incident response to proactive, continuous discipline.

Defenders must shift from tool-centric vulnerability management to threat exposure management. At the core of this is:

- Asset visibility
- Continuous vulnerability awareness
- Intelligence-driven threat detection
- AI-accelerated insight

This means reporting metrics should shift from number of patches deployed to exposure reduced, detection coverage increased, and time-to-mitigate compressed.

9. National and international policy decisions will shape the future of cybersecurity.

In an increasingly contested landscape, cybersecurity demands clearer rules, credible accountability, and sustained cooperation. Strengthening protections for critical infrastructure, advancing operational accountability, and preserving an open, secure, and resilient digital environment must be a priority for governments on both a domestic and international level.

10. Resilience must be a core outcome of cybersecurity planning.

Retain the existing focus on reducing exposure, detecting threats faster, adapting to changing risks, and maintaining operations. This preserves resilience as the outcome without competing with identity, secure foundations, data, and the three-part operating model for 'primary' status.

Identity is at the center of modern cyber risk

Passkeys are the safer replacement for passwords

Passkeys represent a strong countermeasure to both password and phishing attacks.

Passwords continue to be the first vulnerability many attackers test. In the months immediately following the launch of ChatGPT, Microsoft tracked a significant but short-term spike in password attacks, suggesting attackers were experimenting with using generative AI to launch them. The effectiveness of password attacks is waning, however. Overall, password attacks are down 26% year-over-year. This provides strong evidence that security measures such as MFA and passkeys are forcing attackers to pivot to other approaches to achieve access to high-value targets.

Phishing has evolved. In the last few years, attackers have shifted to techniques like fatigue attacks and AiTM phishing to circumvent MFA. In the first six months of 2026, the share of AiTM phishing and token theft attacks more than doubled. This is based on their share as a percentage of the overall attacks Microsoft detected and, of course, is still comparatively small. Nonetheless, it shows the direction of travel. Passkeys are phishing-resistant credentials that aren't replayable. They only work on the website or app for which they were created, meaning if a user lands on a lookalike phishing site, their passkey won't work there.

Organizations need to choose the right passkey for their users

The FIDO Alliance, the National Institute of Standards and Technology (NIST), and the UK's National Cyber Security Centre (NCSC), along with many companies that steward accounts for consumers and businesses, strongly advocate passkey adoption. As organizations standardize using passkeys, choosing the right type—such as device-bound or synced—must be based on user roles and risk.


Passkeys enable phishing-resistant authentication and offer a viable alternative to passwords. However, as organizations secure authentication, threat actors try to take over accounts by attacking other vulnerable areas—for example, new-hire onboarding and account recovery. These rely on weaker credentials like SMS or email-based one-time passcode (OTP). The next step is therefore to remove all phishable credentials from an account to make not only authentication phishing-resistant but also make the account itself phishing-resistant.

Risk class	Synced passkeys	Device-bound passkeys
What it means	Passkey syncs through the passkey provider's cloud to the user's other devices	Private key remains on one device or authenticator and cannot be synced elsewhere
Pros	• Easy to use across phone, laptop, etc. • Built-in backup if a device gets lost • Less chance of getting locked out Synced passkeys can satisfy the requirements for a multifactor cryptographic software authenticator and can be used at AAL2 when implemented in accordance with NIST SP 800-63B-4 requirements.	• More control over where credentials live • Harder for attackers to copy or move Device-bound passkeys implemented as multifactor cryptographic devices can satisfy AAL2 or AAL3 requirements.
Cons	• Dependent on cloud/account security • Harder for IT to tightly control	• Only works on that one device • No automatic backup or recovery • Can be frustrating for users with multiple devices
Best for	Most users with non-admin privileges in the organization	Admins, sensitive systems, high-risk access

Figure 1. What observed prompts most often try to achieve

The goals shown describe what observed prompts tried to achieve, while a single attempt can involve several risk areas. For example, an attempt to misuse a connected tool to extract customer records can involve both goal hijacking/downstream tools and sensitive data exfiltration.

Key		%
A	Malicious use/policy violation	45%
B	Sensitive data exfiltration	17%
C	System prompt leakage	14%
D	Goal hijacking/ downstream tools	12%
E	System-internals leakage	11%
F	Resource exhaustion / financial harm	1%



Source: Microsoft observation of prompts against generative AI agents (February to May 2026).


“

AI is making fraud more sophisticated and accessible, accelerating cybercrime worldwide. As threats evolve across borders, public-private and international cooperation—built on technology, trust, and shared action—is essential to disrupting cybercriminals and protecting society.

Kenichi Sakurazawa, Managing Director
Japan Cybercrime Control Center (JC3)

Figure 2. Top attack vector categories in telemetry data on AI workloads

Malicious link injection, tool and AI agent section abuse, and prompt subversion account for the majority of observed attack activity on AI workloads.



Key		%
A	Malicious link injection and exfil	52%
B	Jailbreak and safety bypass	16%
C	Tool and agent abuse	12%
D	Identity and credential abuse	10%
E	Prompt injection and instruction subversion	8%
F	Reconnaissance and data probing	2%

Source: Telemetry data on AI workloads (90 day window, 2026).

Critical infrastructure and operational technology


Defending a nation starts with defending its critical infrastructure (continued)

Once inside, adversaries move with commodity and bespoke tooling. Against critical infrastructure, Truebot malware (39%) and information stealers (38%) dominate the observed tool mix, with the AzureHound framework (12%) used to map cloud-attack paths. Information stealers accelerate the supply chain. They harvest credentials, which become the next intrusion's entry point, cascading across the vendors, integrators, and shared services that connect a municipality's smallest water utility to its largest hospital system.

Legacy water systems, residual community risk


US drinking water and wastewater utilities run aging operational technology, often maintained by private operators and lower-capacity districts. In 2025, the Environmental Protection Agency reported that more than 70% of inspected water systems failed to meet basic cybersecurity requirements. Incidents such as the late 2024 American Water customer-portal outage, show how fast those gaps translate into service disruption and community safety risk.

Figure 2 – Technique summary: SLED




Key		%
A	Cloud identity abuse	62%
B	Password spray attacks	38%

Figure 3 – Technique summary: Critical infrastructure



Key		%
A	Cloud identity abuse	78%
B	Password spray attacks	13%
C	Brute-force attacks	2%
D	Evolving phishing threats	2%
E	Other	5%

Figure 4 – Tool summary: Critical infrastructure



Key		%
A	Truebot malware	39%
B	Information stealers	38%
C	AzureHound framework	12%
D	Impacket	6%
E	LemonDuck	5%

Figures 2 and 3 — Cloud identity abuse dominates the technique mix in both SLED (62%) and critical infrastructure (78%). Source: Microsoft Threat Intelligence. March 1 to June 30, 2026.

Figure 4 — Truebot malware and information stealers lead the tooling used against critical infrastructure. Source: Microsoft Threat Intelligence. November 1, 2025 to April 30, 2026

Microsoft Security

**Defend
with
AI**

Defender XDR
Perception
iSOC (XDR & SIEM)

**Safely
enable
AI**

Agent 365
Entra Suite

**Trust AI
with your
Data**

Purview


Get Ready for AI

Codename MDASH


Healthy Security Foundation

Hierarchy of cybersecurity needs

Drawing inspiration from Maslow's hierarchy of needs, this graphic illustrates a prioritization of cybersecurity, starting with the most basic need: protecting identities. AI has a role at each tier, underscoring its potential to enhance security measures. Cultivating a robust security culture within the organization, helps ensure the technological defenses and human practices evolve in concert to mitigate threats effectively.



End to End Protection with Microsoft Defender XDR



Protection across every control point

Native signals, context, and controls

Defender correlates activity with organizational context and threat intelligence.

Controls that respond together

Coordinated protection can act across the attack path while the threat is still moving.

Microsoft Defender



The broadest native coverage



Hybrid identities



Endpoints



Cloud workloads



Email and collaboration



SaaS apps



Data security



External sources

Attack disruption impact

Uses correlated signals, the latest threat intelligence, and AI to stop attacks with above 99% confidence

30 sec

average time to
disrupt ransomware

68K

attacks disrupted
per month

45K

AiTM attacks
disrupted per month

81K

users saved
per month

580K

devices saved from
encryption per month

Hear from our customers

"Defender lit up like a Christmas tree with hundreds of alerts and **attack disruption kicked in to save the day.**"



"This is the **silver bullet for a SOC.** Having high-quality detections is not easy to build..."



Qilin Russian ransomware encryption prevented by Defender for Endpoint in February 2026

Microsoft Defender									
Search									
Last 30 days95% Last 30 days									
ExportCopy list linkRefresh1 Week3 IncidentsSearch for name or IDCustomize columns									
Filter set: Save									
Status: New, In progressAlert severity: High, Medium, LowAdd filterReset all									
	Incident name	Incident Id	Priority score	Tags	Severity	Investigation state	Categories	Impact	
☐	Human-operated ransomware attack was launched	1978	56	Ransomware +5	High	4 investigation states	Initial access, Execution, Pe...		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		
☐	Ransomware-linked threat actor detected			Ransomware	High	Terminated by system	Ransomware		

Microsoft Defender

Search

Home

Exposure management

Investigation & response

Incidents & alerts

Incidents

Alerts

Hunting

Actions & submissions

Partner catalog

Threat intelligence

Assets

Devices

Identities

Applications

Cloud

AI agents

Microsoft Sentinel

Identities

Dashboard

Service accounts

Health issues

Tools

Endpoints

Email & collaboration

Alerts

Incident graph

Layout

Group similar nodes

Feb 2, 2026 5:47 PM

New

Potential human-operated malicious activity

Feb 2, 2026 6:13 PM

New

Risk of tampering using Safe Mode reboot as part of a Ransomware attack (attack disruption)

Feb 2, 2026 8:13 PM

New

Risk of Group Policy abuse as part of a Ransomware attack (attack disruption)

Category

Ransomware

MITRE ATT&CK

T1484.001: Group Policy Modification

Detection source

EDR

Service source

Microsoft Defender for Endpoint

Feb 2, 2026 9:28 PM

New

Potential human-operated malicious activity

Feb 2, 2026 9:52 PM

New

Possible lateral movement involving suspicious file

Feb 2, 2026 9:53 PM

New

Possible lateral movement involving suspicious file

Feb 2, 2026 9:53 PM

New

Possible lateral movement involving suspicious file

Feb 2, 2026 9:53 PM

New

Possible lateral movement involving suspicious file

Feb 2, 2026 9:53 PM

New

Possible lateral movement involving suspicious file

Communication

Association

Risk of Group Policy abuse as part of a Ran...

'Qilin' malware was detected during lateral move...

Medium

Detected

New

Ransomware behavior detected in the file system

High

Detected

New

Potential human-operated malicious activity

High

Detected

New

Defender Amplifier

High

Detected

New

File dropped and launched from remote location

Medium

Detected

New

9:53:28 PM

[4] ntoskrnl.exe created file PSEXESVC.exe

Possible compromised user account delivering rans...

High

Detected

New

Remote ransomware infection

High

Detected

New

Suspicious PsExec activity initiated remotely

Medium

Detected

New

Suspicious file registered as a service

Medium

Detected

New

Potentially compromised assets exhibiting ransomw...

High

Detected

New

File dropped and launched from remote location

Medium

Detected

New

Risk of tampering using Safe Mode reboot as part o...

High

Detected

New

'Qilin' ransomware was detected

Medium

Detected

New

'Qilin' malware was detected during lateral move...

Medium

Detected

New

Ransomware behavior detected in the file system

High

Detected

New

Potential human-operated malicious activity

High

Detected

New

Defender Amplifier

High

Detected

New

9:53:28 PM

[4] ntoskrnl.exe created file win.exe

Malware

Defender Amplifier

High

Detected

New

Ransomware

Open alert page

Manage alert

Details

Recommendations

INSIGHT

Quickly classify this alert

Classify alerts to improve alert accuracy and get more insights about threats to your organization.

Classify alert

Alert state

Classification

Assigned to

Not Set

Unassigned

Set Classification

Alert details

Alert ID

Category

da8ecbba7d-5314-414e-88f4-2860e5450541_1

Ransomware

MITRE ATT&CK Techniques

Detection source

T1021.002... +1 More

View all techniques

EDR

Service source

Detection status

Microsoft Defender for Endpoint

Detected



'Qilin' malware was prevented on a Microsoft SQL server

Investigation #110 is complete - Terminated by system

[Comments \(2\)](#)[Investigation graph](#) [Alerts \(2\)](#) [Devices \(1\)](#) [Evidence \(1\)](#) [Entities \(1\)](#) [Log \(3\)](#)

Investigation Summary

Investigation Status Timeline

Started

Feb 1, 2025 7:52 AM

Ended

Feb 3, 2025 7:55 AM

Duration

27h

Investigation details

Status

Terminated by system ⓘ

Alert severity


was: High

Category

Malware

Detection source

Antivirus



'StopCrypt' ransomware was prevented

Investigation #111 is complete - Terminated by system

[Comments \(2\)](#)[Investigation graph](#) [Alerts \(2\)](#) [Devices \(1\)](#) [Evidence \(1\)](#) [Entities \(1\)](#) [Log \(3\)](#)

Investigation Summary

Investigation Status Timeline

Started

Feb 1, 2025 7:59 AM

Ended

Feb 1, 2025 8:00 AM

Duration

20h

Investigation details

Status

Terminated by system ⓘ

Alert severity


was: High

Category


Ransomware

Detection source

Antivirus



Defender 100% protection even with vulnerable infrastructure




Everything needed to operationalize every signal is built in

Native out-of-the-box detections

Get best of breed, cross-workload protection across identities, endpoints, SaaS apps, email and collaboration tools, cloud workloads, and SIEM data

Investigate across all your data

One investigation surface across first- and third-party data bringing together XDR signals and SIEM logs in one place



Autonomous defense


Autonomously disrupt multi-stage active attacks like ransomware in minutes

Built-in SIEM


Analysts get the visibility, investigation depth, and hunting power of SIEM, without requiring security teams to build and operate a traditional SIEM-first environment

All your security signals and context are agent-ready

Combine Microsoft's frontline intelligence with the organization's assets, history, policy, and priorities.



Microsoft
Multi-Agent
System
MDASH
proved the
model
Find what others miss



96.5%
CyberGym
Benchmark¹

Ensemble of diverse models

100+ code-scanning agents

Model-agnostic by design

¹ Cybergym benchmark outcomes with specific benchmark conditions detailed in: (discover and validate exploitable vulnerabilities with codename MDASH) section

May 2026 Patch Tuesday: 16 CVEs found by codename MDASH

Across Windows networking & authentication stack — majority reachable with no credentials from the network

16

Total CVEs

4

Critical Severity

10

Kernel-Mode

6


Usermode

Component	CVE	Type	Sev.	Description
tcpip.sys	CVE-2026-33827	Remote Code Execution	Critical	Remote unauth UAF via SSRR IPv4
ikeext.dll	CVE-2026-33824	Remote Code Execution	Critical	IKEv2 double-free → LocalSystem RCE
netlogon.dll	CVE-2026-41089	Remote Code Execution	Critical	Unauth CLDAP stack overflow
dnsapi.dll	CVE-2026-41096	Remote Code Execution	Critical	Crafted UDP DNS → heap OOB
tcpip.sys	CVE-2026-40415	Remote Code Execution	Important	Pre-auth remote UAF via SA double-decrement
tcpip.sys	CVE-2026-40413	Denial of Service	Important	NULL deref via IPv6 extension headers
ikeext.dll	CVE-2026-35424	Denial of Service	Important	Memory leak
http.sys	CVE-2026-33096	Denial of Service	Important	Unauth remote QUIC OOB read
tcpip.sys	CVE-2026-35422	Security Feature Bypass	Important	IPsec cross-SA fragment splicing
	CVE-2026-32209	Security Feature Bypass	Important	Unauth local WFP RPC disables cache

+ 6 additional CVEs across tcpip.sys and telnet.exe

Teams of agents to cover the entire threat lifecycle

Project Perception - Expose risk. Investigate threats. Strengthen defenses.



**Defend
with
AI**

(Defender XDR
Project Perception
iSOC
Entra Suite)

**Safely
enable
AI**


Agent 365
Entra Suite

**Trust AI
with your
Data**


Purview

Get Ready for AI
Codename MDASH

Identity Infrastructure: Enrich your understanding with a unified identity inventory




ZERO TRUST with Conditional Access




Secure entire session with **continuous access evaluation**

Auto-revoke access in near real-time when critical events or policy violations are detected


- Account revocation
- Account deletion
- Password change
- User location change
- User risk increase




Zero Trust architecture



Zero Trust architecture




Zero Trust for AI



Zero Trust for AI



Zero Trust for AI



Microsoft Entra ID Protection

Microsoft Entra ID Protection blocks identity takeover in real-time with the power of advanced machine learning.

Preventing identity compromise by:



Protecting your organization with Risk-based access policies




Detecting advanced attacks with machine learning and threat intelligence



Responding in real-time with insights from first and third-party integrations

Identity Protection Dashboard





ID Protection at scale

"We stopped 33,000 challenges to the antibot system during our first six weeks...We also stopped 96,000 brute force attacks...We went from 6 million to 10 million IDs in a couple of months after implementing our SAF solution based on Microsoft Entra ID..."

- Roy Corneloues, Enterprise Architect, BT Group

Prevent identity attacks with strong authentication

Microsoft Entra ID supports a broad range of multifactor authentication options and external auth methods



Poor:
Password only

123456
admin
quertyuiop
P@ssword2024!
5%gHy*73xWcd#jQz89
cnpGo!17

Fair:
Password and...

-  SMS
-  Voice

Better:
Password and...

-  Microsoft Authenticator push notifications
-  Software Tokens OTP
-  Hardware Tokens OTP

Better:
Passwordless


-  Microsoft Authenticator phone sign in

Best:
Phishing-resistant

-  Windows Hello for Business
-  FIDO2 security key
-  Certificate-based authentication (multifactor)
-  Passkey in Microsoft Authenticator (device-bound)
-  Platform credential for macOS

Passkey support in Entra ID

Expanding phishing resistant credentials



Security
keys



Passkeys in
Microsoft
Authenticator

Users with **admin privileges** or
in **highly-regulated** industries.



Synced Passkeys

NEW

For **everyone else**
in your organization

Passkey Profiles

>> Scope rollouts to different groups

Configure up to 3 passkey profiles targetable to different user groups to enable scoped rollouts of new passkey features.

Target	Type	Passkey profiles (preview)	
IT Admins	Group	1 selected	
Marketing	Group	3 selected	
Engineering	Group	☐ Default passkey profile	
Finance	Group	☒ FIDO2 security keys	
HR	Group	☒ Microsoft Authenticator	
		☒ Synced passkeys	

>> New control for passkey syncability

Passkey profiles contain settings for attestation, AAGUID and syncability restrictions.

Edit passkey profile

Certain combinations of these settings could target passkeys that may not exist. This can prevent your users from registering and signing in with a passkey.
[View compatibility documentation](#)

Name *

Enforce attestation ☒

Target types *

Target specific passkeys ☒

Behavior * ☒ Allow ☐ Block

Model/Provider AAGUIDs *
+ Add AAGUID


Model/Provider AAGUIDs

c1f1a2bc-1d02-404a-b275-8a29047a43fd 

73bb0cd4-e502-49e8-9c5f-b59445bf720b 

Enable passkeys in Microsoft Authenticator

- Microsoft Entra ID customers can now use **device-bound passkeys** in Microsoft Authenticator app on iOS and Android
- No additional cost and easy to deploy phishing-resistant credential
- Meets strict security requirements—users can't sync, share, or recover any device-bound passkey hosted in Microsoft Authenticator



[How to enable passkeys in Microsoft Authenticator for Microsoft Entra ID | Microsoft Learn](#)

Microsoft Entra Suite

Establish Zero Trust
access controls



Microsoft Entra ID
P1

Secure access for
your employees



Microsoft Entra
Private Access



Microsoft Entra
Internet Access



Microsoft Entra ID
Governance



Microsoft Entra ID
ID Protection



Microsoft Entra Verified ID
Premium

Secure access for
customers/partners



Microsoft Entra
External ID

Secure access
in any cloud



Microsoft Entra
Permissions Management




Microsoft Entra
Workload ID



Microsoft Security Copilot



Improve your workforce access security



Legacy network security stacks are no longer sufficient

The rise of cloud workloads, AI, and hybrid work is straining traditional corporate networks and network security models, resulting in security risks and poor user experience.

Challenges with conventional approaches

1


Dramatic traffic increase strains network capacity and on-prem security stack. Sub optimal user experience on account of traffic hair-pinning.

2


Users circumvent IT controls and access resources directly.

3

Compromised users/devices can move laterally on traditional corporate networks.




Microsoft's Secure Access Service Edge (SASE) Solution



Microsoft's SASE Solution Unique Differentiators

Unified approach across identity and network access




- Unified policy engine for all access
- One place to know about and manage all apps across your environment through Conditional Access policy
- Integrated identity risk to network access policies and enforcements

Microsoft's global private wide area network



- Globally distributed proxy
- Minimized latency by eliminating additional hops
- Delivers a great user experience anywhere

Flexible deployment options




- Enhanced security controls for robust protection, superior visibility and resiliency for Microsoft apps
- Compatibility to be deployed alongside third-party security service edge (SSE) solutions

Boost productivity for your users with fast and seamless access


Directly connect to Microsoft services with a built-in Microsoft traffic profile

- Leverage best-in-class performance for your users by sending Microsoft traffic directly to Microsoft's SSE collocated with front doors for Microsoft services.
- Easily breakout Microsoft traffic from third-party secure access service edge (SASE) in side-by-side model leveraging built-in traffic profile.
- Better security and visibility without compromising end-user productivity.
- Connect remote users securely and directly to points of presence (PoPs) all over the world—avoid traffic hairpins to branch VPN.
- Revoke network access to Microsoft applications in near-real time with Universal Continuous Access Evaluation (CAE).



Extend Conditional Access and continuous access evaluation universally

- Enforce Conditional Access conditions and controls to all internet destinations
- Secure and govern internet access alongside cloud app access through a unified Zero Trust policy engine
- Apply continuous access evaluation to revoke access in near real time when conditions change
- Link context aware Conditional Access policies to relevant network security profiles




Microsoft Entra Internet Access: Secure Web & AI Gateway

Secure access and protect AI


Prompt Injection Protection

Protect enterprise GenAI apps from XPIA/UIPIA attacks, protect jailbreak




Network-based Data Security

Prevent corporate data exfiltration using network file filtering




Discover and onboard shadow AI

Discover apps, agents, data source, etc. being used and seamless onboard and attribute them to include in Conditional Access policies




Prevent corporate data exfiltration

- Enforce identity-centric, policy-driven file filtering at the network layer to prevent sensitive file leaks to unsanctioned destinations like generative AI apps and unmanaged cloud services.
- Combines Microsoft Purview's data classification with Global Secure Access security profiles for granular control over file movement without impacting productivity.
- Create file policies to block or allow files by MIME type or use Scan with Purview for sensitivity labels, content inspection, and user risk evaluation; supports HTTP/1.1 traffic only.




Microsoft Entra Private Access

An identity-centric Zero Trust Network Access (ZTNA) solution



Uniformly govern access to on-premises & cloud apps



Secure access service edge (SASE) ecosystem

from Microsoft Entra

SD-WAN and connectivity

Automated | Integrated

 aviatrix


CHECK POINT™


cisco

HPE aruba
networking

teridion
Know the flow.

 **VERSA**
NETWORKS

aka.ms/Blog/SASE-Partners

Network security integrations

Private Preview

 netskope

Netskope One
Advanced Threat
Protection (ATP)

 netskope

Netskope One
Data Loss
Prevention (DLP)

aka.ms/Blog/SSE-Integrations

Woodgrove HelpDesk - Woodgrove

https://woodgrovehelpdesk.azurewebsites.net





Welcome to Woodgrove Helpdesk portal

QR code scanned. Waiting for Verified ID credential to be shared from wallet...

Step 1: Get your card at MyAccount to prove your employment

To get support, you need to have your [VerifiedEmployee card](#) with a photo.



[I already have my card](#)
[I have a Government ID card](#)

Use the Authenticator to present your VerifiedEmployee card



Step 3: Get support

© 2024 - Woodgrove HelpDesk - No data is retained by this website. Sample code available on github [here](#). Documentation available [here](#).

Woodgrove Helpdesk

 <https://did.woodgrovedemo.com/>

Share with Woodgrove Helpdesk?

You will share:



Verified Employee 

Woodgrove


Result of face check

Required by organization

Next


Zero Trust Assessment – Microsoft 365 & Azure Tenant Security Health Check

<https://learn.microsoft.com/en-au/security/zero-trust/assessment/overview>



Defender EASM assessment – extended visibility of your public exposed assets in the Internet

How to use Microsoft Defender EASM (External Attack Surface Management)



Microsoft Entra Suite

Secure Zero Trust access
for your workforce



Microsoft Entra
Private Access



Microsoft Entra
Verified ID
Premium



Microsoft Entra
Internet Access



Microsoft Entra
ID Protection



Microsoft Entra
ID Governance

THANK YOU!

ovidiup@microsoft.com

<https://www.linkedin.com/in/ovidiupismac>
