



The 2026 State of MSP Threat Report

The evolving landscape of AI-powered threats, NHI, AiTM, BEC, M365, GWS, malware and a whole lot more.



Table of Contents

MSP Threat Overview & Key Findings	3
AI Attacks and Defenses: The New Battleground in 2026	5
Identity Threat Detection and Response (ITDR)	8
• Trends	
• AiTM	
• Cloud Productivity Tools at Risk	
Email Security	19
• Inbox Rule Manipulation	
• 2025 BEC Risk Profile	
• The Story of a \$1.5M BEC Scam	
• Post-Compromise	
Endpoint Security	22
• Threat Volume & Trend Analysis	
• Threat Campaign Analysis	
• Risk Assessment Matrix	
H2 2026 Threat Predictions	28
Connecting the Security Dots with Guardz	30

MSP Threat Overview

The threat landscape facing MSPs is undergoing a fundamental shift, from isolated incidents to continuous, identity-driven attacks operating at scale. Over the past 180 days, Guardz telemetry reveals a sharp increase in both attack volume and sophistication. Traditional entry points like malware are being replaced by stealthier techniques, including credential compromise, session hijacking, and the abuse of legitimate tools. Attackers are no longer relying on breaking in—they are logging in, persisting, and operating undetected across cloud environments.

At the same time, financial impact continues to rise, with business email compromise (BEC) incidents costing businesses millions of dollars, while ransomware activity surged dramatically across monitored environments. This evolving threat landscape highlights a critical reality: identity, access, and trusted tools have become the primary battlegrounds for MSP security.

2026 Key Findings By Attack Vector

AI-Powered Threats

- AI-generated phishing emails are now contextually accurate, grammatically flawless, and personalized to the individual target
- Automated Credential Stuffing with AI-Optimized Password Lists: **~31% of users have compromised passwords every month**
- There were **14,000+ unique spray IPs per month** targeting 10+ accounts each

Identity Security

- Session hijacking incidents increased by **~23% over a 180-day observation period**
- **89% of monitored SMBs** have at least one user with confirmed credential compromise at any point in time
- OAuth consent events soared by **45% between October 2025 and January 2026**

Email Security

- Inbox rules remain the **number one persistence mechanism** in BEC attacks
- Confirmed BEC incidents analyzed this year **ranged between \$140k to \$1.5M**
- The average requested wire transfer of a BEC campaign in early 2025 was **\$40,000** for broad-scale campaigns

Endpoint Security

- Ransomware behavioral detections surged **190% over a 50-day** observation period
- January 2026 saw the largest spike in monthly threat distribution (**+121.8%**) from monitored endpoint activity during analysis
- November 2025 had the largest volume of endpoint threats observed over a 180-day window (**2,911 incidents**)
- Remote Management Tool (RMM) Abuse comprised 26% of total endpoint-related threats

Foreword

The **2026 State of MSP Threat Report** provides a comprehensive analysis of key threats across Identity Threat Detection and Response (ITDR), email security, endpoint security, cloud, and AI-powered attack vectors observed by the Guardz research team from September 2025 through H1 2026. The report also includes an in-depth analysis of each vector, along with emerging threat predictions for H2 2026.

About the Research

This report is based on telemetry from the Guardz platform data lake, analyzing activity across SMB environments managed by MSPs globally. Observations span a 180-day window (September 2025 through February 2026), covering authentication, email, endpoint, and cloud productivity data from Microsoft 365 and Google Workspace tenants across three monitored regions: North America, EMEA, and APAC.

The Guardz research team applied dozens of purpose-built detection rules mapped to the MITRE ATT&CK framework, correlating signals to identify multi-stage attack patterns. Findings are drawn from billions of audit events, representing real-world attacks and defensive activity observed across the MSP-managed SMB population.



Dor Eisner
CEO & Co-Founder

“I’ve spent over two decades in security intelligence, and the pace of what we’re seeing now has no precedent. Attacks that required nation-state resources five years ago are being executed at scale by commodity toolkits today.”



Elli Shlomo
Head of Security Research

“What stands out in this data isn’t any single attack vector. It’s how identity, email, endpoint, and cloud signals are chaining together into multi-stage campaigns that move faster than traditional detection models were designed to handle.”



Doni Brass
SVP Product Strategy

“We’re building agentic security systems using the most advanced AI available. So are the attackers. The toolsets are converging, and the gap between defense and offense is narrowing faster than most people realize.”

AI Attacks and Defenses: The New Threat Surface in 2026

AI is the New Battleground for MSPs

The greatest threat to an MSP is a prompt. AI-generated phishing campaigns are contextually accurate, leverage sentiment analysis, and are dynamically personalized at unprecedented scale. This section analyzes the top AI-enabled attack techniques observed in 2026 and the defensive capabilities that counter them.

Top 3 AI-Powered Attack Techniques

We observed three separate AI-generated attacks and the attributed threat levels between September 2025 to February 2026.

1 Attack Technique: AI-Enhanced Phishing And Social Engineering

Threat Level: **CRITICAL** Active in the wild

Large language models (LLMs) have eliminated the traditional indicators of phishing: grammatical errors, awkward phrasing, and generic content. AI-generated phishing emails are now contextually accurate, grammatically precise, and personalized to the target recipient. Adversary-in-the-Middle (AiTM) attacks are at the forefront of this growing trend.

AiTM session hijacking events are increasingly driven by AI-generated phishing lures. Attackers are leveraging bulk SendAs delegated permissions from compromised accounts to distribute context-aware BEC emails that incorporate the victim's actual email history and communication patterns. Additionally, the clustering of Chrome OS user agents suggests the use of automated, infrastructure-level phishing campaigns built on standardized tooling and scalable attack frameworks.

2 Attack Technique: Automated Credential Stuffing With AI-Optimized Password Lists

Threat Level: **HIGH** Observed at scale

AI models trained on breached credential databases generate smarter password lists that account for per-user patterns: common mutations, corporate password policies, and regional language preferences.

Evidence from this dataset

~**31% of users have compromised passwords every month**. The persistent high rate suggests attackers are using increasingly effective password generation, +13% month-over-month growth in spray IPs, scaling infrastructure to match AI-generated attack lists, and **14,000+ unique spray IPs per month** targeting 10+ accounts each.

Defense

Continuous credential monitoring with leaked password detection, combined with passwordless authentication (passkeys, FIDO2) that eliminates the password attack surface entirely.

3 Attack Technique: Deepfake Voice And Video For BEC

Threat Level: MEDIUM-HIGH Emerging in SMB targeting

AI-generated voice clones of C-level executives are being used to authorize fraudulent wire transfers and override security controls via phone calls. Video deepfakes are also appearing on collaboration platforms like Microsoft Teams and Copilot. A notable example occurred in 2024, when a Hong Kong-based employee at Arup was deceived during a sophisticated deepfake video conference. Believing he was interacting with the company’s CFO, the employee unknowingly authorized a series of transactions totaling approximately \$25 million over a series of phone calls. Copilot access provides the contextual data (writing style, relationships, and terminology) needed to generate convincing deepfakes. With **8.7 million emails sent across the U.S.** alone, attackers have ample training data for voice and writing style cloning. BEC persistence mechanisms further highlight the operational infrastructure supporting AI-enhanced BEC campaigns.

AI-Powered Defenses: How Defenders Fight Back

Defense 1: AI-Powered Threat Detection and Triage

Agentic MDR/SOC platforms use multi-model AI architectures to triage, enrich, and investigate alerts at machine speed:

Capability	Human Analyst	Agentic MDR/SOC
Alert triage speed	15-30 minutes per alert	Seconds per alert
Consistency	Varies by analyst, shift, fatigue	Deterministic scoring framework
Context integration	Manual lookup across 5+ tools	Automated enrichment from all data sources
Scale	50-100 alerts/day per analyst	Thousands of alerts/day are continuously
Accuracy	67% (human baseline)	92.4% (AI-SOC benchmark, 172 incidents)


Defense 2: Behavioral Baselining and Anomaly Detection

AI establishes 30-day behavioral baselines for every identity, human and non-human, tracking authentication patterns, application access, email volume, file operations, and geographic behavior. Deviations from baseline trigger risk-scored alerts.

Baseline Dimension	Normal Pattern	Anomaly Trigger
Authentication location	2-3 cities, weekday hours	New country, weekend 3 AM
Application access	Outlook, Teams, SharePoint	Graph API, Azure PowerShell
Email volume	20-50 emails/day	500+ emails in 1 hour
File downloads	5-10 files/day	200+ files in 30 minutes
MFA patterns	Push approval, 1-2x/day	15+ push denials in 10 minutes

Defense 3: AI-Driven Investigation and Root Cause Analysis

When an alert fires, AI investigators automatically:



Defense 4: Continuous NHI Monitoring

AI platforms monitor non-human identities for:

- Credential usage from new IP ranges (data center → data center movement)
- Permission scope changes (least privilege violations)
- Authentication pattern deviations (time, frequency, target apps)
- Dormant SPN reactivation (orphaned accounts being exploited)

Traditional SIEM tools cannot effectively monitor NHIs because they lack the behavioral context to distinguish legitimate automation from compromise. They also cannot correlate identity intent across systems, track baseline service behavior over time, or interpret machine-to-machine interactions at the scale of modern cloud environments, where API keys and tokens continuously interact across distributed workloads.

Defense 5: Predictive Threat Intelligence

AI analyzes attack patterns across the entire customer fleet to predict emerging threats before they reach individual organizations.

Here is an example of how predictive threat intelligence might appear during a suspicious incident.

- **Cross-organization correlation:** If the same spray IP targets 100 orgs, alert all orgs proactively
- **Attack campaign clustering:** Group related incidents by TTP, infrastructure, and timing
- **Vulnerability prediction:** Identify organizations most likely to be targeted based on their configuration gaps

Traditional SIEMs and legacy tools lack this level of advanced threat intelligence.

When implemented collectively, these five defense mechanisms can greatly benefit SOC analysts and security teams in combating or containing AI-driven threats.

Identity Threat Detection and Response (ITDR)

Identity is the most attacked layer in SMBs. Guardz analyzed authentication and post-authentication activity across Microsoft 365 and Google Workspace, applying dozens of purpose-built detection rules spanning multiple MITRE ATT&CK tactics in a 180-day observation window (September 2025 - February 2026).

The data reveals persistent, escalating identity-based attacks against SMBs and the critical role of layered detection in stopping them, including a seemingly invisible attack surface, better referred to as Non-human identities (NHIs), where AI agents operate at machine scale without governance.

Trends

Non-Human Identities (NHI): The Invisible Attack Surface

Non-human identities outnumber humans by a ratio of 25:1

Non-human identities have become one of the fastest-growing and least visible attack surfaces in modern cloud environments. A single agent can have unauthorized or overprivileged access to hundreds of resources, API tokens, and services across multiple platforms, often without defined ownership attribution or visibility. Service principals, system accounts, managed identities, OAuth applications, and guest accounts now outnumber human users in many Microsoft 365 tenants. These identities are designed to enable automation, integrations, and machine-to-machine communication, but they often operate with elevated privileges and limited security oversight.

Unlike human accounts, non-human identities authenticate continuously and at a large scale through APIs and automated workflows. Because they rarely trigger traditional user-based security monitoring, compromised service principals or leaked API credentials can remain active for extended periods without detection.

This makes them an attractive target for attackers seeking persistent and stealthy access to cloud environments. Analysis of authentication activity shows that all observed non-human identity access from known malicious IP addresses originates from the data center and hosting infrastructure. This pattern strongly suggests the exploitation of compromised service principal credentials or exposed API keys through attacker-controlled virtual servers. One particularly concerning observation includes Copilot access initiated from malicious infrastructure, indicating the potential use of automated tools to systematically query and extract sensitive organizational data from compromised environments.

Authentication Attack Rate



Analysis of authentication and account activity telemetry over the past 180 days revealed a sustained pattern of authentication abuse across monitored SMB cloud environments. Approximately 28-30% of authentication events fail, meaning nearly one in three sign-in attempts is unauthorized across all monitored regions. This ratio has remained consistently stable throughout the observation period, an indicator that the activity is not solely driven by isolated spikes but rather a reflection of a persistent, continuous attack posture targeting cloud identity systems.

Nearly one in three sign-in attempts is unauthorized

The stability of this failure rate strongly suggests ongoing automated credential attacks, commonly associated with techniques such as password spraying, credential stuffing, and distributed authentication probing using bot or proxy infrastructure. Rather than appearing as short-lived campaigns, these attacks form constant background pressure against authentication endpoints.

The composition of identity incidents over the six-month period also demonstrates a notable shift in attacker behavior. Early in the observation window, direct authentication attacks accounted for the largest share of incidents, representing approximately 55% of identity events in September. However, this category gradually decreased approximately 20-22% by February 2026, likely reflecting the impact of stronger identity protections such as multi-factor authentication (MFA), lockout policies, and improved identity threat detection mechanisms.

Conversely, Impossible Travel events increased significantly over the same period, eventually representing the majority of all identity incidents. These events typically indicate authentication activity from geographically inconsistent locations within a short time window, often associated with the use of distributed proxy infrastructure, residential VPN networks, or credential reuse across attacker-controlled systems.



While Adversary-in-the-Middle (AiTM) phishing and token theft incidents remain relatively low in overall volume, their presence indicates the emergence of more advanced post-authentication compromise techniques designed to bypass traditional MFA protections. Similarly, SharePoint data exfiltration events remain infrequent but carry high impact following successful account compromise.

The collective datasets indicate that SMB cloud environments face a persistent identity-focused threat landscape characterized by continuous credential attack attempts combined with a gradual shift toward session-based abuse and anomalous authentication activity.

Although defensive measures appear to be reducing the success rate of traditional credential attacks, adversaries are increasingly adapting their techniques to exploit session tokens, distributed infrastructure, and post-authentication access vectors.

These findings highlight the need for organizations to prioritize identity-centric security defenses, including advanced sign-in anomaly detection, conditional access enforcement, token monitoring, and continuous evaluation of authentication risk.


The Rise of Session Hijacking Activity

Session hijacking is spiraling. During an attack, a threat actor can intercept or take over a valid user session, typically by stealing session cookies or authentication tokens, allowing them to bypass login controls and access or resume a legitimate user’s session without credentials or MFA.

Session hijacking incidents increased by ~23% over the 180-day observation period

Over the 180-day observation period, session hijacking incidents increased by approximately 23%, making it the fastest-growing attack category among identity compromise techniques. Interestingly enough, despite the sudden spike in session hijacking incidents, the total number of affected users has remained relatively stable at approximately 20,000 accounts.

This pattern suggests that adversaries are not significantly expanding their targets, but are instead deepening their level of access within already compromised accounts. Attackers appear to be generating multiple session events per victim, likely through repeated session reuse, token replay, or the establishment of concurrent authenticated sessions across distributed infrastructure.




These key findings highlight the importance of session-level monitoring and token lifecycle controls, including the detection of anomalous session reuse, abnormal device or location changes, and enforcement of short-lived session tokens to reduce the operational window available to attackers. In modern environments, this is most often executed through Adversary-in-the-Middle (AiTM) phishing frameworks, which proxy authentication flows in real time.

The U.S. accounts for over 3/4 (75.4%) of recorded AiTM phishing incidents

Adversary-in-the-Middle (AiTM)

Malicious Infrastructure Authentication

Logins from IPs with high abuse reputation scores indicate either active compromise or attackers using known-bad infrastructure. The table below compares phishing proxies with session replays within a 120-day observation window.



Authentication events from known-malicious infrastructure grew **50% over the 120-day period**. This is the most concerning trend in the dataset. It indicates either a massive expansion of compromised infrastructure or an increase in activation attempts by threat actors who reuse flagged IPs.


Authentication events from known-malicious infrastructure rose by 50% over a 120-day observation period.

OAuth and Application Layer Attacks: The Known Frontier

OAuth abuse is an emerging threat vector that allows attackers to maintain persistent access through authorized application grants, often bypassing multiple password resets and MFA changes. One such example is via device code flow, which is a legitimate authentication method for devices without browsers (e.g., smart TVs). Attackers exploit it by manipulating users into entering a device code on an attacker-controlled authorization page.

We noticed a spike of +2,000% for Google Workspace (GWS) OAuth Abuse between September 2025 and February 2026.

180-Day Trend: OAuth2:Token requests consistently represent the single largest category of authentication events (~38% of total), reflecting both legitimate use and high attacker interest in token-based access.



Growth of OAuth Consent Activity

OAuth consent phishing enables attackers to trick users into granting permissions to **malicious applications**, providing persistent API access to resources such as **email, files, and contacts** without requiring ongoing credential use.

Over the 180-day observation period, **OAuth consent events increased significantly**, rising **45% between October and January** and an **additional 24% from January to February**. During the same period, the number of **unique users granting consent grew from approximately 2,300 to 3,500 per month**, indicating expanding exposure to OAuth-based phishing campaigns.

The Guardz research team also found a **consent failure rate of approximately 15-17%**, suggesting that a portion of potentially malicious consent attempts are being blocked by security controls or user rejection, highlighting both the growing prevalence of OAuth abuse and the importance of consent governance and application monitoring.

Cloud Productivity Tools at Risk

Global Top 10 Threat Operations

The table below highlights the top 10 security threats we found in M365 cloud productivity tools, which range from **Medium to Critical**.

Rank	Operation	Service	Risk
1	AnonymousLinkUsed	OneDrive	CRITICAL
2	AnonymousLinkCreated	OneDrive	CRITICAL
3	DLP Violations	OneDrive	HIGH
4	MessagesExported	Teams	HIGH
5	MemberAdded	Teams	MEDIUM
6	SharingPolicyChanged	SharePoint	HIGH
7	AnonymousLinkCreated	SharePoint	HIGH
8	SiteCollectionAdminAdded	SharePoint	HIGH
9	DLP Violations	Teams	HIGH
10	FileMalwareDetected	SharePoint	MEDIUM


Microsoft 365

Microsoft 365 cloud applications such as SharePoint, Teams, OneDrive, and Copilot form the operational backbone of SMB organizations. They also present security considerations.

These platforms store intellectual property, financial records, client communications, and strategic documents that can be easily compromised by threat actors. The telemetry reveals the massive scale of cloud activity and the attack surfaces hidden within it.


SharePoint Online

Public exposure of anonymous links allows anyone to access SharePoint content without authentication. The risk is when anonymous links bypass all authentication, conditional access policies, and DLP controls. The table below shows a spike in security threat indicators from December to February, which suggests either weak sharing policies or compromised accounts creating backdoor access for malicious actors.



Microsoft Teams

Over **3.1 million messages containing links** were sent via Teams in 180 days. Messages bypass traditional email security controls (no SPF/DKIM/DMARC), making them a highly targeted vector for phishing attacks.



Message export abuse occurs because export capabilities enable the bulk extraction of entire Teams conversations, making them high-value targets for insider threats and compromised accounts.

Application installation risks, specifically see a "Apps Installed in Teams", are a critical piece of this high-risk security puzzle.

Why App Installs Matter: Applications in Microsoft Teams are integrated services that can request and obtain delegated or application-level permissions via OAuth. Once granted, these permissions often extend beyond Teams into the broader Microsoft 365 ecosystem.

Common permission scopes include the ability to:

- Read channel and chat messages
- Access and exfiltrate files stored in SharePoint and OneDrive
- Retrieve user profiles, directory data, and contact lists via Microsoft Graph
- Maintain persistent access through refresh tokens or background services

These permissions introduce significant risk because they:

- Bypass traditional user-centric controls, operating via trusted application identities
- Enable large-scale data exfiltration without direct user interaction
- Are often granted once and rarely reviewed
- Can be abused by malicious or compromised third-party apps

Microsoft Copilot

Copilot has access to all content a user can access across Microsoft 365, including SharePoint, OneDrive, Teams, and Exchange. As a result, it acts as a force multiplier for existing permission misconfigurations and other security risks, such as:


- **Overshared SharePoint sites** become accessible via Copilot queries
- **Stale or excessive permissions** on OneDrive folders get surfaced to Copilot users
- **Sensitive Teams conversations** can be summarized, extracted, and recombined across threads
- **Exchange data** (emails, attachments, calendars) can be queried and correlated in a single output

Without proper AI governance and least privilege access controls, this creates a new class of risk for MSPs to manage.

Legacy Protocols

Legacy authentication protocols remain a critical attack surface across monitored SMB environments. Despite Microsoft’s ongoing deprecation of Basic Authentication, our 180-day analysis reveals persistent exploitation of WS-Trust, Basic Auth, and ROPC flows, alongside significant legitimate use of ActiveSync and EWS protocols, creating a range of detection blind spots.

The table below represents the legacy protocol landscape, consisting of 7.39M recorded events, and spread across five popular authentication protocols.



The current data indicates a high-risk environment characterized by massive automated brute-force activity and successful circumvention of security controls.

Identity Breach Risk


- **MFA Bypass Success:** 114,827 successful logins (4.5% of BAV2ROPC traffic) occurred where MFA was bypassed. This represents a confirmed breach of the primary security perimeter.
- **Validated Credentials:** 24,838 attempts (1.0%) resulted in a "Password OK" status but were halted by MFA requirements. These accounts are considered compromised at the first factor (password) and present an immediate risk of takeover if MFA is ever downgraded or misconfigured.

Operational & Productivity Risk

- **Account Lockout Volume:** 1,913,412 events (75% of BAV2ROPC traffic) resulted in account lockouts.
- **Denial of Service (DoS):** The high volume of lockouts suggests an automated attack that may inadvertently cause a "distributed denial of service" for legitimate employees, leading to significant helpdesk overhead and productivity loss.

Expanding Attack Surface

- **Growth in Attack Volume:** SMTP AUTH traffic, the largest legacy vector, increased by 51.3% between December 2025 and February 2026.
- **Targeting US Infrastructure:** The monthly trend shows a steady escalation in attempts within the US Region, particularly targeting **SMTP AUTH** (858K events/month) and **BAV2ROPC** (613K events/month).
- **Protocol Diversity:** While BAV2ROPC is a primary focus, threat actors are simultaneously leveraging **IMAP** (1.38M events) to probe for vulnerabilities, indicating a multi-vector approach to compromising legacy systems.



Google Workspace

There are currently over 3 billion active users of Google Workspace per month, as per Google’s own estimate. Google Workspace serves as the primary productivity suite for many SMBs. It also serves as a rising threat vector for malicious actors.

The Guardz research team found 125,983 confirmed suspicious logins

Analysis of recent telemetry indicates a sustained, targeted campaign against SMBs, specifically in the tech and startup sectors. The scale of unauthorized login attempts remains the primary entry attack vector, with 125,983 confirmed suspicious logins. The data also suggests that automated credential stuffing and "MFA fatigue" attacks are achieving high success rates, further compounded by non-human identities (NHIs) that retain overpermissive and unregulated access to cloud environments.

BAV2ROPC (Basic Authentication Version 2 Resource Owner Password Credential)

Why 114,827 successful sign-ins are critical? Every successful BAV2ROPC sign-in means:

- The attacker guessed or cracked the user’s password
- The user had no **MFA configured** (or MFA was disabled for legacy auth)
- The attacker received a valid **Exchange Online access token**

The Results

The attacker can now read email, send email, create inbox rules, and exfiltrate data and all without triggering interactive MFA.

Email Security


Email quarantine activity has surged by 240%

An analysis of 1.396 billion Exchange Online audit events across three regions showed a rapidly growing email environment with significant security trends. Email volume is expanding month-over-month (MoM), inbox rule modifications have doubled, quarantine activity has surged 240%, and nearly 2 million SendAs operations indicate widespread email impersonation, both legitimate delegation and potential BEC abuse.

An unsecured inbox also provides a treasure trove of data for a threat actor, including conversation histories, financial workflows, and embedded credentials. AI has further amplified this risk by enabling BEC attacks to be executed at scale with sentiment-aware and context-driven precision.


Client Access Distribution

The telemetry data represents a global distribution of **1.36 billion access events**. While the majority of traffic (55.6%) has migrated to modern REST APIs, legacy protocols (MAPI, EWS, IMAP, etc.) still account for a substantial portion of the environment's footprint, representing a persistent risk for credential stuffing and MFA bypass attacks.



Inbox Rule Manipulation - The BEC Persistence Engine

Inbox rules remain the **number one persistence mechanism** in BEC attacks (MITRE T1098.003). Attackers create rules to intercept financial communications, redirect emails to attacker-controlled accounts, and delete security alerts.



Common Malicious Rule Patterns

Pattern	Description
Forward-and-delete	Forward emails matching “invoice”/“payment” to the external, then delete
Move-to-RSS	Move inbound emails to RSS Feeds/Archive folder
Mark-as-read	Mark targeted emails as read to prevent victim notification
Block-sender	Block emails from the IT security team or password reset services


2025 BEC Risk Profile

Confirmed financial fraud incidents analyzed in 2026 ranged from \$140k to \$1.5M

While the average requested wire transfer in early 2025 hovered near **\$40,000** for broad-scale campaigns, targeted high-value individual loss events showed a much higher ceiling. Confirmed incidents analyzed this year ranged from **\$140,000 to \$1.5M**, highlighting the critical need for secondary out-of-band authentication for all high-value transfers.

The Story of a \$1.5M BEC Scam

The timeline below illustrates an AiTM email thread injection, where threat actors monitored legitimate communications before inserting themselves into the workflow.



Technical & Behavioral Analysis of the Attack

Domain Spoofing & Thread Injection

The attackers did not rely on a simple "cold" phishing email. By spoofing a known domain and inserting themselves into an active thread, they were able to leverage existing trust. This tactic bypasses standard "external sender" warnings if the spoofing is sufficiently sophisticated (e.g., look-alike domains or typosquatting).

Social Engineering "Race Condition"

The report highlights a "race condition" in human workflows, where attackers exploit the gap between a legitimate business need (e.g., payment) and the human verification process. Through persistence and repeated follow-ups, they mimic legitimate administrative behavior.

Voice Verification Failure

A critical takeaway is the completion of voice verification on October 8. This suggests one of two high-risk scenarios:

- **Social Engineering:** The attacker successfully impersonated the legitimate party over the phone.
- **Deepfake/AI:** The attacker leveraged audio synthesis to mimic a known contact's voice, a rising trend in high-value wire fraud.

Post-Compromise: The BEC "Silent Monitor" Phase

Once a threat actor bypasses the login layer, their behavior shifts from "noisy" brute force to "silent" persistence.

Persistence via Identity Manipulation

Attackers are prioritizing **2SV (MFA) changes and recovery info changes**. This ensures that even if the legitimate user changes their password, the attacker retains a "backdoor" via a secondary recovery email or a hardware key they control.

Exfiltration & Monitoring (BEC Indicators)


The most critical indicators of Business Email Compromise (BEC) are changes to Email Forwarding and Delegate Grant.

- **Delegate Grants:** Allows an attacker to read, send, and delete messages on behalf of the victim. This is preferred over forwarding because it leaves no trace in the "Sent" folder of the attacker’s external mail.
- **Blocked Senders:** The high volume of blocked sender changes indicates that attackers are proactively blocking security alerts from Google or emails from the IT department, keeping the victim in the dark.

Endpoint Security

Ransomware behavioral detections surged 190% over a 50-day observation window


Over a 50-day observation window, the platform processed **thousands of endpoint threats** across all three regions, with particular focus on RMM tool abuse, ransomware behavioral detections, and network reconnaissance tools, attack patterns that directly target the MSP ecosystem.



Key trend: Ransomware behavioral detections **surged 190%**, driven primarily by a confirmed lateral movement campaign at a US-based organization and by increased RMM tool flagging. Malware detections declined **55%** in the same period, suggesting a shift in attacker tooling from traditional malware to living-off-the-land techniques.

Endpoint Threat Volume & Trend Analysis

The tables below highlight endpoint threat volume over 180 days across all regions, from September 2025 to February 2026, including a monthly threat distribution of trends.



Monthly Threat Distribution

The threat volume follows a bimodal pattern with peaks in November 2025 and January 2026. The November spike correlates with increased hacktool activity (4.0% vs 1.3% baseline), suggesting a coordinated reconnaissance campaign. The spike in January was driven by a return of high-volume malware (**81.9% classification rate**), potentially indicating a post-holiday campaign restart by threat actors.

Month	Threats	Percentage	Trend
September 2025	2,292	19.1%	Baseline
October 2025	1,708	14.2%	-25.5% decline
November 2025	2,911	24.3%	+70.4% spike
December 2025	1,089	9.1%	-62.6% sharp decline
January 2026	2,415	20.1%	+121.8% surge
February 2026	1,585	13.2%	-34.4% decline

Global Classification Analysis

The following chart represents our key threat classification findings over a 180-day period.

Key Observations:

- **Malware dominates** at 66.3% (7,960 threats), with static analysis engines catching the majority before execution
- **Ransomware at 4.7%** (560 threats) remains the highest-impact category despite lower volume; each instance represents potential business-ending data loss
- **Hacktools at 1.6%** (192 threats) are particularly concerning as they indicate active adversary operations, tools like Advanced IP Scanner and Nmap suggest internal network reconnaissance
- **Benign at 4.7%** (563 threats), appearing only in the Australia region, an indicator of regional variation in detection tuning


Classification	Count	Percentage	Risk Level
Malware	7,960	66.3%	HIGH
General	2,579	21.5%	MEDIUM
Benign	563	4.7%	LOW
Ransomware	560	4.7%	CRITICAL
Hacktool	192	1.6%	HIGH
PUA	97	0.8%	LOW
Cryptominer	46	0.4%	MEDIUM

Critical Trend:

December Ransomware Spike: Ransomware proportion peaked at **8.2% in December 2025**, nearly double the 180-day average. This correlates with known threat actor behavior of targeting organizations during holiday staffing shortages when incident response capacity is reduced.

Threat Campaign Analysis

The following chart represents our key threat classification findings over a 180-day period.



Campaign 1: Remote Management Tool (RMM) Abuse

MITRE ATT&CK: T1219 (Remote Access Software), T1021 (Remote Services) **Severity:** CRITICAL

Tool	Detections	Behavior
MeshAgent.exe	190+	Persistence with an invalid signature, unauthorized deployment
ScreenConnect.ClientSetup.exe	294+	Remote access installation via command line
AteraAgent	105+	RMM agent deployment for persistent access
ISLOnline Agent	63+	Remote support tool abuse
NinjaRMMAgentPatcher	26+	Patching/modification of existing RMM

Guardz Analysis: RMM abuse is the #1 threat campaign, accounting for 26.2% of all threats. Attackers are deploying legitimate remote management tools (e.g., MeshCentral, ScreenConnect/ConnectWise, Atera, NinjaRMM) to establish persistent, encrypted command-and-control channels that blend with legitimate IT management traffic. This technique is particularly effective in MSP environments where multiple RMM tools may be expected.

RMM tools appear in our behavioral detection telemetry not because they are malicious, but because their normal operation is indistinguishable from attacker tooling at the process and network level, and legacy signature-based endpoint detection cannot make the distinction.

Risk: RMM tools provide full remote desktop access, file transfer, command execution, and persistence equivalent to a full-capability RAT but with trusted digital signatures.

Campaign 2: Network Scanning & Reconnaissance

MITRE ATT&CK: T1046 (Network Service Discovery), T1018 (Remote System Discovery) **Severity:** HIGH

Tool	Detections	Purpose
nmap.exe / domotz_nmap.exe	514+	Port scanning, service enumeration
Advanced IP Scanner	701+	Network host discovery
Advanced Port Scanner	144+	Port scanning
netscan.exe	18+	Network scanning

Guardz Analysis: Network scanning tools represent 14.8% of all threats, indicating active internal reconnaissance. The combined use of Nmap and Advanced IP Scanner is a common pattern among initial access operators mapping internal networks prior to lateral movement.

Risk: The presence of domotz_nmap.exe suggests abuse of the Domotz network monitoring platform as cover for scanning activity.

Campaign 3: Commodity Malware (Meson/Player Cluster)

MITRE ATT&CK: T1204 (User Execution), T1059 (Command and Scripting Interpreter) **Severity:** HIGH


Variant	Detections	Classification
player.exe / player (2).exe	345+	General/Malware
meson.exe	88+	Malware
0.exe / 1.exe / 2.exe	343+	Malware
cmd.exe (CLI variants)	427+	General

Guardz Analysis: This cluster represents commodity malware using generic filenames to evade initial detection. The player.exe family and numerically-named executables (0.exe, 1.exe, 2.exe) are characteristic of dropper stages in multi-stage malware delivery chains.

Risk: The high volume of malware suggests automated distribution via phishing or drive-by download campaigns, which is exceptionally difficult to detect at this scale.

Campaign 4: Ransomware Operations

MITRE ATT&CK: T1486 (Data Encrypted for Impact), T1490 (Inhibit System Recovery) **Severity:** CRITICAL



Guardz Analysis: **560 ransomware threats** were detected across 180 days. The December 2025 spike to 8.2% of all threats (89 incidents) correlates with known holiday-period targeting patterns. Key ransomware targeting includes

Risk: Several ransomware variants abuse legitimate process names (OfficeClickToRun.exe, SRService.exe) through DLL sideloading and process injection, making detection more challenging and bypassing application whitelisting.

Campaign 5: Driver-Level Exploitation

MITRE ATT&CK: T1068 (Exploitation for Privilege Escalation), T1014 (Rootkit) **Severity:** CRITICAL

Driver	Detections	Abuse Type
DBUtilDrv2.sys	231	Dell BIOS utility - BYOVD
LeCrud64.sys	33	Vulnerable driver exploitation
ALSysIO64.sys	14	Vulnerable signed driver
wsddprm.sys	62	Suspicious kernel driver

Guardz Analysis: BYOVD (Bring Your Own Vulnerable Driver) is an escalating threat. Attackers deploy legitimately signed but vulnerable drivers to gain kernel-level code execution, disable EDR, and establish rootkit-level persistence. DBUtilDrv2.sys (CVE-2021-21551) remains the most abused driver in this dataset.

Risk: Direct interaction with kernel structures allows attackers to terminate security processes, disable telemetry, and bypass critical EDR protections.

Prediction #2: AI Attacks at Scale

The 2.65 million invalid password attempts we observed suggest that attackers are already using automated tools. In H2 2026, we predict the following trends:

- **LLM-powered social engineering** will make phishing emails indistinguishable from legitimate communications
- **AI-generated password lists** based on leaked data will increase first-attempt success rates
- **Deepfake voice calls** will target SMB finance teams for wire fraud

 **Guardz Recommendation:** Implement passwordless authentication and Zero Trust conditional access policies.

Prediction #3: Attacks Against the MSP Supply Chain Will Intensify

The detection of **spoofed AteraAgent MSI** files and **ScreenConnect** abuse signals a growing trend:

- Attackers will increasingly impersonate RMM tools used by MSPs to establish persistent access
- Compromising a single MSP can provide access to **hundreds of downstream SMB clients**
- Nation-state groups will adopt MSP-targeting techniques previously limited to cybercriminal groups

 **Guardz Recommendation:** MSPs should implement application allowlisting and monitor RMM tool deployments with hash verification.

Prediction #4: Business Email Compromise Will Evolve Beyond Inbox Rules

With the new inbox rules and mailbox configuration changes by the attackers, we predict that in H2:

- BEC actors will shift from **inbox rules to Microsoft Graph API abuse** for email manipulation
- **Consent phishing** through **OAuth application grants** will replace traditional credential phishing
- Multi-stage BEC attacks will combine **identity compromise, email persistence, and data exfiltration** into a single campaign

 **Guardz Recommendation:** Monitor OAuth consent grants and Graph API activity alongside traditional mailbox audit logs.

Threat Risk Assessment Matrix

The table below provides an overview of the top five campaign threats and business impact by severity level and evasion capabilities.

Threat Risk Assessment Matrix — Category x Dimension				
RMM Tool Abuse	CRIT+	HIGH	CRIT	CRIT
Network Scanning	CRIT	MED	MED	MED
Commodity Malware	CRIT	HIGH	HIGH	HIGH
Ransomware	MED	CRIT+	CRIT	CRIT+
Driver Exploitation	MED	CRIT	CRIT+	CRIT
Script-Based Attacks	MED	HIGH	CRIT	HIGH
Info Stealers	LOW	CRIT	CRIT	CRIT
Cryptomining	LOW	MED	HIGH	MED
	Prevalence	Severity	Evasion Capability	Business Impact

H2 2026 Threat Predictions

Based on the current trends observed in this reporting period and broader threat intelligence, Guardz anticipates the following threat developments for the second half of 2026:

Prediction #1: AiTM Phishing Kits Will Become Commodity Tools

Session-hijacking events represent a sharp increase over previous periods. As AiTM phishing kits, such as Tycoon 2FA, Evilginx 3, and Greatness, become more accessible through phishing-as-a-service (PhaaS) marketplaces, we predict that:

- **Session token theft to surpass password theft** as the primary compromise vector for cloud identities
- Attackers will increasingly target **OAuth tokens and refresh tokens** rather than session cookies

87% of U.S. and UK workforces are deploying passkeys for employee sign-ins

 **Guardz Recommendation:** Deploy phishing-resistant MFA (FIDO2/passkeys) to prevent session token theft and enhance workforce authentication security.

Prediction #5: Ransomware Targeting SMBs Will Adopt Double-Extortion by Default

Attackers will exploit anonymous SharePoint links as exfiltration channels

Ransomware detections and confirmed lateral movement activity indicate that operators continue to target SMB infrastructure. Here's what we see moving forward in H2:

- **Cloud-native ransomware** targeting SharePoint and OneDrive will emerge, encrypting files at the cloud storage layer rather than the endpoint
- SMBs will be targeted for **data-theft-only extortion** (no encryption), which is harder to detect and cheaper to execute
- Attackers will exploit anonymous SharePoint links as exfiltration channels

 **Guardz Recommendation:** Enable versioning and retention policies on all SharePoint document libraries; monitor for bulk anonymous link creation.

Prediction #6: Google Workspace Will See Increased Targeting

With 126,000 suspicious GWS logins already observed, we believe that:


- Attackers will expand from Microsoft 365-focused tooling to **dual-platform attack kits** targeting both M365 and GWS
- **Google OAuth token abuse** through malicious Chrome extensions will increase.
- The gap between Microsoft 365 and GWS security tooling will be exploited by attackers who recognize that GWS tenants often have weaker security controls.

 **Guardz Recommendation:** Enable Google Advanced Protection for admin and high-value accounts; deploy ITDR coverage across both Google Workspace and Microsoft 365 environments.

Connecting the Security Dots with Guardz

Every client environment an MSP manages is a potential attack vector and opportunity waiting to be exploited. The same remote access capabilities that let a security professional resolve a client issue in minutes are what an attacker uses to deploy ransomware across a managed portfolio in hours.

This is why MSPs must take a proactive approach to securing all attack vectors.



Guardz helps MSPs connect the security dots across all endpoints, identities, cloud services, email platforms, and data infrastructure through a unified cybersecurity platform.

Guardz MDR unifies SentinelOne EDR, ITDR, and other platform detections into a single, contextualized system of normalized incidents, backed by 24/7 expert security coverage.

Guardz AI-SOC deploys a **multi-model architecture** for fast triage and enrichment, mid-tier models for investigation, and frontier models for senior review of complex cases, achieving **92.4% accuracy across benchmarked incidents**.

Correlate signals to real identities, continuously monitor user behavior and activity across Microsoft 365 and Google Workspace environments, and deliver contextual threat intelligence with Guardz ITDR.

Discover the power of the Guardz platform for your MSP business.

Book a Demo

Watch On Demand Demo