Consultiply

• AI

MIP RESOURCE CENTER

Check your GDAP, service principals and leaked secrets after Microsoft's agentic AI warning

Microsoft Threat Intelligence reports that threat actors are using the agentic AI framework ARTEX to find exposed secrets on GitHub and use them against cloud tenants. Here is what MSPs should check across Partner Center, GDAP and customer tenants this week.

MICROSOFT SECURITY BLOG

ARTEX

NOV 5, 2026

Source: Microsoft Security Blog, Threat actors use agentic AI tooling to find and abuse exposed secrets, Nov 5, 2026. Independent reading; not affiliated with or endorsed by Microsoft.

What Microsoft saw

On Nov 5, 2026, Microsoft Threat Intelligence reported that threat actors are using an agentic AI offensive framework, ARTEX, to search public GitHub repositories for exposed secrets. These include credentials, application credentials and AI inference-related API keys.

Searching public code for leaked keys is not new. What changes, according to Microsoft, is speed and autonomy. The agent finds the secret, analyzes it and runs the follow-on attack with little need for a human at the keyboard. In at least one case, Microsoft observed an AI agent go from scanning GitHub for secrets to exfiltrating data from Azure Key Vault in just five minutes.

Key points from Microsoft's write-up:

- **Starting in August 2026**, Microsoft observed an agentic AI framework submitting large volumes of GitHub queries for secrets. Queries often combined a target organization's name with keywords for secrets, cloud services or AI providers.
- **As early as Oct 28, 2026**, Microsoft began observing successful service principal sign-ins from **199.19.77[.]84** involving multiple application IDs across some impacted customers. This included cases where a multitenant, vendor-owned application credential was compromised at the source and used to reach multiple downstream tenants.
- The sign-ins primarily came from two IPs, per Microsoft: **199.19.77[.]84** and **199.19.77[.]18**. Microsoft notes the second is a Tor exit node and may carry unrelated malicious traffic.
- Microsoft reports the agent successfully conducted service principal sign-ins on **dozens of associated applications across multiple tenants**.
- Microsoft also identified malicious operations from **65.87.223[.]26**, potentially associated with this campaign, including signs of vulnerability exploitation via shell commands embedded in a user agent and attempts to obtain App Service authentication tokens. Separately, Microsoft observed reconnaissance from **157.22.126[.]128** targeting AI workloads on Azure AI services.

After sign-in, Microsoft describes the attacker going after anything that returns credentials: storage account keys, App Service settings, connection strings, publishing credentials, AI and ML connections, and AKS user credentials. Compromised apps then accessed SharePoint and OneDrive files and mailbox items, ran Microsoft Graph searches, enumerated Azure subscriptions and permissions, searched Azure DevOps for secrets, and used Run Command on Azure VMs. Microsoft also notes the agents showed a strong focus on credentials for AI services.

Facts table

Claim	Figure	Status
Agentic AI framework (ARTEX) used to search public GitHub repos for exposed secrets	Observed by Microsoft Threat Intelligence	Microsoft Security Blog
Time from GitHub secret scanning to Azure Key Vault exfiltration, in at least one case	Five minutes	Microsoft Security Blog
Start of large-volume agentic GitHub queries for secrets	August 2026	Microsoft Security Blog
Earliest successful service principal sign-ins from 199.19.77[.]84	Oct 28, 2026	Microsoft Security Blog
Primary sign-in IPs	199.19.77[.]84 and 199.19.77[.]18 (Tor exit node)	Microsoft Security Blog
Applications with successful service principal sign-ins	Dozens, across multiple tenants	Microsoft Security Blog
Additional IPs potentially associated	65.87.223[.]26 (malicious operations), 157.22.126[.]128 (reconnaissance of Azure AI workloads)	Microsoft Security Blog
MSP-specific guidance	Review Partner Center, GDAP, customer tenants; notify customers per agreements	Microsoft Security Blog

Why this matters to an MSP

Microsoft called out MSPs by name in its containment guidance. The way MSPs work maps closely to the attack path Microsoft describes.

In plain language:

- **Service principals are app identities.** When a script, integration or automation tool signs in to Microsoft Entra ID, it does so as a service principal using a client secret or certificate from its app registration. There is no MFA prompt and no user to notice. If the secret leaks, whoever holds it is your app.
- **MSPs create a lot of these.** Backup connectors, RMM and PSA integrations, reporting scripts, onboarding automation and AI pilots all tend to come with an app registration and a secret. Many were set up quickly, and secrets sometimes end up in a script, a config file or a repo.
- **Multitenant apps multiply the blast radius.** Microsoft describes cases where a vendor-owned multitenant application credential was compromised at the source and used to reach multiple downstream tenants. If your MSP runs a multitenant app consented into customer tenants, one leaked secret can open many doors.
- **GDAP and Partner Center are the keys to the estate.** Granular delegated admin privileges and Partner Center access define what your team, and anything acting as your team, can do in customer tenants. They need to be reviewed alongside app credentials, not separately.
- **Deleting the leak is not a fix.** Microsoft is explicit that publicly exposed credentials stay usable until revoked or rotated, even if the original file or commit was edited or deleted.

- **Speed changes the response window.** If an agent can move from a leaked secret to Key Vault in minutes, the realistic defense is to not have long-lived secrets lying around, and to rotate fast when one is found.

Step-by-step checklist for MSPs

These steps come from Microsoft's own mitigation, detection and hardening guidance in the Nov 5, 2026 post, organized for an MSP running many tenants.

1. Review Partner Center access, GDAP relationships and customer tenants

- List every active GDAP relationship and the roles granted in each. Remove relationships and roles that are no longer needed.
- Review who in your organization has Partner Center access and which security groups map to customer roles.
- Note which customer tenants have your multitenant apps consented. Those tenants move to the top of the hunt list.

2. Hunt service principal and non-interactive sign-ins back to Oct 28, 2026

- In each tenant, review Entra service principal sign-in logs and non-interactive sign-in logs for **199.19.77[.]84** and **199.19.77[.]18**, including failed attempts. Microsoft's guidance starts its hunting window at Oct 28, 2026, and also advises searching as far back as retention allows.
- Watch for invalid or expired secret errors on your apps. Microsoft says these indicate someone holds a credential for your app, and that a failure followed by a success is urgent.
- Extend the search to Azure Activity, Key Vault and Storage diagnostics, Entra audit, Azure DevOps audit and Microsoft 365 unified audit logs. Look for secret reads, blob and file downloads, and changes to apps, permissions or pipelines. Include **65.87.223[.]26** and **157.22.126[.]128** in the indicators you check.
- Look for behavior, not just IPs. Microsoft highlights one address authenticating to several service principals with mixed success, or one app requesting tokens across Azure Resource Manager, Key Vault, Storage and Graph in a short window. Microsoft notes attackers can and do change IPs and user agents.
- Microsoft publishes advanced hunting queries in the post. Use them from the source page directly.

3. Audit app registrations and credentials, then remove and rotate

- Inventory app registrations and enterprise apps in your own tenant and in customer tenants you manage. Note every client secret, certificate and federated credential.
- For any affected application, Microsoft advises removing all client secrets, certificates and federated credentials, including any added by the attacker, and disabling the service principal if needed. Tokens already issued stay valid until they expire.
- Rotate everything stored in any Key Vault the affected application could read, plus any credentials found in documents it accessed.
- For affected users, revoke Entra sessions, Azure DevOps PATs and SSH keys.

- In Azure DevOps auditing, review PAT creation and first use, federated service connections, and pipeline runs and edits. Turn auditing on if it is off.

4. Scan public repos for secrets, and rotate even if the leak was deleted

- Scan your own GitHub organizations and any repos your engineers have published. Microsoft points to GitHub secret scanning, public monitoring for leaks in other public repos, and GitHub Advanced Security for Azure DevOps with push protection.
- Microsoft's position is clear: treat any publicly exposed secret, and the application behind it, as compromised. Removing the commit does not remove copies in history, caches, archives or logs.
- Longer term, Microsoft recommends moving apps off secret-based authentication and onto managed identities where possible.

5. Notify affected customers per your agreements

- If your hunt finds sign-ins, credential access or data access in a customer tenant, follow the notification terms in your contracts and any regulatory obligations that apply.
- Be specific: which app, which time window, which resources, and what you have rotated.

6. Optional containment and hardening items Microsoft lists

These are Microsoft's recommendations. Check licensing and fit for each tenant.

- **Conditional Access block:** Microsoft recommends blocking **199.19.77[.]84/32** as a named location in a Conditional Access policy targeting workload identities, assigned directly to service principals, plus a user-scoped block policy to cover replayed sessions. Microsoft notes that workload identity policies require Workload ID Premium licensing and cover only single-tenant apps registered in your tenant. Block the exact IP, not the whole ASN.
- **Key Vault rotation:** rotate all secrets in any vault an affected app could read, as above.
- **Least privilege:** review Azure RBAC on service principals and restrict them to only what each app needs. Limit who can register apps by setting user app registration to "No" and limiting the Application Developer role.
- **Defender for Cloud plans:** Microsoft suggests enabling relevant plans for critical workloads, including Defender for Resource Manager, Storage, Key Vault, App Service and Databases.
- **Protect backups:** restrict access to backup and recovery resources and monitor attempts to change their protection.
- **Deception:** Microsoft reports that honeypots, decoys, tarpits and fake wins helped reduce attack traffic reaching real assets in its testing.

Closing

Microsoft's report is a reminder that an app secret in a public repo is now a live incident, not a hygiene backlog item. For MSPs, the practical response is a short, repeatable routine: know every app identity you own, know every GDAP relationship you hold, scan for leaks, rotate fast, and tell customers when it touches them.

Reach us at transform@consultiply.ai

Sources

- Microsoft Threat Intelligence and Microsoft Security Research, Microsoft Security Blog, Nov 5, 2026: Threat actors use agentic AI tooling to find and abuse exposed secrets

Verified quotes

All quotes are verbatim from Microsoft Security Blog, Nov 5, 2026.

- "In at least one case, Microsoft observed an AI agent advance from scanning GitHub for secrets to exfiltrating data from Azure Key Vault in just five minutes."
- "MSPs: Review Partner Center access, GDAP relationships, and customer tenants, and notify affected customers per your agreements."
- "Publicly exposed credentials remain usable until revoked or rotated; removing the original disclosure alone does not remediate the exposure."
- "Organizations that find their secrets exposed in public repositories must assume that the exposed secrets and associated applications are compromised."
- "The agent successfully conducted service principal sign-ins on dozens of associated applications across multiple tenants."

Availability caveats

This guide is Consultiply's independent reading of a single Microsoft Security Blog post published Nov 5, 2026. Consultiply is not affiliated with or endorsed by Microsoft. Indicators, licensing requirements and recommended controls may change as Microsoft updates its research. Microsoft notes that 199.19.77[.]18 is a Tor exit node that may carry unrelated traffic, and that attackers can change IPs and user agents, so IP matches are a starting point, not proof. Always confirm current guidance on the source page and validate changes in a test tenant before rolling them out to customers.
