CSAI Foundation | Cloud Security Alliance

Entra Connect Health Exposure: Tenant Metadata as Identity-Attack Targeting Data

Security Implications and Guidance for Hybrid Identity Defenders

2026-11-24

AI-assisted Rapid Research

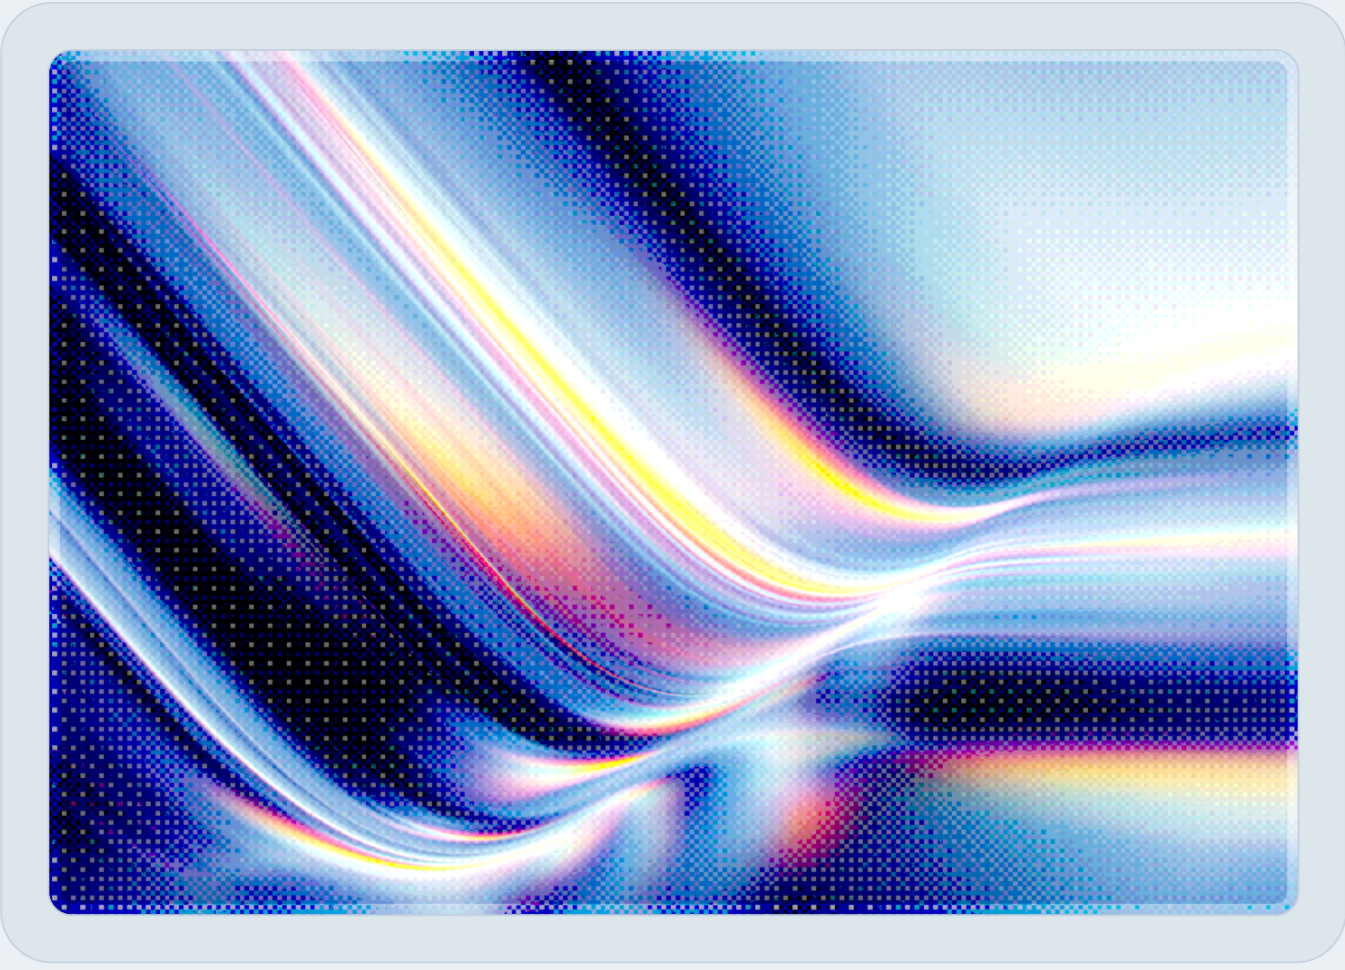

CSAI

CSA cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

Microsoft has notified customers that a vulnerability in the Microsoft Entra Connect Health API, introduced on July 23, 2026 and detected on November 9, allowed an unauthorized party to read tenant-level metadata. Microsoft fixed it service-side on November 19, 2026 [1][2]. For tenants that received the notification, Microsoft stated that "unauthorized read-only access to tenant-level metadata associated with your organization had occurred," and that it "did not identify any unauthorized configuration changes" [1][2]. The flaw went undetected for roughly 109 days (July 23 to November 9) and remained exploitable for about 119 days until the November 19 fix [4][5]. Microsoft has not disclosed who accessed the data, how many tenants were affected, when access began, or what class of flaw was involved, and none of the published coverage identifies a CVE [1][4][5].

Credentials were not exposed: Microsoft reports that no passwords, hashes, tokens, or valid keys were returned [1][3]. The exposed fields are still significant. They included administrator and technical-contact email addresses, tenant identifiers and initial domains, country and region, Entra license tier, directory synchronization status and last sync time, and Connect Health onboarding status [1][4][5]. For any tenant whose record was accessed, these fields amount to a pre-qualified target profile. The record names the people who hold privilege, shows whether the tenant is hybrid, and hints at which defensive features it has licensed.

Microsoft says no software update or customer remediation is required [1][3]. CSA's assessment is that this advice covers the vulnerability, not the downstream risk. Affected organizations should prepare for targeted phishing, password spraying, help-desk social engineering, and consent-phishing aimed at the named administrators, although no such campaigns have been publicly reported as of this writing [5]. They should confirm that every privileged account uses phishing-resistant authentication. They should also review identity logs back to July 23, 2026.

Background

Microsoft Entra Connect Health is a cloud monitoring service for the on-premises side of hybrid identity. Lightweight agents installed on Entra Connect Sync servers, Active Directory Federation Services (AD FS) servers, and Active Directory Domain Services (AD DS) domain controllers report health, performance, and audit data to a Microsoft-operated service. Administrators view that data in the Entra

admin center [6]. Agents are licensed per server against Entra ID P1 or P2. Installing the Connect Health for Sync agent requires a Global Administrator [6]. As a result, the service's tenant records are tied to the most privileged layer of a hybrid identity deployment. Anyone who can query those records learns something about that layer.

According to the notification text quoted by several outlets, Microsoft detected a vulnerability in the Connect Health API on November 9, 2026, traced it to a change made on July 23, 2026, and deployed a fix on November 19 [2][3][4]. On November 20, Microsoft concluded its investigation and began notifying affected organizations individually through Azure Service Health, with email copies to administrators [1][4][5]. A Microsoft spokesperson described the issue as one "that could have resulted in unauthorized access to tenant-level contact and configuration metadata" and declined to give further detail [1]. Administrators reported the notices in public forums. Some managed service providers said they received separate notices for several customer tenants, arriving hours apart [1].

The notification says the flaw could also have allowed an unauthorized party to "potentially make changes to Connect Health monitoring configuration," although Microsoft reports finding no such changes [2]. Microsoft published read-only PowerShell scripts that query the Connect Health APIs and show the *types* of metadata that could have been returned for a tenant [1][7]. These scripts report current values. They do not show which records were actually accessed [1][5].

The table below consolidates the exposed and non-exposed data categories as reported across the published coverage. BleepingComputer's list is the most complete; the other outlets report subsets [1][2][3][4].

Data category	Reported as exposed	Attacker utility
Administrator and technical-contact email addresses	Yes	Named, high-value phishing and spraying targets
Tenant ID, display name, initial `onmicrosoft.com` domain	Yes	Links organizations to tenants; supports tenant-specific lures
Country code and region	Yes	Language, time-zone, and regulatory tailoring of lures
Entra license tier	Yes	May indicate whether P1/P2 controls (Conditional Access, PIM, risk-based policies) are available

Data category	Reported as exposed	Attacker utility
Directory sync status and last sync timestamp	Yes	Marks tenant as hybrid; points toward on-premises sync infrastructure
Connect Health onboarding status	Yes	Indicates AD FS / sync monitoring footprint
Expired certificate for an internal Microsoft agent-registration service decommissioned in late 2024 [1][4]	Yes	Microsoft states no risk; no longer valid [1][4]
Passwords, hashes, tokens, valid keys	No	—
On-premises server or AD topology; Azure subscription data	No	—

Security Analysis

Why Tenant Metadata Matters to Attackers

Treating this incident as low severity because only "metadata" was exposed would miss its value to an attacker. Identity-plane intrusions against Microsoft cloud tenants typically begin with reconnaissance, in which the attacker must work out which tenant belongs to which organization, who the privileged users are, whether the environment is hybrid, and which controls are likely in place. Much of this is normally gathered piecemeal from public DNS, OpenID discovery documents, and social media. The Connect Health API returned all of it as one structured record per tenant. If access extended across many tenants, which Microsoft has neither confirmed nor ruled out, records of this kind would let an adversary build target lists across organizations at once and prioritize those that appear most valuable or least defended. Even taking Microsoft's description of the access as read-only at face value, a read-only reconnaissance dataset can still be the first stage of a write-capable attack.

The individual fields help an attacker in specific ways. The administrator email addresses are the most direct asset. Technical contacts in Microsoft tenants are often Global Administrators, a point Secureworks' Counter Threat Unit made in research conducted in 2021 and published in April 2022, after

a similar Microsoft API exposed technical-contact names and phone numbers for arbitrary tenants [8]. The license tier is more subtle. A tenant without P1 licensing likely lacks Conditional Access entirely, while a P2 tenant probably has Privileged Identity Management and risk-based policies available. Because Connect Health itself requires P1 or P2, the more useful signal for tenants actively using the service is likely the distinction between P1 and P2, and thus whether PIM and risk-based policies are available. This is an inference about capability, not configuration, but it lets an attacker decide where to spend effort. Sync status and last-sync time confirm a hybrid deployment. In a hybrid deployment, the on-premises Entra Connect server is a known high-value target, and compromising it has historically yielded cloud-privileged credentials [9].

Hybrid Identity Raises the Stakes

Knowing a tenant is hybrid gives an attacker a second route into the cloud tenant, through the on-premises sync and federation infrastructure that Connect Health monitors. Security researchers have shown repeatedly that the Entra Connect server is effectively a Tier 0 asset. SpecterOps documented in 2025 that after Microsoft moved Connect Sync from a password-based connector account to certificate-based application authentication, an attacker with access to the sync server's ADSync context could still request tokens. That attacker could also add their own certificate to the Entra Connect application, giving persistent off-host access without disrupting synchronization [9]. The Connect Health exposure did not reveal server names or AD topology [1]. It did tell an attacker which organizations run this infrastructure and when it last synced, which narrows where an on-premises intrusion would pay off most.

A Recurring Pattern in Shared Identity Services

This incident continues a pattern of authorization flaws in Microsoft-operated identity service planes, several of which returned tenant data or granted access across tenant boundaries. In research conducted in 2021 and published in April 2022, Secureworks' Counter Threat Unit reported Microsoft diagnostic and admin-center APIs that exposed licensing, directory synchronization status, and technical-contact details for arbitrary tenants. Microsoft remediated these in stages through April 2022 [8]. In 2025, researcher Dirk-jan Mollema disclosed CVE-2025-55241, in which undocumented service-to-service "actor tokens" combined with a validation flaw in the legacy Azure AD Graph API could allow impersonation of any user, including Global Administrators, in any tenant. Microsoft fixed that issue without evidence of exploitation [10][11]. The table below summarizes these cross-tenant cases alongside the current incident. Related context comes from a flaw that was not cross-tenant: earlier in 2026, CSA analyzed an Entra Agent ID

Administrator scoping flaw that allowed service principal takeover within a single tenant [15]. That case is excluded from the comparison, but it shows that authorization errors in Entra's service plane are not limited to legacy APIs.

Incident	Year	Exposure	Vendor-confirmed unauthorized access
Secureworks Azure AD diagnostic/admin APIs [8]	2021–2022	Licensing, sync status, technical-contact name and phone, cross-tenant	No
Actor tokens / CVE-2025-55241 [10][11]	2025	Potential cross-tenant Global Admin impersonation	No
Entra Connect Health API [1][2]	2026	Tenant metadata and administrator contacts	Yes (read-only, per Microsoft)

The Connect Health incident differs from these predecessors in one important way: Microsoft confirmed that unauthorized access actually occurred [1][2]. The flaw went undetected for 109 days. The length of that interval raises questions about provider-side monitoring of this API, which Microsoft has not addressed; it has not said how the issue was eventually found. Customers cannot see API activity within Microsoft's own service plane. They therefore depend entirely on provider disclosure, and that disclosure has so far left out who the actor was and how wide the access went.

What Attacks to Expect

No phishing campaigns using the exposed data have been publicly reported as of this writing [5]. However, the exposed data fits several established attack techniques. Administrators at affected tenants are plausible targets for adversary-in-the-middle phishing kits and device-code phishing, which CSA analyzed in its July 2026 research note on the Forg365 phishing-as-a-service platform [14]. The notification itself creates a ready-made lure: an email that imitates a Microsoft "Important Security Notification Regarding Microsoft Entra Connect Health" sent to the same administrators who just received the real one [5]. Password spraying against the named accounts, help-desk impersonation requesting MFA resets, and OAuth consent requests aimed at privileged users are all lower-effort follow-ons. Administrators who still authenticate with passwords plus SMS or push-based MFA remain the most exposed, because adversary-in-the-middle kits can relay or bypass those factors [14]. Generative AI lowers the cost of tailoring these lures to an organization's language, region, and tenant name, all of which were in the exposed record [14].

Recommendations

Immediate Actions

First, confirm whether your organization received a notification. To avoid the phishing risk described above, check by going to the Azure portal directly and opening Service Health security advisories, not by following links in email [2][5]. Managed service providers should check each customer tenant individually, because notifications were issued per tenant [1][4]. Organizations with a notice should then run Microsoft's read-only Connect Health scripts after reviewing their contents [7]. This lists the administrator and technical-contact addresses, license data, and sync attributes that were potentially readable. Those named individuals are now the priority population for the steps below.

Next, verify that every account on that list, and every account holding Global Administrator, Hybrid Identity Administrator, Privileged Role Administrator, or Application Administrator, authenticates only with phishing-resistant methods such as FIDO2 security keys, passkeys, or certificate-based authentication. Enforce this through Conditional Access authentication strengths [1][2][5]. Tenants that permit device-code flow for administrators should block it unless there is a documented business need [5]. Brief IT, help-desk, and security operations staff that the named administrators are likely phishing and social-engineering targets. Require out-of-band verification before any MFA reset or credential change for privileged users [2][5].

Short-Term Mitigations

Review Entra sign-in and audit logs and Azure Activity Logs from July 23, 2026 forward [2][3][5]. Look for password-spray patterns against the exposed accounts, device-code or legacy-protocol sign-ins by administrators, administrator sign-ins that did not satisfy phishing-resistant MFA, unexpected MFA method registrations, new OAuth application consents or credentials added by privileged accounts, and new key credentials on the Entra Connect Sync application [5][9]. In Connect Health itself, confirm that the registered Sync, AD FS, and AD DS servers belong to your organization, that alert notification recipients are expected, and that alerting has not been disabled. The notification acknowledged that configuration changes were possible, even though Microsoft found none [2][5]. Published hunting guidance indicates that credential rotation is not required for this incident alone [5]. Where log review turns up anything anomalous, however, rotating affected credentials remains a prudent precaution, following Microsoft's documented procedures to avoid disrupting synchronization or federation.

Organizations should also reduce what any future metadata leak can reveal. Secureworks' 2022 guidance is still relevant: populate the tenant's technical-contact details with a generic, monitored role (for example, "IT Department") rather than an individual administrator identity [8]. Privileged accounts should be cloud-only, separate from daily-use identities, and ideally without a mailbox, so that their addresses are neither guessable nor useful as phishing delivery targets [5]. Using Privileged Identity Management for just-in-time elevation [2] shrinks the set of standing privileged accounts an attacker can usefully target.

Strategic Considerations

The broader lesson concerns shared responsibility in identity-as-a-service. Customers cannot monitor or log the provider-side APIs that store data about their tenant. They should plan as though tenant metadata, administrator identities, and configuration signals will eventually become known to adversaries. That argues for designing identity defenses that do not depend on keeping metadata secret. Phishing-resistant authentication, least-privilege role assignment, and Tier 0 treatment of Entra Connect servers and AD FS infrastructure [9] all remain effective whether or not an attacker knows who the administrators are.

Organizations should also raise transparency expectations with their identity providers through contract and assurance channels. Useful asks include timelines for detection and notification, disclosure of the scope of access and the threat actor where known, and provider-side access logs available to affected tenants. The 109-day detection gap, the roughly 119-day exposure window, and the absence of a CVE or published root cause make independent risk assessment difficult [4][5]. Finally, hybrid organizations should revisit whether their on-premises sync and federation footprint can shrink. Each migration step, such as moving off AD FS federation or consolidating sync servers, reduces the infrastructure that this kind of metadata points attackers toward.

CSA Resource Alignment

CSA's *Zero Trust Principles and Guidance for Identity and Access Management (IAM)* [12] is the most directly applicable framework. It sets out risk-based, contextual authentication and authorization decisions that do not depend on the secrecy of identity attributes. The Connect Health exposure shows why that principle matters: once administrator identities and tenant configuration are known, continuous verification and phishing-resistant authentication become the primary controls standing between reconnaissance and compromise. The guidance's emphasis on strong authentication for privileged identities and on contextual policy enforcement corresponds to the Conditional Access authentication-strength and device-code-blocking measures recommended above.

CSA's research note on the Forg365 phishing service [14] documents attack techniques that the exposed data would make easier to target. It covers AI-generated lures, adversary-in-the-middle session theft, and OAuth device-code phishing aimed at Microsoft 365 users. Defenders at affected tenants should treat the detection and mitigation guidance in that note as directly applicable to the administrators named in the Connect Health exposure. CSA's analysis of the Entra Agent ID Administrator service principal takeover [15] provides context on the recurring pattern of Entra service-plane authorization flaws discussed in this note. It also covers the importance of monitoring credential additions to applications and service principals, a detection that applies equally to the Entra Connect Sync application.

At the control level, the CSA AI Controls Matrix (AICM) v1.1 [13] provides relevant requirements in its Identity and Access Management domain, covering strong authentication, privileged access management, and least privilege. Its Supply Chain Management, Transparency, and Accountability domain is also relevant, covering provider incident notification and transparency expectations. Organizations using AICM to evaluate cloud identity providers can use this incident as a test case for their supplier assessment criteria on detection timelines and disclosure scope.

References

[1] Abrams, Lawrence. "Microsoft Entra Connect Health flaw exploited to steal tenant data." BleepingComputer, November 21, 2026.

[2] Paganini, Pierluigi. "Microsoft Entra Connect Health Breach Exposes Tenant Metadata and Admin Details." Security Affairs, November 23, 2026.

[3] Noureen, Rabia. "Security Flaw Exposed Microsoft Entra Connect Health Tenant Metadata." Petri, November 23, 2026.

[4] Burnel, Florian. "Entra Connect Health Flaw Exposed Microsoft Tenant Admin Emails." IT-Connect, November 24, 2026.

[5] Zhang, Roche. "Microsoft Entra Connect Health Breach: Checks and Hunting." ThreatDefence, November 24, 2026.

[6] Microsoft. "Microsoft Entra Connect Health frequently asked questions." Microsoft Learn, updated June 2026.

[7] Microsoft. "Entra Connect Health PowerShell scripts." GitHub (AzureAD/EntraIDGovernance), November 2026.

[8] Secureworks Counter Threat Unit. "Azure Active Directory Exposes Internal Information." Secureworks, April 2022.

[9] Heinsen, Daniel. "Update: Dumping Entra Connect Sync Credentials." SpecterOps, June 9, 2025.

[10] Townsend, Kevin. "All Microsoft Entra Tenants Were Exposed to Silent Compromise via Invisible Actor Tokens: Researcher." SecurityWeek, September 23, 2025.

[11] Mollema, Dirk-jan. "One Token to rule them all - obtaining Global Admin in every Entra ID tenant via Actor tokens." dirkjanm.io, September 2025.

[12] Cloud Security Alliance. "Zero Trust Principles and Guidance for Identity and Access Management (IAM)." CSA Zero Trust Working Group, 2024.

[13] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." CSA, 2025.

[14] Cloud Security Alliance. "Forg365: AI Lure Generation in an M365 Phishing Kit." CSA AI Safety Initiative, July 10, 2026.

[15] Cloud Security Alliance. "Entra Agent ID Administrator Flaw: Service Principal Takeover." CSA AI Safety Initiative, April 28, 2026.