

Microsoft delivers comprehensive solution to battle rise in consent phishing emails

By Microsoft Threat Intelligence

Published: 2021-07-14 · Archived: 2026-04-06 01:49:37 UTC

Microsoft threat analysts are tracking a continued increase in consent phishing emails, also called [illicit consent grants](#), that abuse OAuth request links in an attempt to trick recipients into granting attacker-owned apps permissions to access sensitive data.

This blog offers a look into the current state of consent phishing emails as an initial attack vector and what security administrators can do to prevent, detect, and respond to these threats using advanced solutions like Microsoft Defender for Office 365. Consent phishing attacks aim to trick users into granting permissions to malicious cloud apps in order to gain access to user's legitimate cloud services. The consent screen displays all permissions the app receives; and because the cloud services are legitimate, unsuspecting users accept the terms or hit 'enter,' which grants the malicious app those requested permissions.

Consent phishing attacks are a specialized form of phishing, so they require a comprehensive, multi-layer defense. It's important for system administrators to gain visibility and control over apps and the permissions these apps have in their environment. User consent settings with consent policies in Azure Active Directory enable administrators to manage when end users can grant consent to apps. A new [app governance add-on feature in Microsoft Defender for Cloud Apps](#) (previously Microsoft Cloud App Security) provides organizations the visibility to enable them to quickly identify when an app exhibits anomalous behavior.

Microsoft has previously [warned against these application-based attacks](#) as many organizations shifted to remote work force at the onset of the COVID-19 pandemic. Microsoft's Digital Crimes Unit (DCU) has in the past also taken steps to [disrupt cybercriminal infrastructure](#) used for a particular consent phishing campaign.

The state of consent phishing attacks

Consent phishing attacks abuse legitimate cloud service providers, including Microsoft, Google, and Facebook, that use OAuth 2.0 authorization—a widely used industry protocol that allows third-party apps to access a user's account and perform actions on their behalf.

The goal of these attacks is to trick unsuspecting users into granting permissions (consent) to malicious attacker-owned applications. This is different from a typical credential harvesting attack, where an attacker looking to steal credentials would craft a convincing email, host a fake landing page, and expect users to fall for the lure. If the attempt is successful, user credentials are then passed on to the attacker.

In a consent phishing attack, the user sign-in takes place at a legitimate identity provider, rather than a fake sign-in page, in an attempt to trick users into granting permissions to malicious attacker-controlled applications. Attackers use the obtained access tokens to retrieve users' account data from the API resource, without any further action by

the user. Targeted users who grant the permissions allow attackers to make API calls on their behalf through the attacker-controlled app. Depending on the permissions granted, the access token can also be used to access other data, such as files, contacts, and other profile details.

Microsoft Defender for Office 365 data shows an increasing use of this technique in recent months.

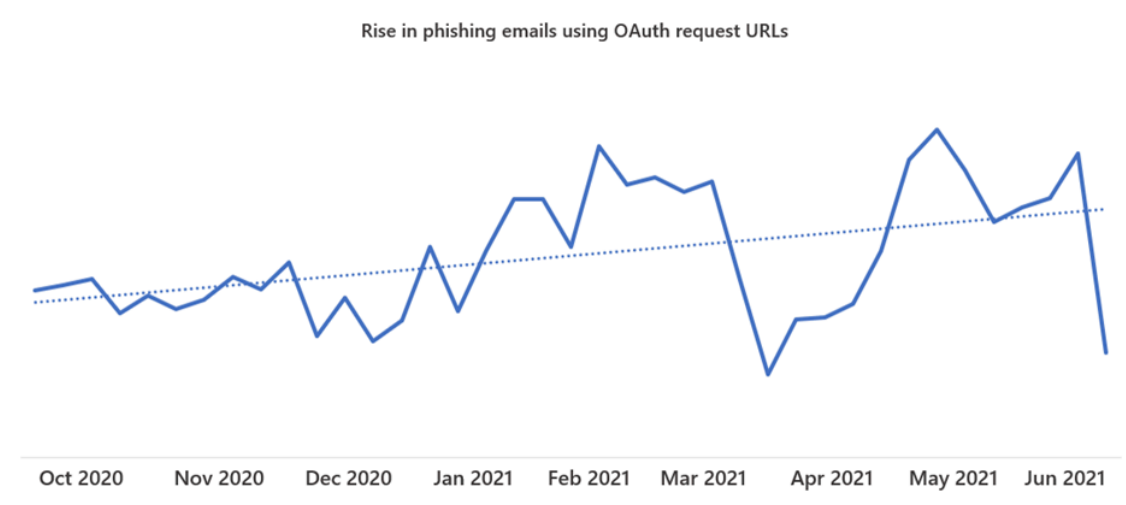


Figure 1. OAuth phishing URL trend from October 2020

In most cases, consent phishing attacks do not involve password theft, as access tokens don't require knowledge of the user's password, yet attackers are still able to steal confidential data and other sensitive information. Attackers can then maintain persistence in the target organization and perform reconnaissance to further compromise the network.

A typical consent phishing attack follows this attack chain:

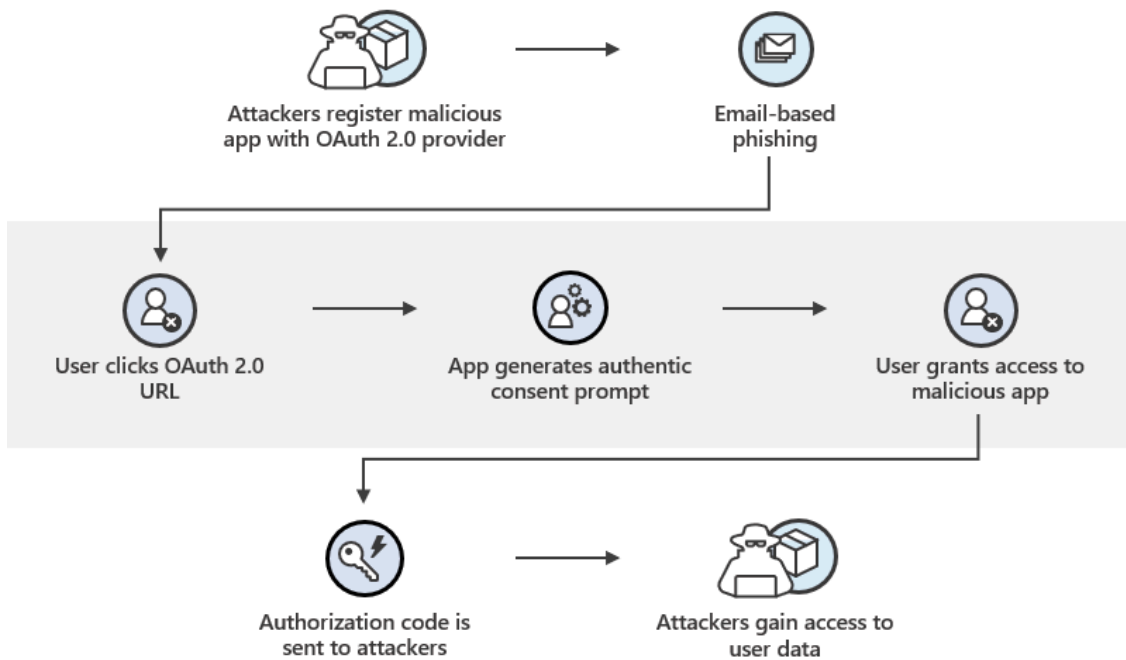


Figure 2. Consent phishing attack flow

Attackers typically configure apps so that they appear trustworthy, registering them using names like “Enable4Calc”, “SettingsEnabler”, or “Settings4Enabler,” which resemble legitimate business productivity app integrations. Attackers then distribute OAuth 2.0 URLs via conventional email-based phishing attacks, among other possible techniques.

Clicking the URL triggers an authentic consent prompt, asking users to grant the malicious app permissions. Other cloud providers, such as Google, Facebook, or Twitter, display consent prompts or dialog boxes that request for users’ permissions on behalf of third-party apps. The permissions requested vary depending on the app.

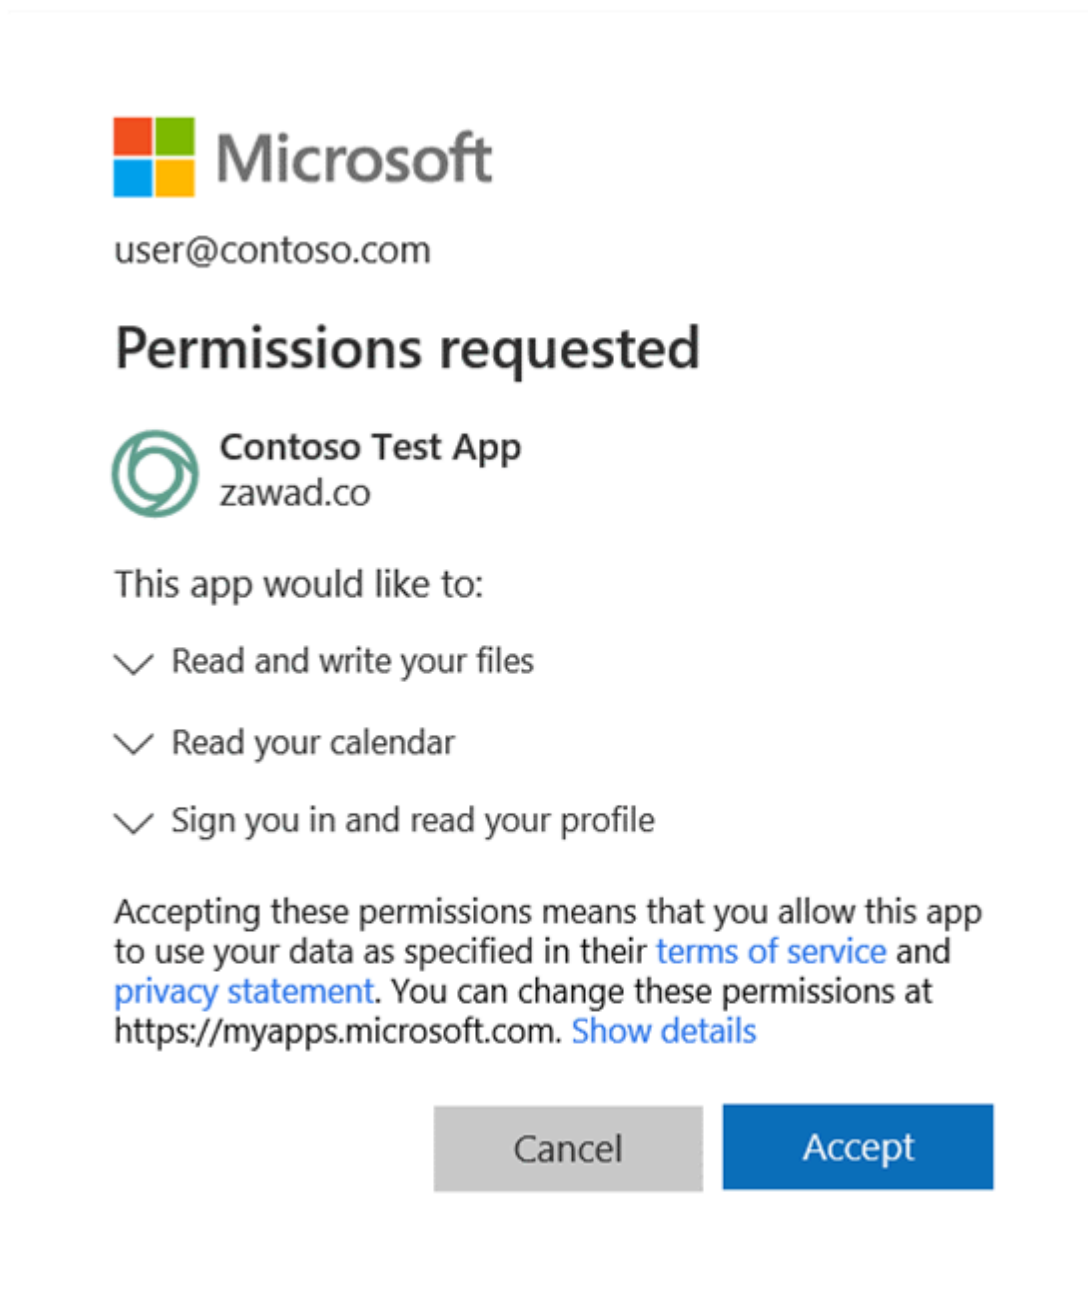


Figure 3. OAuth apps gain permission by displaying a “Permissions requested” dialog that shows what permissions the third-party is requesting

When users click “accept” or “allow”, the app obtains an authorization code that it redeems for an access token. This access token is then used to make API calls on behalf of the user, giving attackers access to the user’s email, forwarding rules, files, contacts, and other sensitive data and resources.

Consent phishing campaign: A case study

A recent consent phishing attack we tracked employed social engineering techniques to craft an email that impersonates a business growth solutions company. The message falsely claims to instruct users to review and sign a document, signaling a sense of urgency for the user—a tactic that is apparent in most phishing emails.

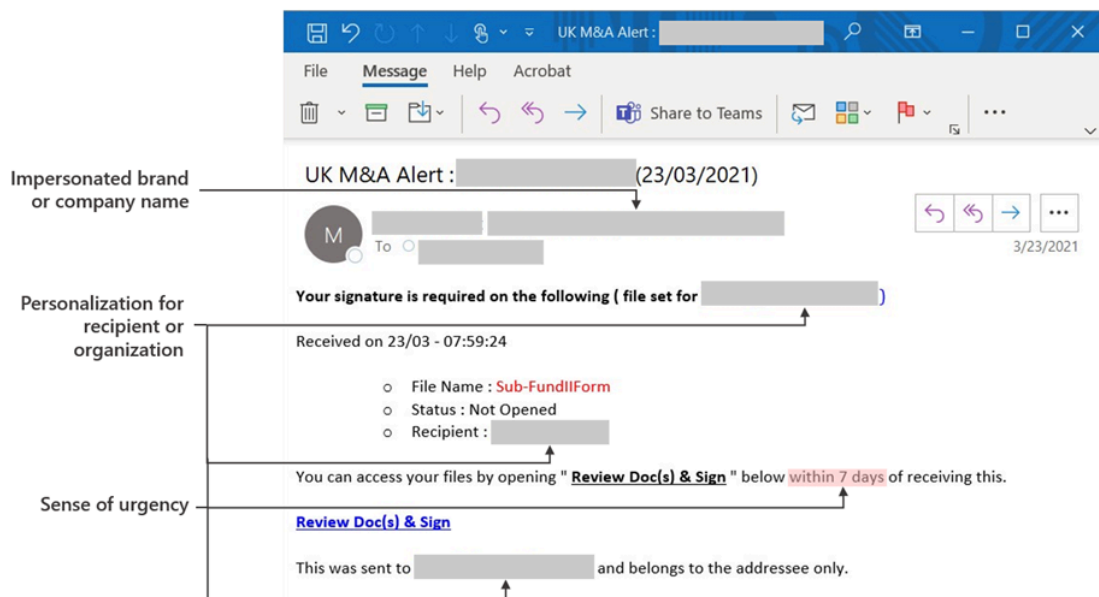


Figure 4. Sample email campaign with a Review Doc(s) & Sign link pointing to an OAuth URL

There are several phishing techniques in this email campaign: brand impersonation, personalized email text specific to the recipient or organization, and a recognizable sense of urgency as a social engineering lure.

What differentiates this attack from others is how the [OAuth](#) URL serves malicious content. To the email recipient, the “Review Doc(s) & Sign” OAuth URL appears legitimate, while URL is formatted with the identity provider URL as well.

The pattern we observed in this instance displays the the OAuth URL as “login.microsoftonline.com.” Other providers, such as Google, also format OAuth URLs in a similar manner.

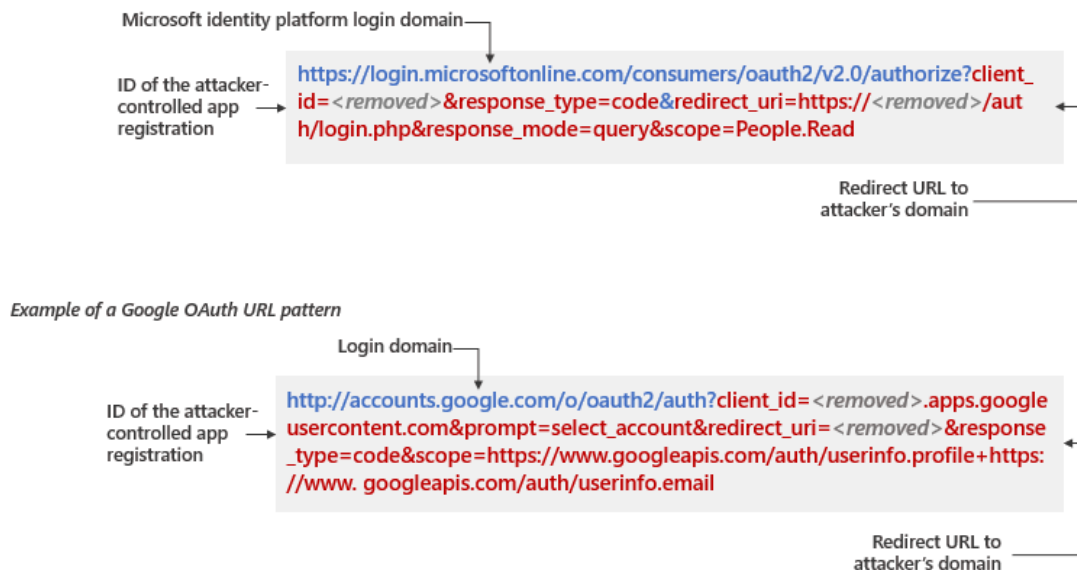


Figure 5. Observed patterns in OAuth URLs pointing to attacker's domain

Given the recent trend in OAuth abuse, we encourage organizations to look into and prevent this critical threat, beyond what traditional security measures offer.

How Microsoft delivers comprehensive, coordinated defense against consent phishing

The sophisticated and dynamic threat landscape exemplified by consent phishing attacks demonstrates the importance of employing a [Zero Trust security model](#) with a multi-layer defense architecture.

[Microsoft 365 Defender](#) provides comprehensive protection against consent phishing by coordinating defense across domains using multiple solutions: [Microsoft Defender for Office 365](#), [Microsoft Defender for Cloud Apps](#), and [Azure Active Directory](#).

The Microsoft identity platform helps prevent consent phishing in a few ways.

With risk-based step-up consent, Azure Active Directory (Azure AD) blocks end users from being able to grant consent to apps that are considered potentially risky. For example, a newly-registered multi-tenant app that has not been publisher-verified might be considered risky, and end users would not be allowed to grant consent, even if they visit the OAuth phishing URL.

Azure AD puts admins in control over when users are allowed to grant consent to apps. This is a powerful mechanism for preventing the threat in the first place, and Microsoft recommends that organizations review [settings for when users can grant consent](#). Microsoft recommends choosing the out-of-the-box option where users are only allowed to consent to apps from verified publishers, and only for chosen, lower risk permissions. For additional granularity, admins can also [create custom consent policies](#), which dictate the conditions for allowing users to grant consent, including for specific apps, publishers, or permissions.

Blocking consent phishing emails with Microsoft Defender for Office 365

Microsoft Defender for Office 365 uses advanced filtering technologies backed by machine learning, IP and URL reputation systems, and unparalleled breadth of signals to provide durable protection against phishing and other malicious emails, helping to block consent phishing campaigns out of the gate. Anti-phishing policies in Defender for Office 365 help protect organizations against [impersonation-based phishing attacks](#).

Microsoft researchers are constantly tracking OAuth 2.0 URL techniques and use this knowledge to provide feedback to email filtering systems. This helps ensure that Microsoft Defender for Office 365 is providing protection against the latest OAuth phishing attacks and other threats. Signals from Microsoft Defender Office 365 helps identify malicious apps and prevent users from accessing them, and provides rich threat data that organizations can query and investigate using [advanced hunting](#) capabilities.

Identifying malicious apps with Microsoft Defender for Cloud Apps

Microsoft Defender for Cloud Apps policies such as [activity policies](#), [anomaly detection](#), and [OAuth app policies](#) help organizations manage apps connected to their environment. The new [app governance add-on feature to Microsoft Defender for Cloud Apps](#) helps organizations:

- Define appropriate Microsoft 365 app behavior with data, users, and other apps
- Quickly detect unusual app behavior activity that varies from the baseline, and
- Disable an app when it behaves differently than expected

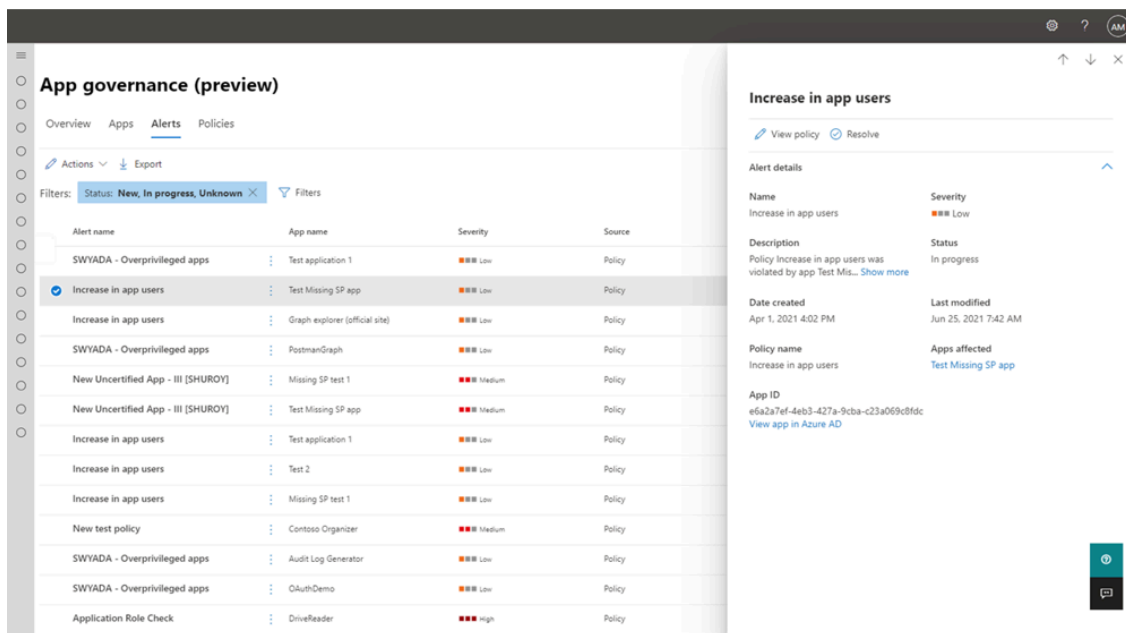


Figure 6. App governance in Microsoft 365 Compliance

To give organizations and users confidence in using apps in the Microsoft 365 ecosystem, the [Microsoft 365 App Compliance Program](#) enables app developers to establish authenticity of their applications. The program includes [publisher verification](#), publisher attestation, and Microsoft 365 certification.

Investigating and hunting for consent phishing attacks

Security operations teams can use [advanced hunting](#) capabilities in Microsoft 365 Defender to locate consent phishing emails and other threats. Microsoft 365 Defender consolidates and correlates email threat data from Microsoft Defender for Office 365, app signals from Microsoft Defender for Cloud Apps, and intelligence from other Microsoft services to provide a comprehensive end-to-end view of attacks. Security operations teams can then use the rich tools in Microsoft 365 Defender to investigate and remediate attacks.

OAuth URL pattern redirects to domain with unusual TLD

The consent phishing campaigns we described in this blog used a variety of unusual TLDs for communication with the attacker infrastructure. Use query below to find inbound emails with suspicious OAuth patterns. The suggested TLDs are based on our investigations. Security teams can modify the TLDs to expand the search. [Run query in Microsoft 365 Defender](#).

Best practices for protecting organizations against consent phishing

In addition to taking full advantage of the tools available to them in Microsoft 365 and Microsoft Azure, administrators can further strengthen defenses against consent phishing by following these measures:

- [Configure user consent settings](#) to only allow user consent for apps from verified publishers, for specific low-risk permissions.
- Increase end user awareness on consent phishing tactics as part of security training. Training should include checking for poor spelling and grammar in phishing mails or the application's consent screen as well as spoofed app names and domain URLs that are made to appear to come from legitimate applications or companies.
- Educate the organization on how permissions and consent frameworks work. Understand the data and permissions an application is asking for and understand how [permissions and consent work within our platform](#). Ensure administrators know how [to manage and evaluate consent requests](#) and [investigate and remediate risky OAuth applications](#).
- [Audit apps and consented permissions](#) in your organization to ensure applications being used are accessing only the data they need and adhering to the principles of least privilege.
- [Create proactive app governance policies](#) to monitor third party app behavior on the Microsoft 365 platform since policy driven and machine-learning initiated remediations address app behaviors both for common and emerging threat scenarios.

Additional resources

App governance add-on feature for Microsoft Defender for Cloud Apps is initially available as a public preview to existing Microsoft Defender for Cloud Apps customers in North America and Europe with other regions being added gradually the next few months.

To get started with app governance, visit our [quick start guide](#). To learn more about app governance, visit our [documentation](#). To launch app governance portal in Microsoft 365 Compliance center, go to <https://aka.ms/appgovernance>.

Refer to our documentation for reference on [configuring and managing user consent and app permissions](#) in Azure AD. For more information on Microsoft Defender for Cloud Apps refer to our [blog](#) and [Microsoft Defender for Cloud Apps explainer video](#).

Source: <https://www.microsoft.com/security/blog/2021/07/14/microsoft-delivers-comprehensive-solution-to-battle-rise-in-consent-phishing-emails/>