Detecting Post-Compromise Threat Activity in Microsoft Cloud Environments

us-cert.cisa.gov/ncas/alerts/aa21-008a

Summary

This Advisory uses the MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK®) framework. See the ATT&CK for Enterprise for all referenced threat actor tactics and techniques.

Updated April 15, 2021: The U.S. Government attributes this activity to the Russian Foreign Intelligence Service (SVR). Additional information may be found in a statement from the White House. For more information on SolarWinds-related activity, go to https://us-cert.cisa.gov/remediating-apt-compromised-networks and https://www.cisa.gov/supply-chain-compromise.

This Alert is a companion alert to AA20-352A: Advanced Persistent Threat Compromise of Government Agencies, Critical Infrastructure, and Private Sector Organizations. AA20-352A primarily focuses on an advanced persistent threat (APT) actor's compromise of SolarWinds Orion products as an initial access vector into networks of U.S. Government agencies, critical infrastructure entities, and private network organizations. As noted in AA20-352A, the Cybersecurity and Infrastructure Security Agency (CISA) has evidence of initial access vectors in addition to the compromised SolarWinds Orion products.

This Alert also addresses activity—irrespective of the initial access vector leveraged—that CISA attributes to an APT actor. Specifically, CISA has seen an APT actor using compromised applications in a victim's Microsoft 365 (M365)/Azure environment. CISA has also seen this APT actor utilizing additional credentials and Application Programming Interface (API) access to cloud resources of private and public sector organizations. These tactics, techniques, and procedures (TTPs) feature three key components:

- Compromising or bypassing federated identity solutions;
- Using forged authentication tokens to move laterally to Microsoft cloud environments; and
- Using privileged access to a victim's cloud environment to establish difficult-to-detect persistence mechanisms for Application Programming Interface (API)-based access.

This Alert describes these TTPs and offers an overview of, and guidance on, available open-source tools—including a CISA-developed tool, Sparrow—for network defenders to analyze their Microsoft Azure Active Directory (AD), Office 365 (O365), and M365 environments to detect potentially malicious activity.

Note: this Alert describes artifacts—presented by these attacks—from which CISA has identified detectable evidence of the threat actor's initial objectives. CISA continues to analyze the threat actor's follow-on objectives.

Technical Details

Frequently, CISA has observed the APT actor gaining *Initial Access* [TA0001] to victims' enterprise networks via compromised SolarWinds Orion products (e.g., Solorigate, Sunburst).[1] However, CISA is investigating instances in which the threat actor may have obtained initial access by *Password Guessing* [T1110.001], *Password Spraying* [T1110.003], and/or exploiting inappropriately secured administrative or service credentials (*Unsecured Credentials* [T1552]) instead of utilizing the compromised SolarWinds Orion products.

CISA observed this threat actor moving from user context to administrator rights for *Privilege Escalation* [TA0004] within a compromised network and using native Windows tools and techniques, such as Windows Management Instrumentation (WMI), to enumerate the Microsoft Active Directory Federated Services (ADFS) certificate-signing capability. This enumeration allows threat actors to forge authentication tokens (OAuth) to issue claims to service providers—without having those claims checked against the identity provider—and then to move laterally to Microsoft Cloud environments (*Lateral Movement* [TA0008]).

The threat actor has also used on-premises access to manipulate and bypass identity controls and multi-factor authentication. This activity demonstrates how sophisticated adversaries can use credentials from one portion of an organization to move laterally (*Lateral Movement* [TA0008]) through trust boundaries, evade defenses and detection (*Defense Evasion* [TA0005]), and steal sensitive data (*Collection* [TA0009]).

This level of compromise is challenging to remediate and requires a rigorous multi-disciplinary effort to regain administrative control before recovering.

Mitigations

Detection

Guidance on identifying affected SolarWinds software is well documented.[2] However—once an organization identifies a compromise via SolarWinds Orion products or other threat actor TTPs—identifying follow-on activity for on-premises networks requires fine-tuned network and host-based forensics.

The nature of cloud forensics is unique due to the growing and rapidly evolving technology footprints of major vendors. Microsoft's O365 and M365 environments have built-in capabilities for detecting unusual activity. Microsoft also provides premium services (Advanced Threat Protection [ATP] and Azure Sentinel), which enable network defenders to investigate TTPs specific to the Solorigate activity.[3]

Detection Tools

CISA is providing examples of detection tools for informational purposes only. CISA does not endorse any commercial product or service, including any subjects of analysis. Any reference to specific commercial products, processes, or services does not constitute or imply their endorsement, recommendation, or favoring by CISA.

There are a number of open-source tools available to investigate adversary activity in Microsoft cloud environments and to detect unusual activity, service principals, and application activity.[4] Publicly available PowerShell tools that network defenders can use to investigate M365 and Microsoft Azure include:

- CISA's Sparrow,
- Open-source utility Hawk, and
- CrowdStrike's Azure Reporting Tool (CRT).

Additionally, Microsoft's Office 365 Management API and Graph API provide an open interface for ingesting telemetry and evaluating service configurations for signs of anomalous activity and intrusion.

Note: these open-source tools are highlighted and explained to assist with on-site investigation and remediation in cloud environments but are not all-encompassing. Open source tools can be complemented by services such as Azure Sentinel, a Microsoft premium service that provides comprehensive analysis tools, including custom detections for the activity indicated.

General Guidance on Using Detection Tools

1. Audit the creation and use of service principal credentials. Look for unusual application usage, such as use of dormant applications.
2. Audit the assignment of credentials to applications that allow non-interactive sign-in by the application. Look for unexpected trust relationships added to the Azure Active Directory.
3. Download the interactive sign-ins from the Azure admin portal or use the Microsoft Sentinel product. Review new token validation time periods with high values and investigate whether it was a legitimate change or an attempt to gain persistence by a threat actor.

Sparrow

CISA created Sparrow to help network defenders detect possible compromised accounts and applications in the Azure/M365 environment. The tool focuses on the narrow scope of user and application activity endemic to identity- and authentication-based attacks seen recently in multiple sectors. It is neither comprehensive nor exhaustive of available data. It is intended to narrow a larger set of available investigation modules and telemetry to those specific to recent attacks on federated identity sources and applications.

(Updated April 8, 2021): CISA has also created "Aviary," which is a companion Splunk dashboard that can assist in visualizing and reviewing the output of Sparrow. Network defenders can find Aviary on CISA's Sparrow GitHub page. CISA advises network defenders to perform the following actions to use Sparrow:

1. Use Sparrow to detect any recent domain authentication or federation modifications.
    1. Domain and federation modification operations are uncommon and should be investigated.

2. Examine logs for new and modified credentials applied to applications and service principals; delineate for the credential type. Sparrow can be used to detect the modification of service principals and application credentials.
    1. Create a timeline for all credential changes, focusing on recent wholesale changes.
    2. Review the “top actors” for activity in the environment and the number of credential modifications performed.
    3. Monitor changes in application and service principal credentials.
    4. Investigate any instances of excessive permissions being granted, including, but not limited to, Exchange Online, Microsoft Graph, and Azure AD Graph.
3. Use Sparrow to detect privilege escalation, such as adding a service principal, user, or group to a privileged role.
4. Use Sparrow to detect `OAuth` consent and users’ consent to applications, which is useful for interpreting changes in adversary TTPs.
5. Use Sparrow to identify anomalous Security Assertion Markup Language (SAML) token sign-ins by pivoting on the unified audit log UserAuthenticationValue of 16457, which is an indicator of how a SAML token was built and is a potential indicator for forged SAML tokens.
    1. Note that this TTP has not been the subject of significant published security research but may indicate an unusual usage of a token, such as guest access for external partners to M365 resources.
6. Review the PowerShell logs that Sparrow exports.
    1. Review PowerShell mailbox sign-ins and validate that the logins are legitimate actions.
    2. Review PowerShell usage for users with PowerShell in the environment.
7. Use Sparrow to check the Graph API application permissions of all service principals and applications in M365/Azure AD.
    1. Investigate unusual activity regarding Microsoft Graph API permissions (using either the legacy https://graph.windows.net/ or https://graph.microsoft.com). Graph is used frequently as part of these TTPs, often to access and manipulate mailbox resources.
8. Review Sparrow’s listed tenant’s Azure AD domains, to see if the domains have been modified.
9. For customers with G5 or E5 licensing levels, review MailItemsAccessed for insight into what application identification (ID) was used for accessing users’ mailboxes. Use Sparrow to query for a specific application ID using the app id investigation capability, which will check to see if it is accessing mail or file items.
    1. The MailItemsAccessed event provides audibility for mailbox data accessed via mail protocols or clients.
    2. By analyzing the MailItemsAccessed action, incident responders can determine which user mailbox items have been accessed and potentially exfiltrated by a threat actor. This event will be recorded even in some situations where the message was not necessarily read interactively (e.g., bind or sync).[5]
    3. The resulting suspicious application ID can provide incident responders with a pivot to detect other suspicious applications that require additional analysis.
    4. Check for changes to applications with regards to the accessing of resources such as mail or file items.

(Updated April 8, 2021): Aviary can be used to assist with performing the above tasks. To install Aviary, after running Sparrow:

1. Ingest comma separated values (CSV) output from the Sparrow PowerShell script into Splunk.
    1. Sparrow output will have the following default filenames, which should not be modified: `AppUpdate_Operations_Export.csv`, `AppRoleAssignment_Operations_Export.csv`, `Consent_Operations_Export.csv`, `Domain_List.csv`, `Domain_Operations_Export.csv`, `FileItems_Operations_Export.csv`, `MailItems_Operations_Export.csv`, `PSLogin_Operations_Export.csv`, `PSMailbox_Operations_Export.csv`, `SAMLToken_Operations_Export.csv`, `ServicePrincipal_Operations_Export.csv`
2. Copy and paste the contents of the .xml file (aviary.xml in the root directory) into a new dashboard.
3. Use the data selection filters to point to the indexed Sparrow data (see figure 1)

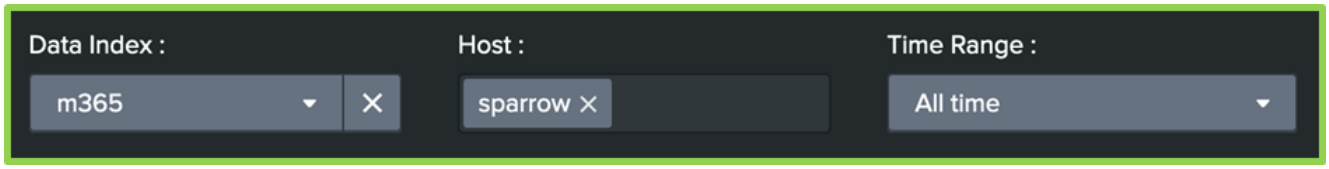

Figure 1: Data Selection Filters

Hawk

Hawk is an open-source, PowerShell-driven, community-developed tool network defenders can use to quickly and easily gather data from O365 and Azure for security investigations. Incident responders and network defenders can investigate specific user principals or the entire tenant. Data it provides include IP addresses and sign-in data. Additionally, Hawk can track IP usage for concurrent login situations.

Hawk users should review login details for administrator accounts and take the following steps.

CrowdStrike Azure Reporting Tool

CrowdStrike's Azure Reporting Tool (CRT) can help network defenders analyze their Microsoft Azure AD and M365 environment to help organizations analyze permissions in their Azure AD tenant and service configuration. This tool has minor overlap with Sparrow; it shows unique items, but it does not cover the same areas. CISA is highlighting this tool because it is one of the only free, open-source tools available to investigate this activity and could be used to complement Sparrow.

Detection Tool Distinctions

Detection Methods

Microsoft breaks the threat actor's recent activity into four primary stages, which are described below along with associated detection methods. Microsoft describes these stages as beginning with all activity after the compromise of the on-premises identity solution, such as ADFS.[6]

Note: this step provides an entry vector to cloud technology environments, and is unnecessary when the threat actor has compromised an identity solution or credential that allows the APT direct access to the cloud(e.g., without leveraging the SolarWinds Orion vulnerability).

Stage 1: Forging a trusted authentication token used to access resources that trust the on-premises identity provider

These attacks (often referred to as "Golden Security Assertion Markup Language" attacks) can be analyzed using a combination of cloud-based and standard on-premises techniques.[7] For example, network defenders can use `OAuth` claims for specific principals made at the Azure AD level and compare them to the on-premises identity.

Export sign-in logs from the Azure AD portal and look at the Authentication Method field.

Note: at portal.azure.com, click on a user and review the authentication details (e.g., date, method, result). Without Sentinel, this is the only way to get these logs, which are critical for this effort.

Detection Method 1: Correlating service provider login events with corresponding authentication events in Active Directory Federation Services (ADFS) and Domain Controllers

Using SAML single sign-on, search for any logins to service providers that do not have corresponding event IDs 4769, 1200, and 1202 in the domain.

Detection Method 2: Identifying certificate export events in ADFS

Look for:

Detection Method 3: Customizing SAML response to identify irregular access

This method serves as prevention for the future (and would only detect future, not past, activity), as it helps identify irregularities from the point of the change forward. Organizations can modify SAML responses to include custom elements for each service provider to monitor and detect any anomalous requests.[8]

Detection Method 4: Detecting malicious ADFS trust modification

A threat actor who gains administrative access to ADFS can add a new, trusted ADFS rather than extracting the certificate and private key as part of a standard Golden SAML attack.[9]
Network defenders should look for:

Stage 2: Using the forged authentication token to create configuration changes in the Service Provider, such as Azure AD (establishing a foothold)

After the threat actor has compromised the on-premises identity provider, they identify their next series of objectives by reviewing activity in the Microsoft Cloud activity space (Microsoft Azure and M365 tenants).

The threat actor uses the ability to forge authentication tokens to establish a presence in the cloud environment. The actor adds additional credentials to an existing service principal. Once the threat actor has impersonated a privileged Azure AD account, they are likely to further manipulate the Azure/M365 environment (action on objectives in the cloud).

Network defenders should take the following steps.

Stage 3: Acquiring an `OAuth` access token for the application using the forged credentials added to an existing application or service principal and calling APIs with the permissions assigned to that application

In some cases, the threat actor has been observed adding permissions to existing applications or service principals. Additionally the actor has been seen establishing new applications or service principals briefly and using them to add permissions to the existing applications or service principals, possibly to add a layer of indirection (e.g., using it to add a credential to another service principal, and then deleting it).[11]

Network defenders should use Sparrow to:

Stage 4: Once access has been established, the threat actor Uses Microsoft Graph API to conduct action on objectives from an external RESTful API (queries impersonating existing applications).

Network defenders should:

Microsoft Telemetry Nuances

The existing tools and techniques used to evaluate cloud-based telemetry sources present challenges not represented in traditional forensic techniques. Primarily, the amount of telemetry retention is far less than the traditional logging facilities of on-premises data sources. Threat actor activity that is more than 90 days old is unlikely to have been saved by traditional sources or be visible with the Microsoft M365 Management API or in the UAL.

Service principal logging is available using the Azure Portal via the "Service Principal Sign-ins" feature. Enable settings in the Azure Portal (see “Diagnostic Setting”) to ingest logs into Sentinel or a third-party security information and event management (SIEM) tool. An Azure Premium P1 or Premium P2 license is necessary to access this setting as well as other features, such as a log analytics workspace, storage account, or event hub.[12] These logs must be downloaded manually if not ingested by one of the methods listed in the Detection Methods section.

Global Administrator rights are often required by tools other than Hawk and Sparrow to evaluate M365 cloud security posture. Logging capability and visibility of data varies by licensing models and subscription to premium services, such as Microsoft Defender for O365 and Azure Sentinel.

According to CrowdStrike, "There was an inability to audit via API, and there is the requirement for global admin rights to view important information which we found to be excessive. Key information should be easily accessible."[13]

Documentation for specific event codes, such as UserAuthenticationMethod 16457, which may indicate a suspicious SAML token forgery, is no longer available in the M365 Unified Access Log. Auditing narratives on some events no longer exist as part of core Microsoft documentation sources.

The use of industry-standard SIEMs for log detection is crucial for providing historical context for threat hunting in Microsoft cloud environments. Standard G3/E3 licenses only provide 90 days of auditing; with the advanced auditing license that is provided with a G5/E5 license, audit logs can be extended to retain information for a year. CISA notes that this license change is proactive, rather than reactive: it allows enhanced visibility and features for telemetry from the moment of integration but does not provide retroactive visibility on previous events or historical context.

A properly configured SIEM can provide:

Built-in tools, such as Microsoft Cloud Services and M365 applications, provide much of the same visibility available from custom tools and are mapped to the MITRE ATT&CK framework and easy-to-understand dashboards.[14] However, these tools often do not have the ability to pull historical data older than seven days. Therefore, storage solutions that appropriately meet governance standards and usability metrics for analysts for the SIEM must be carefully planned and arranged.

1. Ingest comma separated values (CSV) output from the Sparrow PowerShell script into Splunk.
    1. Sparrow output will have the following default filenames, which should not be modified: `AppUpdate_Operations_Export.csv`, `AppRoleAssignment_Operations_Export.csv`, `Consent_Operations_Export.csv`, `Domain_List.csv`, `Domain_Operations_Export.csv`, `FileItems_Operations_Export.csv`, `MailItems_Operations_Export.csv`, `PSLogin_Operations_Export.csv`, `PSMailbox_Operations_Export.csv`, `SAMLToken_Operations_Export.csv`, `ServicePrincipal_Operations_Export.csv`
2. Copy and paste the contents of the .xml file (aviary.xml in the root directory) into a new dashboard.
3. Use the data selection filters to point to the indexed Sparrow data (see figure 1)

4.
    1. Investigate high-value administrative accounts to detect anomalous or unusual activity (Global Admins).
    2. Enable PowerShell logging, and evaluate PowerShell activity in the environment not used for traditional or expected purposes.
        1. PowerShell logging does not reveal the exact `cmdlet` that was run on the tenant.
    3. Look for users with unusual sign-in locations, dates, and times.
    4. Check permissions of service principals and applications in M365/Azure AD.
    5. Detect the frequency of resource access from unusual places. Use the tool to pivot to a trusted application and see if it is accessing mail or file items.
    6. Review mailbox rules and recent mailbox rule changes.
    - Sparrow differs from CRT by looking for specific indicators of compromise associated with the recent attacks.
    - CRT focuses on the tenant's Azure AD permissions and Exchange Online configuration settings instead of the unified audit log, which gives it a different output from Sparrow or Hawk.
    - CRT returns the same broad scope of application/delegated permissions for service principals and applications as Hawk.
    - As part of its investigation, Sparrow homes in on a narrow set of application permissions given to the Graph API, which is common to the recent attacks.
    - CRT looks at Exchange Online federation configuration and federation trust, while Sparrow focuses on listing Azure AD domains.
    - Among the items network defenders can use CRT to review are delegated permissions and application permissions, federation configurations, federation trusts, mail forwarding rules, service principals, and objects with KeyCredentials.
    1. The IP address and Activity_ID in EventCode 410 and the Activity_ID and Instance_ID in EventCode 500.
    2. Export-PfxCertificate or certutil-exportPFX in Event IDs 4103 and 4104, which may include detection of a certificate extraction technique.
    3. Deleted certificate extraction with ADFSdump performed using Sysmon Event ID 18 with the pipe name \microsoft##wid\tsql\query (exclude processes regularly making this pipe connection on the machine).
    4. Event ID 307 (The Federation Service configuration was changed), which can be correlated to relevant Event ID 510 with the same instance ID for change details (Event ID 510 with the same Instance ID could be more than one event per single Event ID 307 event).
    5. Event ID 307 (The Federation Service configuration was changed), which can be correlated to relevant Event ID 510 with the same Instance ID for change details. (Event ID 510 with the same Instance ID could be more than one event per single Event ID 307 event.)
        1. Review events, particularly searching for Configuration: Type: IssuanceAuthority where Property Value references an unfamiliar domain.

6. Possible activity of an interrogating ADFS host by using ADFS PowerShell plugins. Look for changes in the federation trust environment that would indicate new ADFS sources.
7. Audit the creation and use of service principal and application credentials. Sparrow will detect modifications to these credentials.
    1. Look for unusual application usage, such as dormant or forgotten applications being used again.
    2. Audit the assignment of credentials to applications that allow non-interactive sign-in by the application.
8. Look for unexpected trust relationships that have been added to Azure AD. (Download the last 30 days of non-interactive sign-ins from the Azure portal or use Azure Sentinel.).[10]
9. Use Hawk (and any sub-modules available) to run an investigation on a specific user. Hawk will provide IP addresses, sign-in data, and other data. Hawk can also track IP usage in concurrent login situations.
10. Review login details for administrator accounts (e.g., high-value administrative accounts, such as Global Admins). Look for unusual sign-in locations, dates, and times.
11. Review new token validation time periods with high values and investigate whether the changes are legitimate or a threat actor's attempts to gain persistence.
12. Examine highly privileged accounts; specifically using sign-in logs, look for unusual sign-in locations, dates, and times.
13. Create a timeline for all credential changes.
14. Monitor changes in application credentials (the script will export into csv named AppUpdate_Operations_Export).
15. Detect service principal credentials change and service principal change (e.g., if an actor adds new permissions or expands existing permissions).
    1. Export and view this activity via the ServicePrincipal_Operations_Export.
16. Record `OAuth` consent and consent to applications
    1. Export and view this record via the Consent_Operations_Export file.
17. Investigate instances of excessive high permissions, including, but not limited to Exchange Online, Microsoft Graph, and Azure AD Graph.
    1. Review Microsoft Graph API permissions granted to service principals.
    2. Export and view this activity via the ApplicationGraphPermissions csv file.
        1. **Note:** Hawk can also return the full list of service principal permissions for further investigation.
    3. Review top actors and the amount of credential modifications performed.
    4. Monitor changes in application credentials.
18. Identify manipulation of custom or third-party applications.
    1. Network defenders should review the catalog of custom or third-party vendors with applications in the Microsoft tenant and perform the above interrogation principles on those applications and trusts.

19. Review modifications to federation trust settings.
    1. Review new token validation time periods with high values and investigate whether this was a legitimate change or an attempt to gain persistence by the threat actor.
        1. The script detects the escalation of privileges, including the addition of Service Principals (SP) to privileged roles. Export this data into csv called AppRoleAssignment_Operations_Export.
20. In MailItemsAccessed operations, found within the Unified Audit Log (UAL), review the application ID used (requires G5 or E5 license for this specific detail).
21. Query the specific application ID, using the Sparrow script's app ID investigation capability to interrogate mail and file items accessed for that applicationID (Use the application ID utility for any other suspicious apps that require additional analysis.).
22. Check the permissions of an application in M365/Azure AD using Sparrow.
    1. Hawk will return Azure_Application_Audit, and Sparrow will return ApplicationGraphPermissions.
    2. Network defenders will see the IP address that Graph API uses.
    3. Note: the Microsoft IP address may not show up as a virtual private server/anonymized endpoint.
23. Investigate a specific service principal, if it is a user-specific user account, in Hawk. This activity is challenging to see without Azure Sentinel or manually downloading and reviewing logs from the sign-in portal.
24. Longer term storage of log data.
25. Cross correlation of log data with endpoint data and network data (such as those produced by ADFS servers), endpoint detection and response data, and identity provider information.
26. Ability to query use of application connectors in Azure.

Contact Information

CISA encourages recipients of this report to contribute any additional information that they may have related to this threat. For any questions related to this report, please contact CISA at

- 1-888-282-0870 (From outside the United States: +1-703-235-8832)
- central@cisa.dhs.gov (UNCLASS)
- us-cert@dhs.sgov.gov (SIPRNET)
- us-cert@dhs.ic.gov (JWICS)

CISA encourages you to report any suspicious activity, including cybersecurity incidents, possible malicious code, software vulnerabilities, and phishing-related scams. Reporting forms can be found on the CISA/US-CERT homepage at http://www.us-cert.cisa.gov/.

Resources

Azure Active Directory Workbook to Assess Solorigate Risk: https://techcommunity.microsoft.com/t5/azure-active-directory-identity/azure-ad-workbook-to-help-you-assess-solorigate-risk/ba-p/2010718

Volexity - Dark Halo Leverages SolarWinds Compromise to Breach Organizations: https://www.volexity.com/blog/2020/12/14/dark-halo-leverages-solarwinds-compromise-to-breach-organizations/

How to Find Activity with Sentinel: https://www.verboon.info/2020/10/monitoring-service-principal-sign-ins-with-azuread-and-azure-sentinel/

Third-Party Walkthrough of the Attack: https://dirkjanm.io/azure-ad-privilege-escalation-application-admin/

National Security Agency Advisory on Detecting Abuse of Authentication Mechanisms: https://media.defense.gov/2020/Dec/17/2002554125/-1/-1/0/AUTHENTICATION_MECHANISMS_CSA_U_OO_198854_20.PDF

Microsoft 365 App for Splunk: https://splunkbase.splunk.com/app/3786/

CISA Remediation Guidance: https://us-cert.cisa.gov/ncas/alerts/aa20-352a

References

Revisions

Initial version: January 8, 2021

February 4, 2021: Removed link and section for outdated product feedback form

April 8, 2021: Added Aviary Dashboard information

April 15, 2021: Added Attribution Statement

Please share your thoughts.

We recently updated our anonymous product survey; we'd welcome your feedback.