مجلس الأمن السيبراني
CYBER SECURITY COUNCIL
United Arab Emirates

Multiple Vulnerabilities in Azure Data Factory's Apache Airflow Integration

Tracking #:432316696

Date:02-01-2025

TLP: WHITE

EXECUTIVE SUMMARY:

The UAE Cyber Security Council has observed security researchers have uncovered several security vulnerabilities in Microsoft's Azure Data Factory Apache Airflow integration and other Azure services, including Azure Key Vault and Amazon Bedrock, which could expose sensitive data and infrastructure to attack.

TECHNICAL DETAILS:

Security researchers have uncovered several security vulnerabilities in Microsoft's Azure Data Factory Apache Airflow integration and other Azure services, including Azure Key Vault and Amazon Bedrock, which could expose sensitive data and infrastructure to attack.
These flaws, if exploited, could allow attackers to gain persistent access as shadow administrators over the entire Airflow Azure Kubernetes Service (AKS) cluster, potentially leading to data exfiltration, malware deployment, and unauthorized access to Azure-managed internal resources

The vulnerabilities include:

- Misconfigured Kubernetes RBAC in Airflow cluster
- Misconfigured secret handling of Azure's internal Geneva service
- Weak authentication for Geneva

Despite being classified as low severity by Microsoft, these vulnerabilities pose a significant risk to organizations using Azure Data Factory. Attackers could exploit these flaws to gain unauthorized administrative control over clusters, access Azure's internal services, and potentially manipulate critical logs and metrics.

Attack Vector

- The initial access technique involves crafting a malicious directed acyclic graph (DAG) file and uploading it to a private GitHub repository connected to the Airflow cluster, or altering an existing DAG file. The attacker's goal is to launch a reverse shell to an external server as soon as the DAG is imported.

To achieve this, the threat actor needs to:

- Gain write permissions to the storage account containing DAG files by:
- Utilizing a compromised service principal
- Using a shared access signature (SAS) token
- Breaking into a Git repository using leaked credentials1
- Upload the malicious DAG file, which automatically executes when imported by Airflow

Exploitation Process

- Initial Access: The malicious DAG file creates a reverse shell, granting the attacker access to an Airflow worker pod.
- Privilege Escalation: Although the initial shell has minimal permissions, a misconfigured service account with cluster-admin permissions is connected to the Airflow runner pod.
- Full Cluster Control: The attacker can download the Kubernetes command-line tool kubectl and take full control of the entire cluster by deploying a privileged pod and breaking out onto the underlying node.
- Deep Cloud Penetration: With root access to the host virtual machine (VM), the attacker can access Azure-managed internal resources, including the Geneva service.

- Log Manipulation: The attacker can potentially tamper with log data or send fake logs to Geneva, concealing their activities

Impact

If successfully exploited, these vulnerabilities could lead to:

- Unauthorized administrative control over Azure infrastructure
- Data exfiltration from connected storage accounts and databases
- Malware deployment within the Airflow cluster
- Manipulation of critical logs and metrics
- Creation of shadow workloads for cryptomining or further malicious activities
- Long-term persistent access through the creation of new service accounts

RECOMMENDATIONS:

- Implement Strict Access Controls: Carefully manage service permissions to prevent unauthorized access to DAG files and storage accounts.
- Monitor Third-Party Services: Regularly monitor the operations of critical third-party services integrated with your Azure environment.
- Secure Git Repositories: Implement strong access controls and regularly rotate credentials for Git repositories connected to Airflow clusters.
- Implement Least Privilege: Review and adjust Kubernetes RBAC configurations to ensure minimal necessary permissions for service accounts.
- Enhance Authentication: Strengthen authentication mechanisms, especially for internal services like Geneva.
- Regular Security Audits: Conduct frequent security assessments of your Azure Data Factory and Airflow configurations.
- Patch Management: Keep all systems and services up-to-date with the latest security patches

Kindly circulate this information to your subsidiaries and partners as well as share with us any relevant information and findings.

The UAE Cyber Security Council extends its appreciation for the continued collaboration.

REFERENCES:

- https://unit42.paloaltonetworks.com/azure-data-factory-apache-airflow-vulnerabilities/