[redacted]

2026Q2 Cybersecurity Risk Assessment Report

2026Q2 / 1

Project Objective

[redacted] contracted Abacus to perform a scoped cybersecurity risk assessment on corporate systems between the dates of April 06th, 2026, and April 27th, 2026.

Objectives of this Exercise:

- Identify and evaluate controls across three key areas:
  - Management Controls -- Management of the information technology security system and the management and acceptance of risk.
  - Operational Controls -- Security methods focusing on mechanisms implemented and executed primarily by people, including media safeguards and inventory controls.
  - Technical Controls -- Hardware and software controls provide automated protection to the system or applications.

- Collect, scan, and assess identity management, security groups/ACLs, isolation and segregation mechanisms, web application firewall configs, dependency versions, logging and accounting mechanisms, data-at-rest and data-in-transit cryptography, access control mechanisms, and more.

- Outline the potential impact and severity of each vulnerability, weakness, and risk, including observed deviations from standard information security best practices and principles.

- Present information that identifies the root causes behind the vulnerabilities, weaknesses, and risks identified in the assessment while conveying a message of how those problems relate to the inherent security posture of [redacted]

- Provide [redacted] with a full due diligence report, including information such as vulnerability/weakness/risk details, risk mitigation recommendations, assessment methods, information system statistics, and more.

Abacus performed security testing in accordance with NIST SP800-115 Technical Guide to Information Security Testing and Assessment guidelines. The findings in this report are limited to those within the scope of the Rules of Engagement (RoE). This report reflects the security position of [redacted] information systems, as seen during the dates in which the testing was conducted. Future changes to any applications or infrastructure will change the security position of the environment from its current state. Additionally, as time passes, new attacks may arise that were previously unknown to the security industry. Considering this, Abacus always recommends that all information systems undergo routine security testing within a reasonable timeframe.

Table of Contents

Cybersecurity Risk Assessment Methodology

Abacus' cybersecurity risk assessment methodology is modeled from the standards outlined in NIST SP 800-115, Technical Guide to Information Security Testing and Assessment, and CIS. The methodology includes four phases: Planning, Data Collection, Analysis, and Reporting. The process of this assessment aggregates and analyzes huge amounts of data and security findings across all analyzed technologies and services. The assessment methodology systematically organizes, prioritizes, and ranks only the most meaningfully impactful from the perspective of both the magnitude of impact as well as ease of exploitation.

Phase 1: Planning

The Planning phase involves determining the scope of the assessment, identifying the in-scope cloud platforms, and defining the objectives. This phase also outlines the assessment approach, defines the roles and responsibilities of the assessment team, and establishes the timeline for the assessment.

Phase 2: Data Collection

The Data Collection phase involves gathering information about the cloud platforms in scope, including their configuration, architecture, and security controls. This phase also includes identifying potential vulnerabilities, threats, and risks associated with the cloud platforms. Abacus utilizes several complementary techniques to ensure a holistic security picture is understood, including cloud-native features, open-source and commercial vulnerability assessment tools, and manual verification by expert security engineers.

Phase 3: Analysis

The Analysis phase involves reviewing and analyzing the data collected in the previous phase. This phase assesses the effectiveness of the security controls in place and identifies any gaps or weaknesses in the security posture of the cloud platforms. During all analyses specific care is taken to understand the context of the environment, the presence of any compensating controls, and the practical impact to the organization.

Phase 4: Reporting

The Reporting phase involves documenting the findings of the assessment. This phase includes creating a comprehensive report that summarizes the assessment results, identifies improvement areas, and provides recommendations for enhancing the security posture of the cloud platforms.

It is important to note that the security findings and risk ratings generated by automated tools are not necessarily accurate on their own, as automated tools lack the context concerning compensating controls or actual environmental impact. That said, the Security Findings section of this report includes security risks and risk mitigation recommendations with this context in mind and are ranked accordingly. Informational-level security findings are not featured in this report for the sake of brevity and prioritizing security finding analysis.

Risk Scoring

To present the findings in a manner easily digested and prioritized, Abacus uses the traditional aspects of Magnitude of Impact and Ease of Exploitation to determine the Level of Risk associated with individual findings identified in this security assessment.

Magnitude of Impact

The Magnitude of Impact rating is a measure of the damage each weakness could inflict on the organization and is based on access, data, etc., obtained through exploitation (i.e., level of access gained, types of information accessed, the strength of security measures bypassed)

Rating	Definition
HIGH	Exploitation could seriously affect system confidentiality, availability, integrity, or authentication. Resources to mitigate high risks should receive priority.
MODERATE	Exploitation could moderately affect system confidentiality, availability, integrity, or authentication.
LOW	Exploitation could minimally affect system confidentiality, availability, integrity, or authentication. The presence of multiple low impact findings could result in a finding with a higher impact rating.

Ease of Exploitation

The Ease of Exploitation rating measures the degree of difficulty for exploiting each finding. The higher the rating, the easier it is to exploit the finding and, therefore, the more likely it is to be exploited

Rating	Definition
HIGH	Exploitation is easily accomplished using freely and readily available tools, techniques, and exploits.
MODERATE	Exploitation is moderately accomplished using a combination of freely available and custom tools, techniques, and exploits.
LOW	Exploitation is difficult to accomplish and requires a significant investment of time to develop custom tools, techniques, and exploits.

Level of Risk

		MAGNITUDE OF IMPACT		
		HIGH	MODERATE	LOW
EASE OF EXPLOITATION	HIGH	HIGH	MODERATE	MODERATE
	MODERATE	MODERATE	MODERATE	LOW
	LOW	MODERATE	LOW	LOW

Security Findings that are noteworthy but do not have an associated risk are labeled as Informational. Informational Security Findings are not inherently usable as attack vectors but may have associated operational excellence recommendations for implementing additional technical safeguards and controls.

Risk Results Key

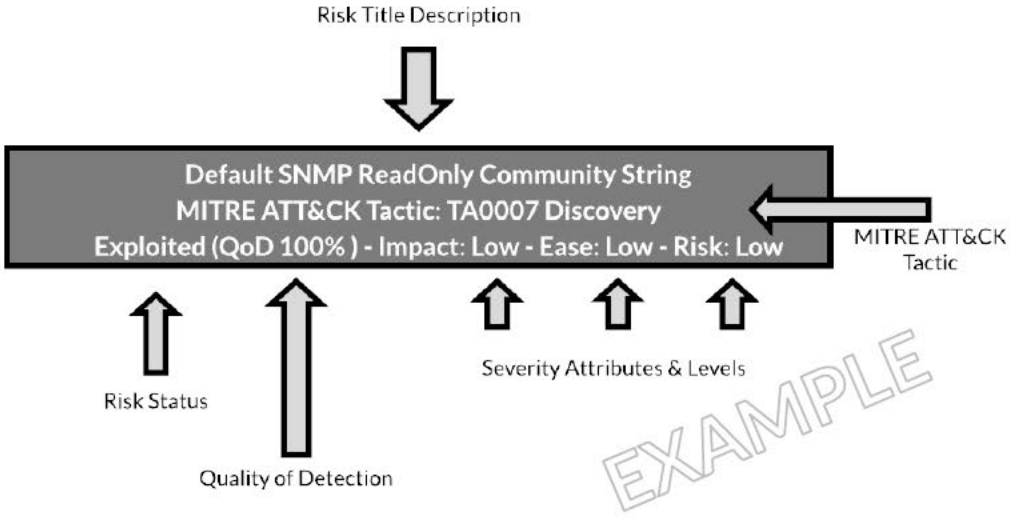

- **Risk Title Bar**
  - ***Risk Title Description***
    - Contains a brief description of the pertinent risk.
  - ***MITRE ATT&CK Tactic***
    - Defines the MITRE ATT&CK Tactic category to which the risk is most closely associated to.
  - ***Risk Status***
    - Specifies if the risk was detected or exploited. In many situations, a high-severity risk is detected but not exploited due to the potential service impact to a production system and subsequent violation of the project Rules of Engagement (RoE).
    - In some cases of a service-impactful risk, a proof-of-concept exploit can be safely performed.
  - ***QoD (Quality of Detection)***
    - For "*Detected*" but not exploited risks, QoD is a value between 0% and 100% (i.e., confidence) that ranks the reliability of the executed network vulnerability test or fingerprinted information system.
  - ***Impact***
    - The Impact rating is a measure of the damage each weakness could inflict on the organization and is based on access, data, etc., obtained through exploitation.
  - ***Ease***
    - The Ease rating measures the degree of difficulty for exploiting each finding. The higher the rating, the easier it is to exploit the finding and, therefore, the more likely it is to be exploited.
  - ***Risk***
    - The Risk rating is based largely upon the potential negative impact to the pertinent environment combined with the ease of exploitation as illustrated in the matrix on the previous page.

Insecure TLS Protocols and Cipher Suites Susceptible to BEAST Attack
MITRE ATT&CK Tactic: TA0001-Initial Access
Exploited – **Impact:** Moderate – **Ease:** Moderate – **Risk:** Moderate

Summary:
The affected endpoints supported outdated TLS versions 1.0 and 1.1, along with weak cipher suites with known attacks. This may allow a malicious actor to sniff webserver traffic and decrypt the content of an encrypted session, gaining access to potentially sensitive information. TLSv1.0 is considered obsolete, and TLSv1.1 is rapidly being retired from production. It was possible to leverage the Browser Exploit Against SSL/TLS (BEAST) attack to defeat TLS on the affected endpoints.

Risk Detection Result(s):

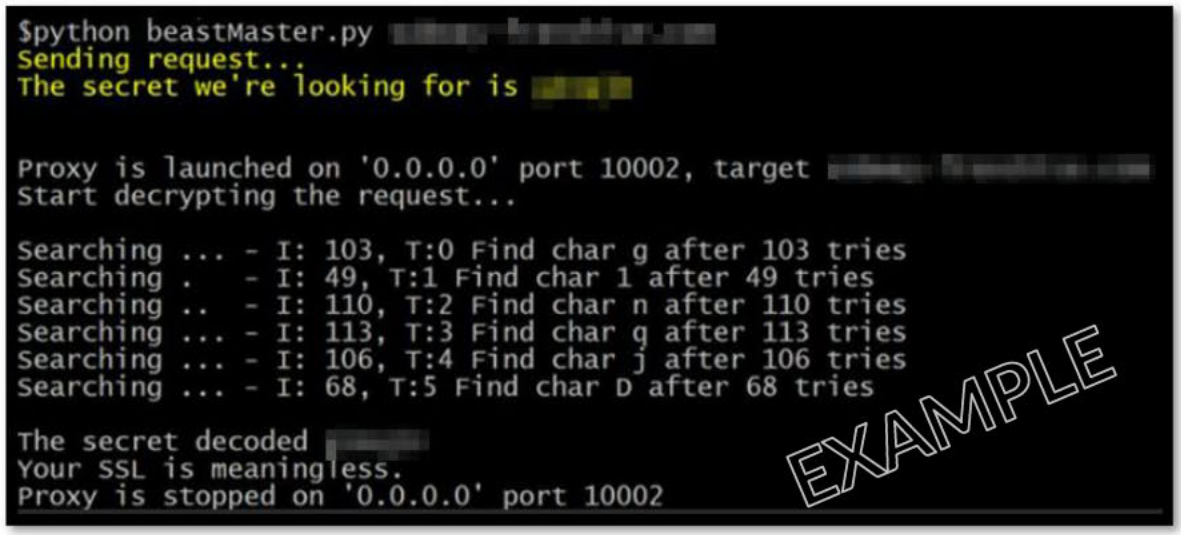

Risk Mitigation Recommendation(s):
- Update TLS configurations for modern or intermediate compatibility. Intermediate is the recommended configuration for the vast majority of services that do not need compatibility with legacy clients, such as Windows XP or old versions of OpenSSL, as it is highly secure and compatible with nearly every client released in the last five (or more) years.

- **Risk Body**
  - ***Summary***
    - Contains a high-level overview of the risk description.
  - ***Risk Detection Result(s)***
    - Demonstrates evidence that the risk exists, showcases exploitation if exploitation was performed.
  - ***Risk Mitigation Recommendation(s)***
    - Gives recommendations on how to best mitigate this risk.
  - ***Reference(s)***
    - Lists useful outside references that elaborate on the risk and/or provide additional mitigation recommendations.
  - ***Affected Endpoint(s)***
    - Lists every endpoint detected as being affected by this risk.

[redacted]

[redacted]

Organization	[redacted]		
Point of Contact Name	[redacted]	Security Testing Timeframe	April 6th, 2026 - April 27th, 2026
Point of Contact Phone	N/A	Point of Contact Email	[redacted]

Total Count of Security Findings:

2	9	2	4
High	Moderate	Low	Informational

Summary of Risk Results	
Risk Title Description	**Affected Endpoint(s)**
Print Spooler Service Enabled on Domain Controllers	[redacted]
Non-Default Principals With DCSync Rights	[redacted]
Lack of Conditional Access Policy to Block Legacy Authentication	Microsoft Entra ID
Insufficient Credential Management in AD	On-premises Active Directory
Insufficient Endpoint Hardening Controls	Windows Devices
Lack of Kerberos Security Hardening	[redacted]
Lack of Conditional Access Policy to Lock Down Global Admin Access	Microsoft Entra ID
Credentials Not Rotated for User, Admin and Privileged Accounts	Microsoft Entra ID
Privileged Group Membership Governance in AD	[redacted]
Lack of Data Loss Prevention Controls	Microsoft Purview
Insufficient Use of Password Manager	Organization-wide

Insecure Microsoft SharePoint Guest Access Policy	Microsoft SharePoint
Insecure Microsoft Teams External Access Policy	Microsoft Teams
GraphQL Query and Schema Exposure	[redacted]
Misconfigured SPF Mail Record	[redacted]
Previously Breached Credentials Identified	[redacted]
Inadequate Governance of Guest Accounts	Microsoft Entra ID
Activity Log Alerts Not Set Within Azure	Microsoft Azure

Summary of Provided System Access

Abacus Group was provided global read-only access to the following system(s) for inclusion in this risk assessment:

- Microsoft 365
- Microsoft Entra ID
- Vulnerability Management as a Service (VMaaS)

Abacus Group reviewed the following systems on a call with [redacted] PoC and [redacted] as direct access was not provisioned. The following systems are listed below:

- SentinelOne
- Active Directory & Entra ID via Purple Knight
- N-Able & Addigy RMM

Cybersecurity Risk Assessment Narrative

Microsoft 365

Abacus conducts risk assessments that analyze the critical security features of environments and compliance with regulatory requirements. As such, this engagement focused on assessing [redacted] systems in accordance with NIST SP800-115 Technical Guide to Information Security Testing and Assessment guidelines, and what is generally recommended by regulatory bodies such as the SEC, FCA, DORA, CSSF, ADGM, and DFSA. Abacus coordinated with the necessary personnel to obtain access to the environments and technical details for the systems examined.

Microsoft 365 environment assessment began with identifying users and associated devices, groups, roles, and services, such as logging facilities, SIEM, virtual network firewall coverage, virtual machines, Microsoft Defender, and multifactor authentication policies (MFA). Then, Abacus engineers holistically evaluated the environment's security posture to determine if any overlapping security controls were in place to identify potentially vulnerable configurations.

Abacus also verified [redacted] licensing level as Microsoft 365 E5. Please note that all recommendations related to Microsoft 365, including Microsoft Entra ID, within this report can be accomplished with the observed current licensing level.

Azure Information Protection Premium P1	1	14/15
Exchange Online (Plan 1)	0	3/3
Microsoft 365 Copilot	0	1/1
Microsoft Defender for Office 365 (Plan 1)	0	2/2
Microsoft Defender for Office 365 (Plan 2)	0	1/1
Microsoft Entra ID P2	0	1/1
Microsoft Power Automate Free	9991	9/10000
Microsoft Teams Audio Conferencing includes dial-out t...	1	0/1
Office 365 E5	0	15/15

Demonstrating M365 Licensing

Domain name ↑	Status
[redacted] (Default)	Healthy
[redacted]	Healthy
[redacted]	Incomplete setup

Validating the other domains within the suite

Microsoft Entra ID

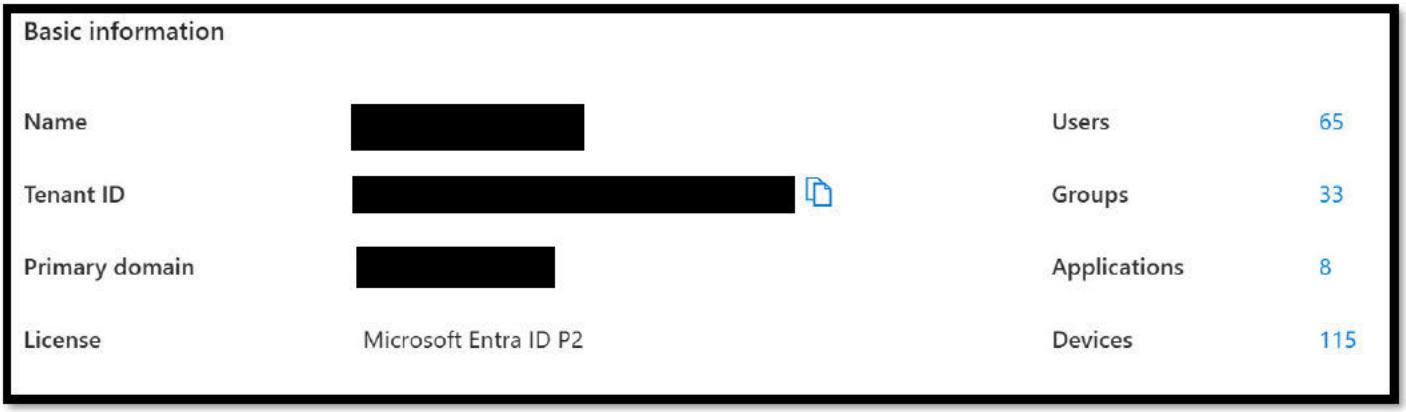

User and device information via Microsoft Entra ID showed 115 devices

Abacus assessed Microsoft Entra ID Conditional Access policies, which are used to enforce access controls through conditional "if-then" logic based on user, device, location, and risk signals. During review, it was identified that no Conditional Access policy is in place to block legacy authentication, and no policy specifically enforces additional restrictions or protections for Global Administrator accounts. Legacy authentication methods represent a well-known attack vector, as they do not support modern authentication controls and can bypass MFA protections even when MFA is otherwise enforced. The absence of a policy to explicitly block legacy authentication increases exposure to credential-based attacks such as password spraying and brute-force attempts, including from mobile and non-browser clients.

Implementing Conditional Access policies to block legacy authentication across all users and to apply stricter access controls to Global Administrator accounts would reduce the attack surface and better align the tenant with Microsoft security best practices.

Policy name	Created by	State	Creation date
Multifactor authentication and reauthentication for risky sign-ins	MICROSOFT	On	8/4/2025, 6:47:52 PM
Multifactor authentication for admins accessing Microsoft Admin Portal	MICROSOFT	Report-only	12/6/2023, 4:01:54 PM
AzureVPN and Apps - Require Hybrid Joined Machine	USER	Report-only	10/8/2024, 8:34:09 PM
Blocked Locations	USER	On	9/27/2023, 3:38:44 PM
Require Hybrid Azure AD joined device	USER	Report-only	1/16/2025, 10:01:12 AM
Require MFA	USER	On	8/8/2023, 8:24:01 PM

Demonstrating implemented conditional access policies

Abacus noted that a conditional access policy is enabled for the requirement of MFA. It is Abacus' belief that the one excluded account would be considered "break-glass" and is therefore an acceptable exclusion.

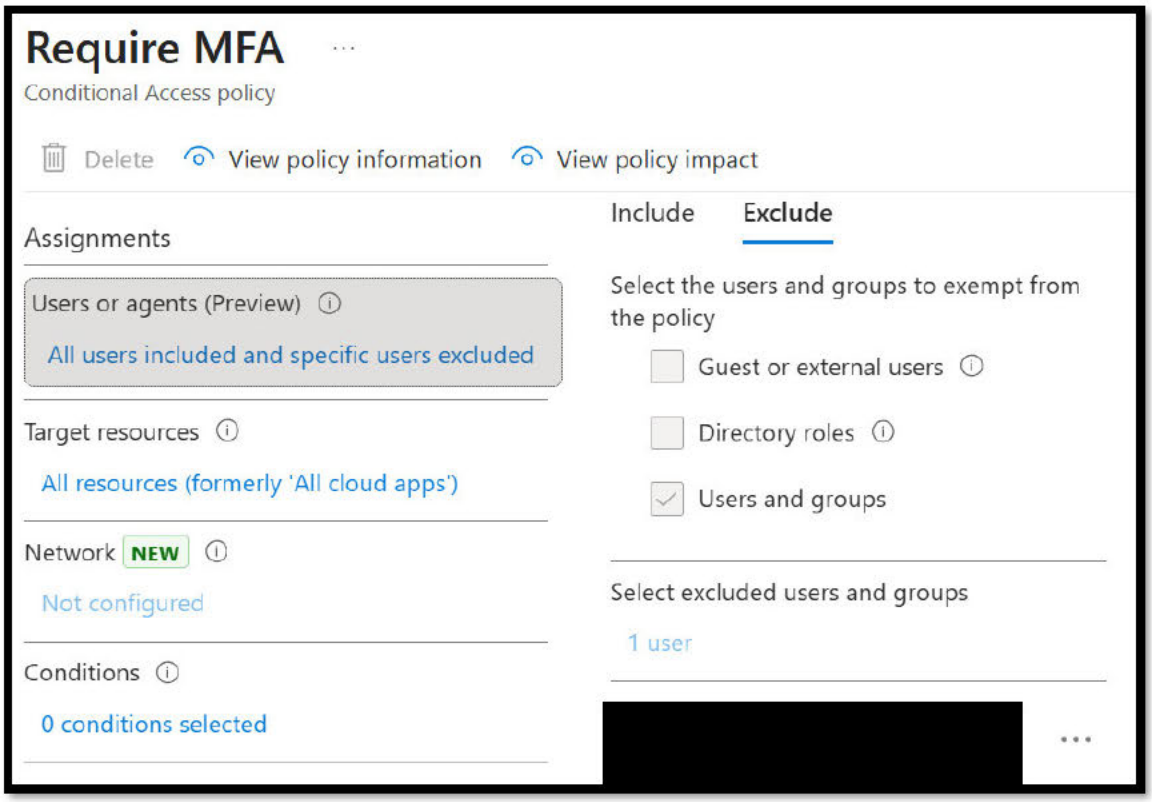

Identifying MFA registration is required for all users

Abacus noted that several conditional access policies were functionally contingent on users having compliant devices as defined by Intune compliance settings. This prompted Abacus to assess [redacted] environment thoroughly.

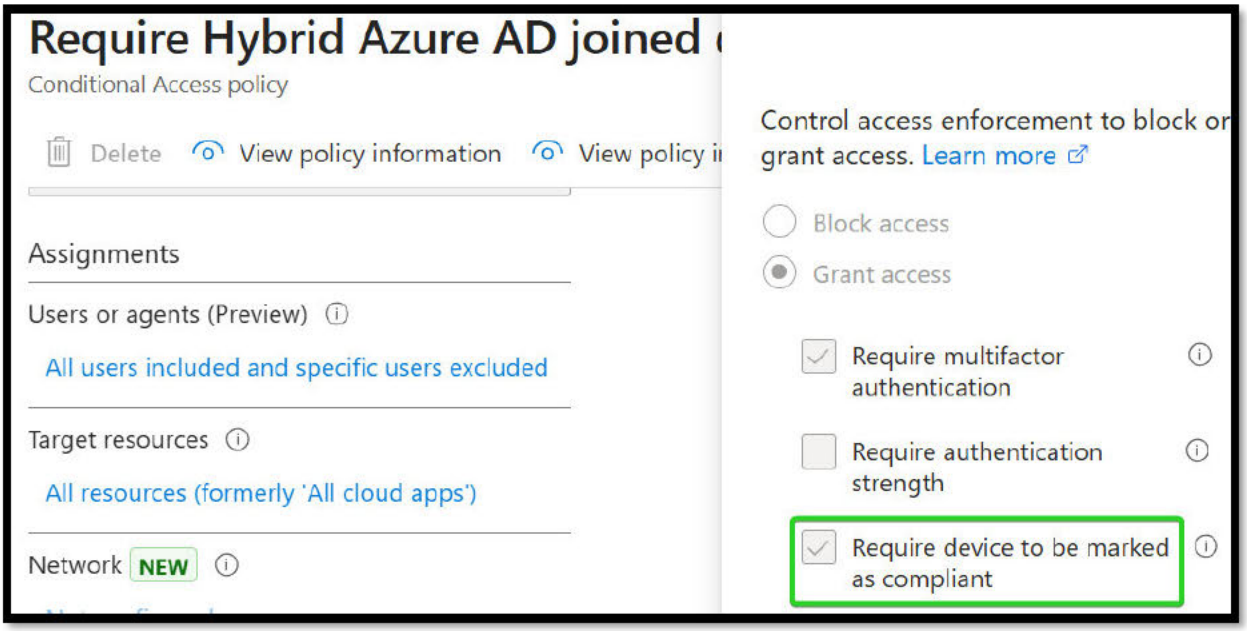

Demonstrating a conditional access policy contingent on compliant versus non-compliant Windows Devices

It was noted that users are not permitted to grant consent to unreliable applications. This restriction mitigates a popular method for securing persistence in compromised email accounts that involves fraudulently granting access to third-party applications.

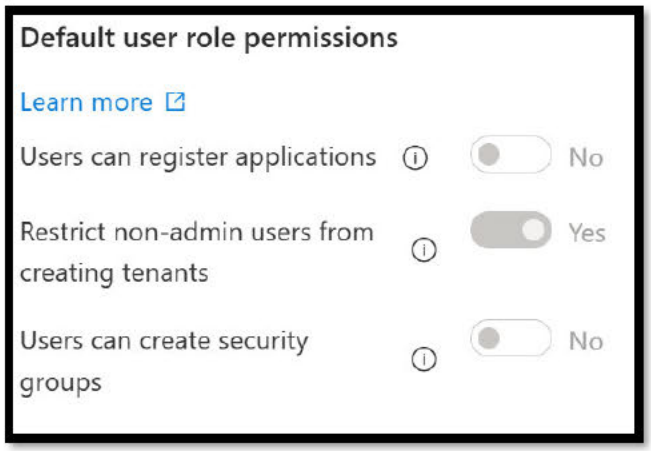

Demonstrating that users are not able to register applications

Demonstrating the guest user access settings

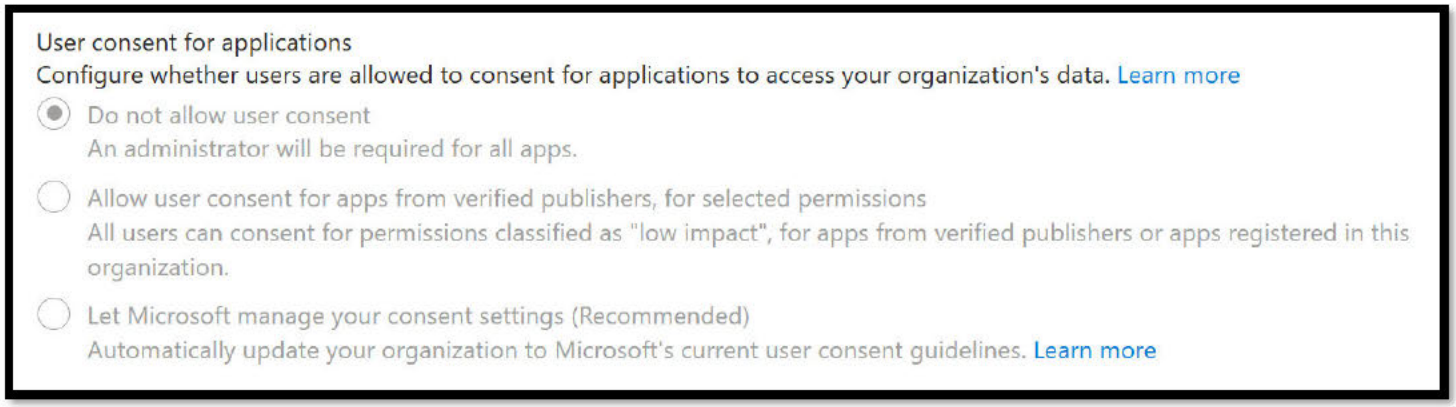

Demonstrating that the user consent for applications is set to "Do not allow user consent"

N- Able & Addigy RMM

While a Microsoft Intune environment is not in use, Abacus was able to review [redacted] device management architecture during a screen-sharing session with their [redacted] At present, all client endpoints are centrally managed through [redacted] remote monitoring and management (RMM) platform, N-able. Endpoint patching, monitoring, and policy enforcement are handled exclusively through this RMM, with no workloads managed directly by the client.

Endpoint provisioning is performed using an automated deployment platform, Emibot, which is used to install and provision new systems during device builds. Post-deployment software installation is largely performed manually, rather than through a centralized application catalog or MDM-driven software deployment workflow.

Apple devices are managed separately using Addigy. At the time of assessment, only two macOS devices were enrolled and managed through this platform.

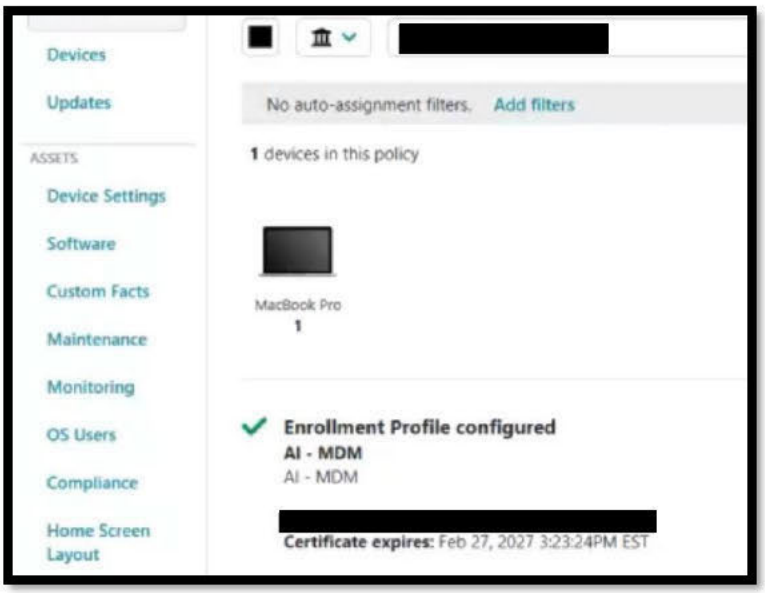

Identifying Addigy-managed devices

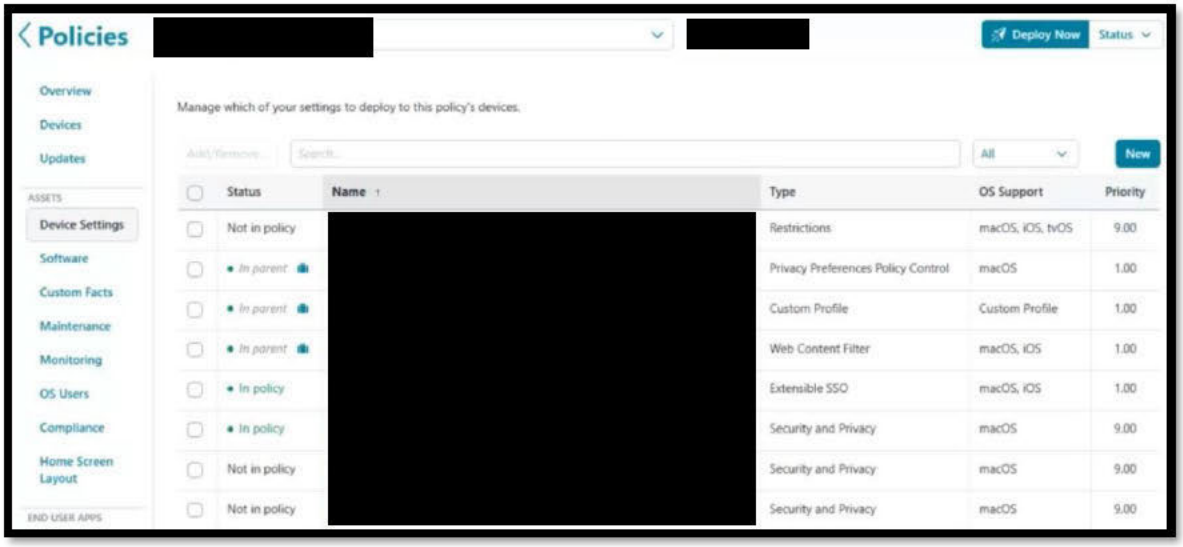

Demonstrating the Device settings within Policies

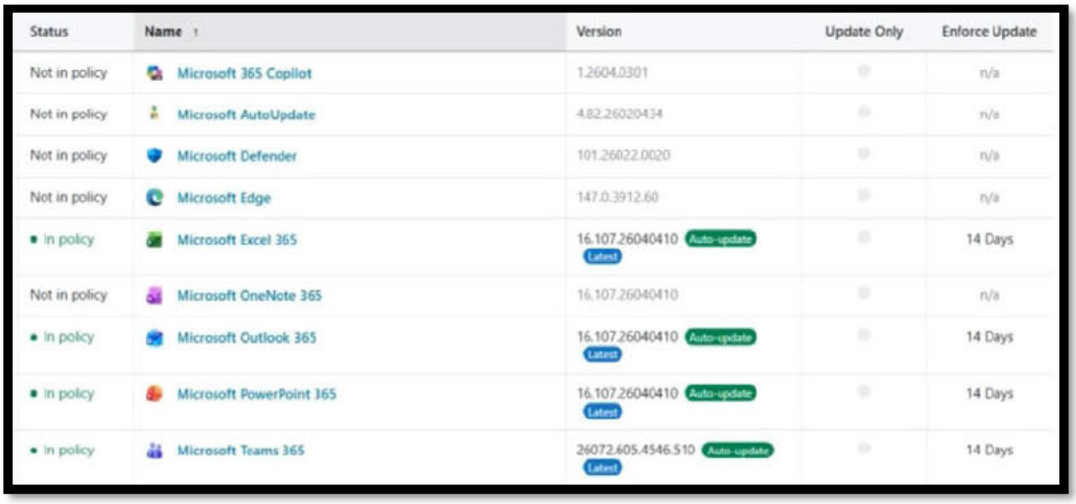

Status	Name	Version	Update Only	Enforce Update
Not in policy	Microsoft 365 Copilot	1.2604.0301		n/a
Not in policy	Microsoft AutoUpdate	4.82.26020434		n/a
Not in policy	Microsoft Defender	101.26022.0020		n/a
Not in policy	Microsoft Edge	147.0.3912.60		n/a
In policy	Microsoft Excel 365	16.107.26040410 Auto-update Latest		14 Days
Not in policy	Microsoft OneNote 365	16.107.26040410		n/a
In policy	Microsoft Outlook 365	16.107.26040410 Auto-update Latest		14 Days
In policy	Microsoft PowerPoint 365	16.107.26040410 Auto-update Latest		14 Days
In policy	Microsoft Teams 365	26072.605.4546.510 Auto-update Latest		14 Days

Demonstrating Software set to Auto-update within Addigy

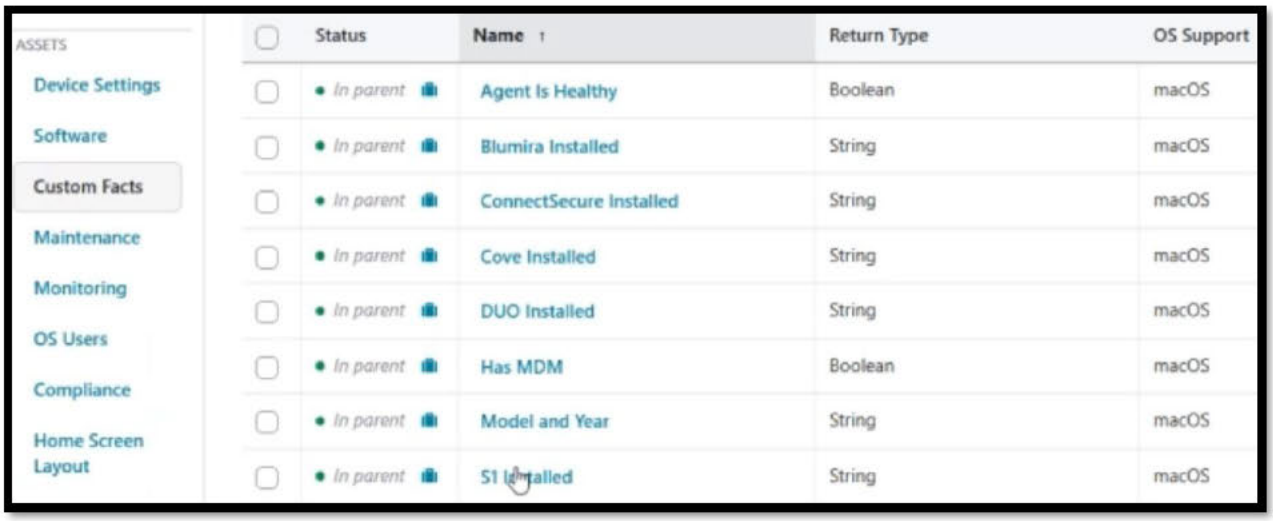

Demonstrating Custom Facts within Addigy

Demonstrating that macOS Updates are enabled

Continued analysis of an N-Able export identified that all devices were running current versions of their respective operating systems (OS). Additionally, zero devices were identified with a lack of disk encryption at rest. Without disk encryption, there is an inherent risk of a malicious actor obtaining access to sensitive data with physical access to a device (such as a stolen workstation).

A	B	C	H	I	J
Endpoint Name	**Site**	**Last Logged In Us**	**Model Name**	**OS**	**OS Version**
[redacted]	[redacted]	[redacted]	Mac14,9	macOS	macOS,26.4.1
			HP - HP EliteBook 840 14 inch G11 Notebook F	Windows	Windows 11 Pro,22631
			HP - HP EliteBook 860 16 inch G11 Notebook F	Windows	Windows 11 Pro,26100
			SAMSUNG ELECTRONICS CO., LTD. - 960XGL	Windows	Windows 11 Pro,26200
			HP - HP EliteBook 860 16 inch G11 Notebook F	Windows	Windows 11 Pro,26100
			SAMSUNG ELECTRONICS CO., LTD. - 960XGL	Windows	Windows 11 Pro,26200
			SAMSUNG ELECTRONICS CO., LTD. - 960XGL	Windows	Windows 11 Pro,26100
			HP - HP EliteBook 860 16 inch G10 Notebook F	Windows	Windows 11 Pro,22631
			HP - HP EliteBook 860 16 inch G10 Notebook F	Windows	Windows 11 Pro,22631

Analyzing OS versions

A	B	C	H	O
Endpoint Name	**Site**	**Last Logged In Us**	**Model Name**	**Disk Encr**
[redacted]	[redacted]	[redacted]	Mac14,9	On
			HP - HP EliteBook 840 14 inch G11 Notebook F	On
			HP - HP EliteBook 860 16 inch G11 Notebook F	On
			SAMSUNG ELECTRONICS CO., LTD. - 960XGL	On
			HP - HP EliteBook 860 16 inch G11 Notebook F	On
			SAMSUNG ELECTRONICS CO., LTD. - 960XGL	On
			SAMSUNG ELECTRONICS CO., LTD. - 960XGL	On
			HP - HP EliteBook 860 16 inch G10 Notebook F	On
			HP - HP EliteBook 860 16 inch G10 Notebook F	On

Analyzing disk encryption statuses

Microsoft Defender

The Microsoft Defender portal (https://security.microsoft.com) is an integrated combination of several products within the Defender suite, including Defender for Office 365, Defender for Endpoint, Defender for Identity, and Defender for Cloud Apps. While many aspects of Microsoft Defender do require additional licensing, such as attack simulation training and Defender Vulnerability Management, there are still numerous key security features available with [redacted] current licensing.

Abacus examined the security settings within Defender for Microsoft 365 to review the configuration of alert and mail flow policies.

Abacus then reviewed the mail flow rules that reduce common email attack methods. Microsoft Defender for Office 365 has settings to inform and protect users from potentially malicious emails, including safety tips for domain and user impersonation, spam and phishing detection, and Safe Links.

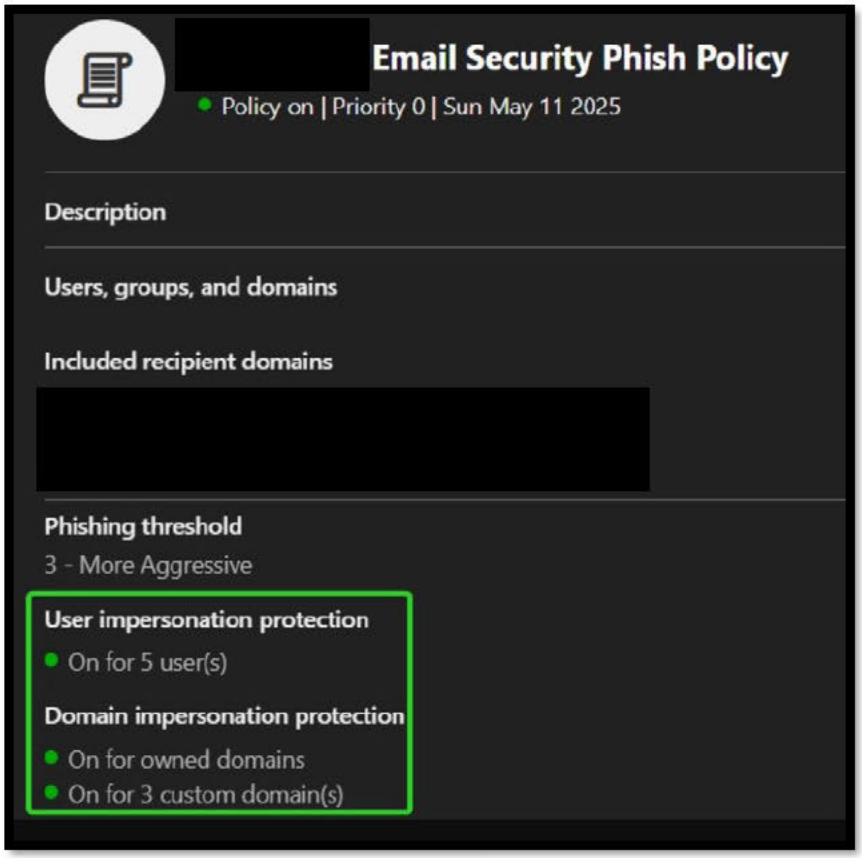

Demonstrating the Office 365 configured AntiPhish policy

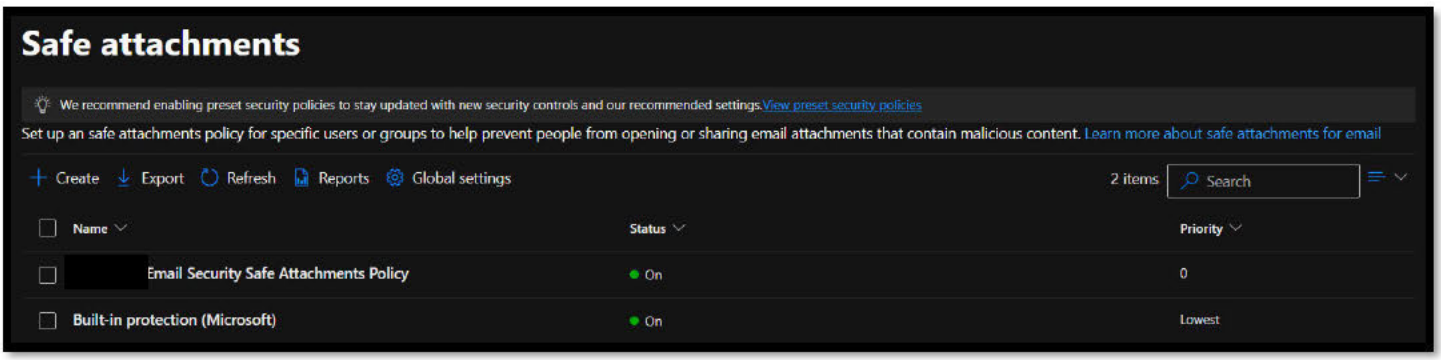

Demonstrating that safe attachments are enabled for all users

Malicious actors are increasingly leveraging enterprise communications platforms such as Microsoft Teams for sophisticated phishing campaigns, as it inherently bypasses most anti-phishing security controls associated with email systems. Abacus assessed the Microsoft Teams settings of [redacted] environment, which confirmed that external domains were permitted to send Microsoft Teams messages to [redacted] personnel, which introduces this potential attack vector.

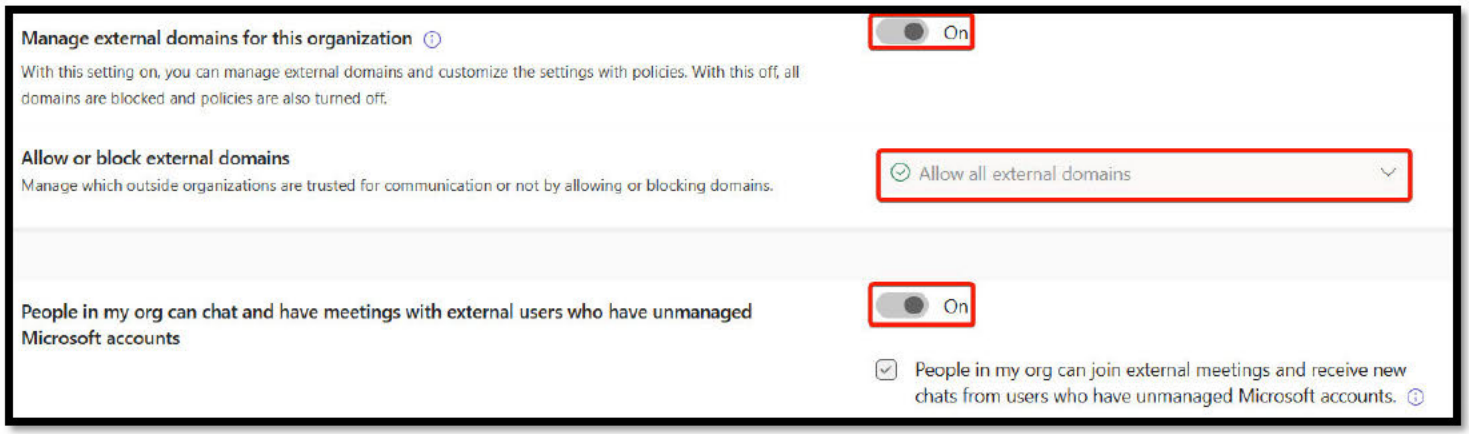

Demonstrating that all external domains are permitted to send MS Teams messages to [redacted] personnel

Microsoft SharePoint/OneDrive

Abacus also reviewed the configurations within Microsoft SharePoint/OneDrive. Specifically, the controls around sharing and sharing permissions were closely analyzed. The current configuration is overly permissive and should be adjusted. Of course, the permissiveness is at the discretion of [redacted] but Abacus recommends, at a minimum, that the sharing configurations be set to "New and existing guests".

Best practices could also be improved with file and folder links. This configuration for the type of link was adequately set to "only people in your organization", and a link expiration was set to 30 days. Additionally, the default sharing permission is set to "View" for sharing links, which is the most secure configuration. The external sharing settings, namely, could result in numerous security-related threats, such as an expanded attack surface, data loss, data integrity implications, and even data leakage. These threats could pose an even greater risk when considering that no data loss prevention controls currently exist. Finally, Abacus reviewed the shares and files, and while there were no noticeable membership misconfigurations, it is recommended that members be reviewed on occasion so that the principles of least privilege are followed.

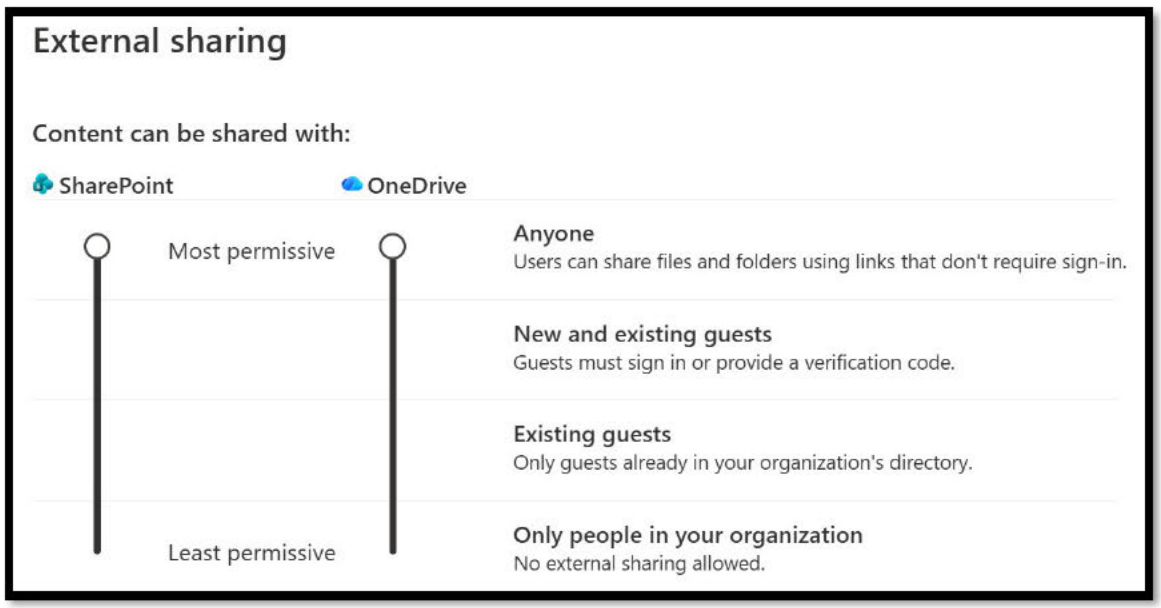

Reviewing the external sharing configurations, which were overly permissive

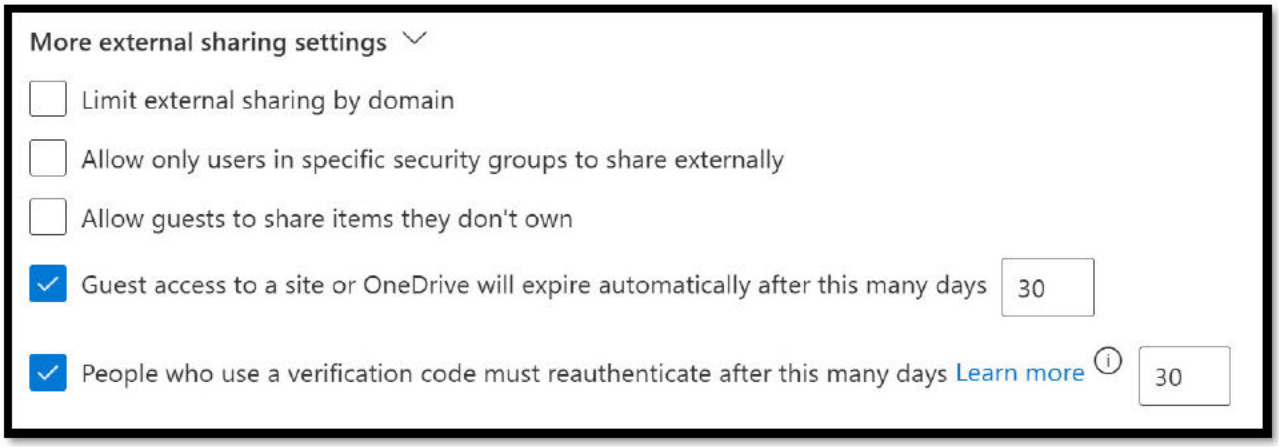

Reviewing additional external sharing configurations, which were in line with security best practices

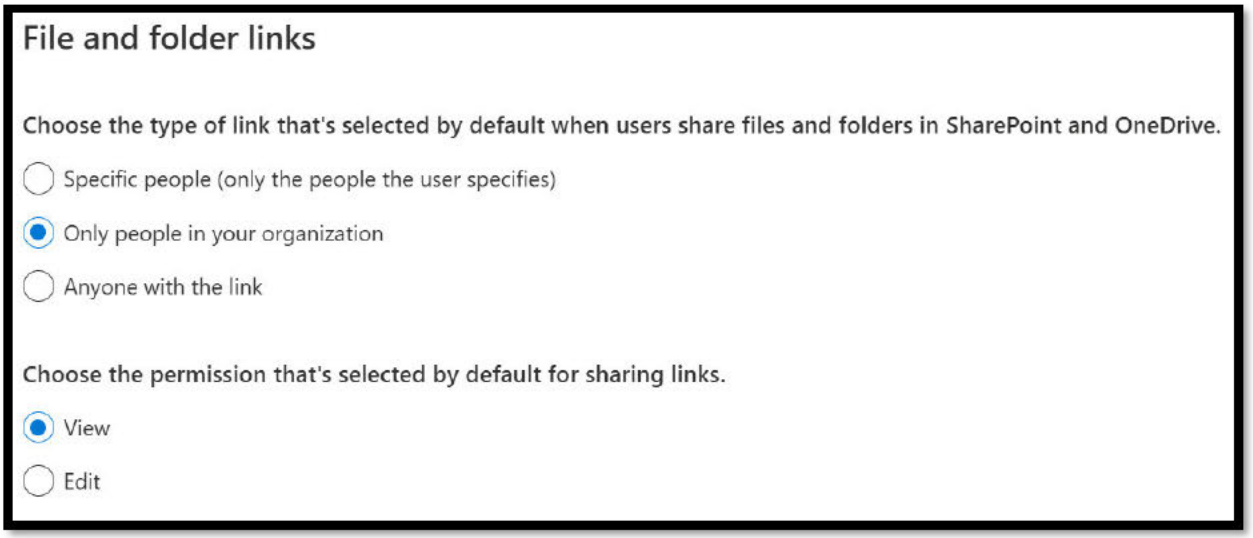

Reviewing the file and folder links settings, which were also in line with security best practices

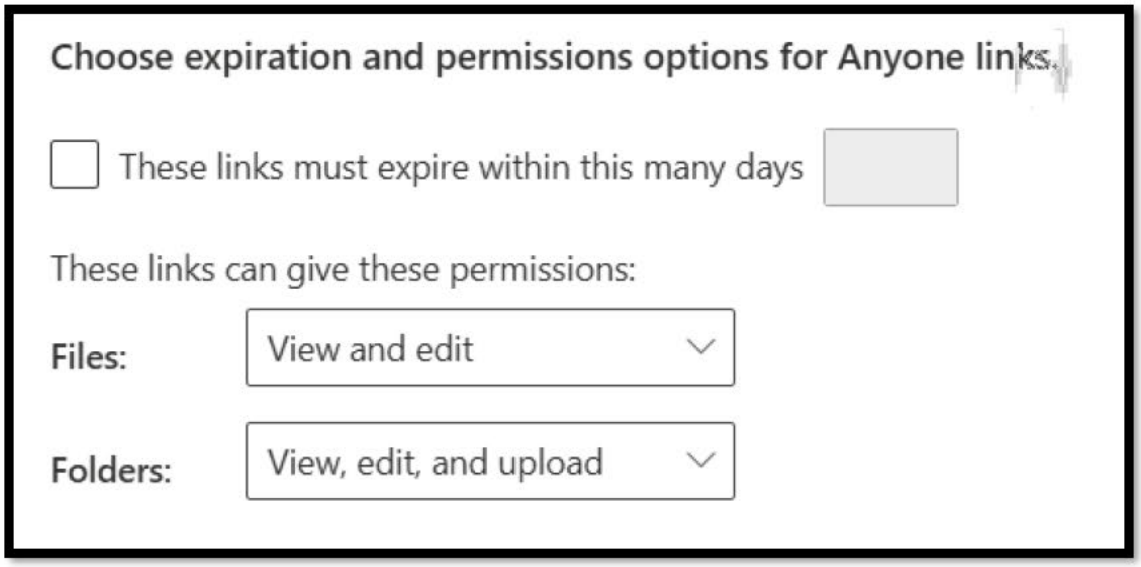

Demonstrating that there was no expiration set for "Anyone links"

Finally, Abacus reviewed other security configurations, such as the "Idle session sign-out." This means that SharePoint and other Microsoft 365 web sessions remain active indefinitely, which increases the risk of unauthorized access on shared or unmanaged devices. It's recommended to enable the idle timeout to protect sensitive data by signing out inactive users automatically.

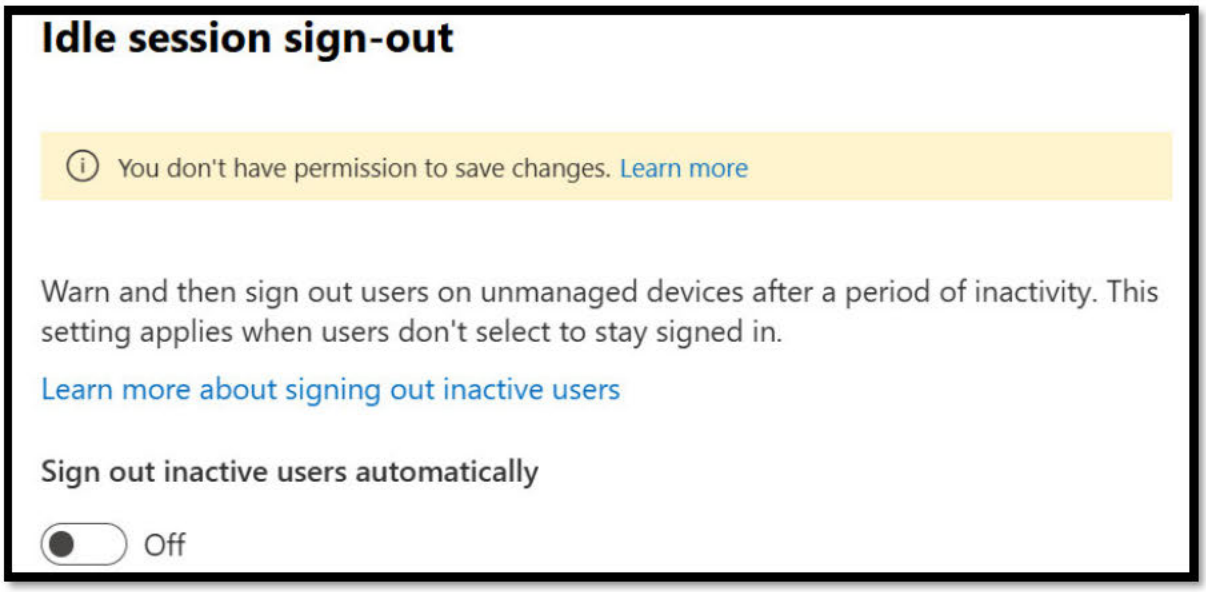

Identifying the "Idle session sign-out" configuration

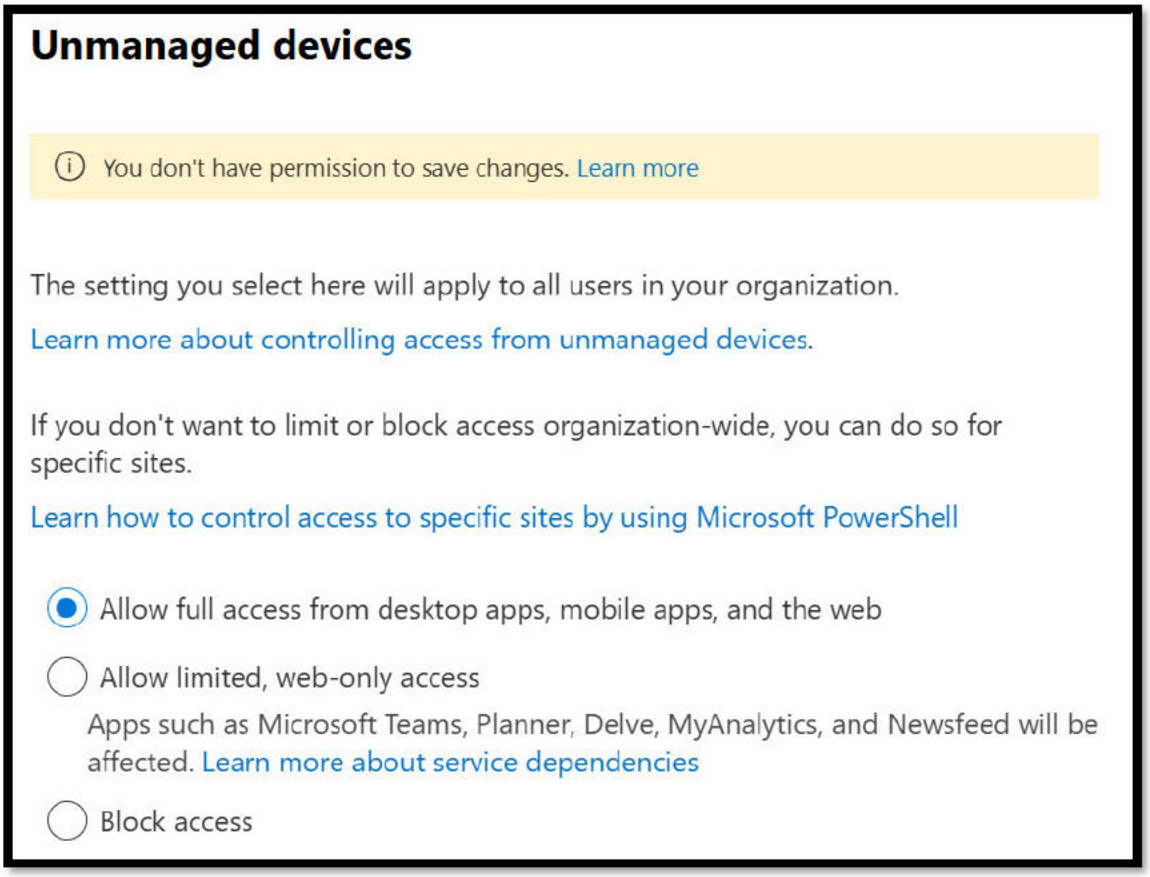

Identifying the "Unmanaged devices" configuration

Microsoft Purview

Microsoft Purview (https://purview.microsoft.com) is a solution within the Microsoft 365 ecosystem that provides integrated tools to increase visibility and nuanced control over data protection and governance. Microsoft Purview was created as a combination of the previous solutions Azure Purview and Microsoft 365 Compliance. While the majority of features within Microsoft Purview do require additional licensing, such as E5 Compliance, several key features such as the ability to apply sensitivity labels to data, data loss prevention controls, and retention policies, can be configured with [redacted] current licensing.

Abacus reviewed the Data Classification and Data Loss Prevention (DLP) policies in Microsoft 365. DLP policies work with data classification policies to create and enforce Access Control Lists (ACLs) to ensure that sensitive information and documents cannot be exfiltrated in an unauthorized manner from storage environments. Abacus Group determined that the DLP policy was in "*simulation with notifications*" mode, meaning it does not actively prevent data loss. Abacus Group recommends enabling enforcement after validation to ensure sensitive data is protected.

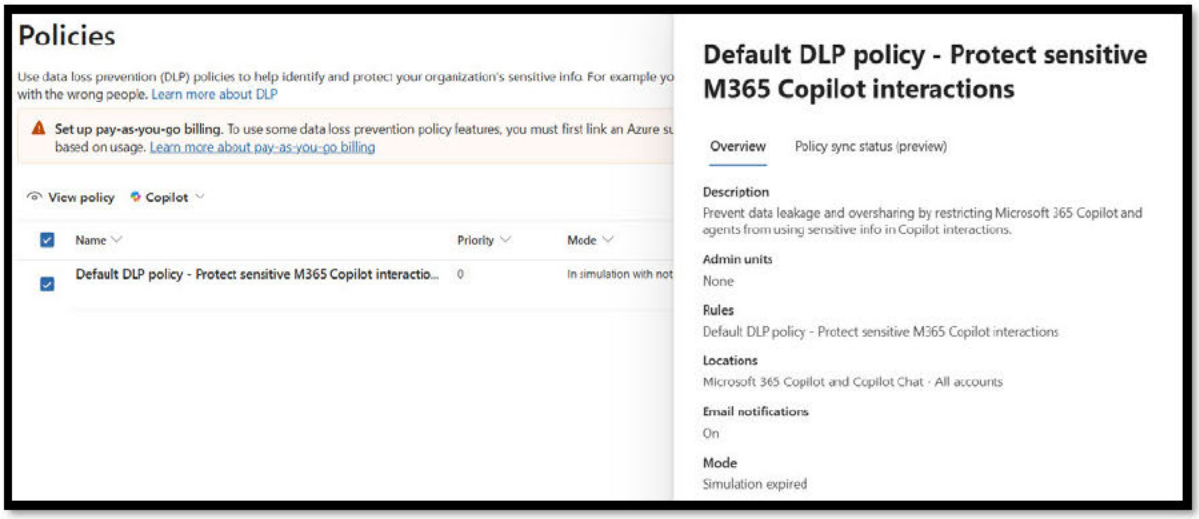

Demonstrating a lack of DLP policies established within Microsoft 365

Abacus noted that while reviewing the data retention policies within Microsoft Purview, the existing policies are based on Messaging Records Management (MRM). Although MRM is a legacy approach primarily used in Exchange Online, these policies are still functional and enforceable. However, Microsoft now recommends using Microsoft Purview retention policies and labels for a more unified and flexible compliance strategy across Microsoft 365 services. Integrating the existing MRM policies into the broader Purview framework could enhance consistency, simplify management, and ensure alignment with modern compliance practices.

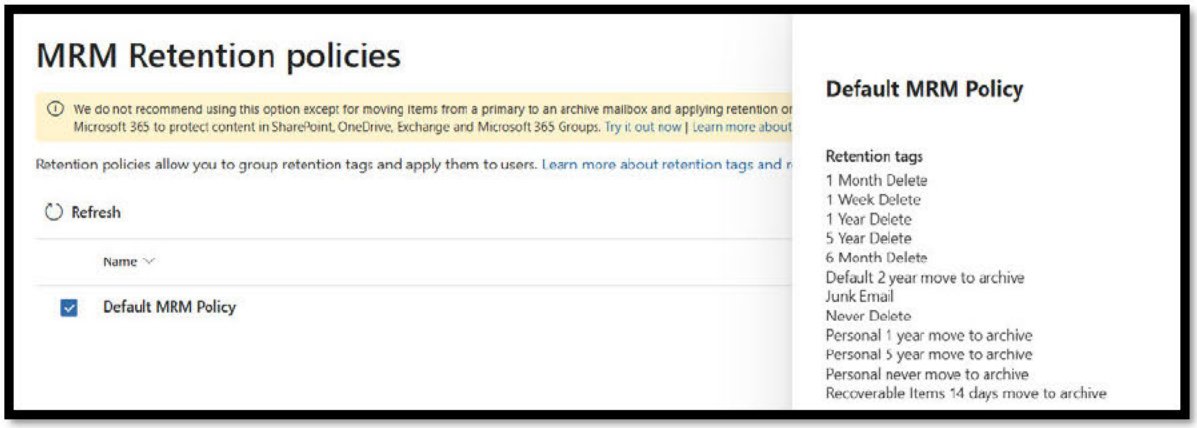

A single MRM Data Retention policy was observed in Microsoft 365

Cloud Infrastructure Analysis

Abacus leveraged public search engines and custom OSINT tooling to determine if [redacted] had any unintentionally exposed cloud resources. This process involved systematically enumerating common Azure Website DNS, storage accounts such as Amazon S3 buckets and Azure Blob, databases, virtual machines, and more.

Once all unauthenticated external enumeration was completed, Abacus leveraged direct access to AWS to further evaluate the security configurations. Abacus leveraged manual enumeration and automated tooling for the review.

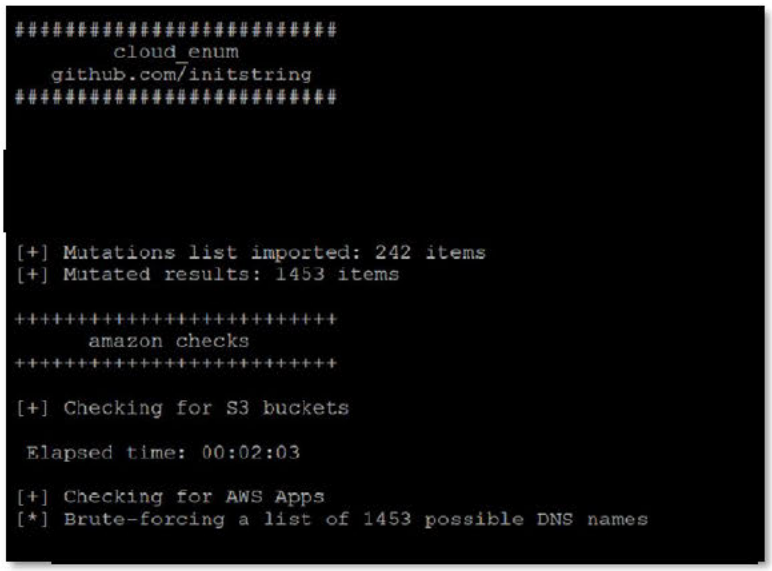

Demonstrating that no exposed Amazon cloud resources were found via cloud_enum

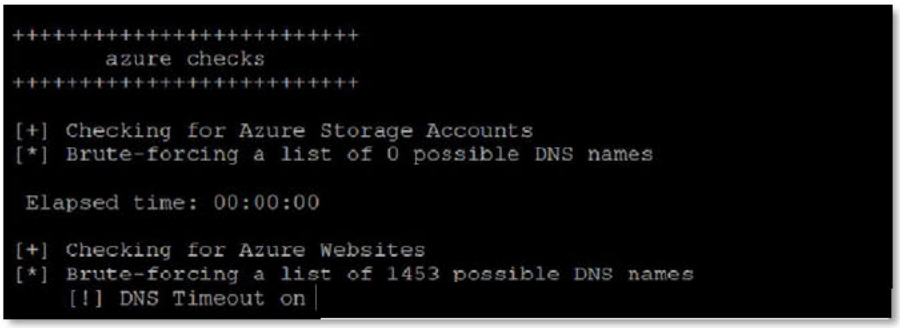

Demonstrating that no exposed Azure cloud resources were found via cloud_enum

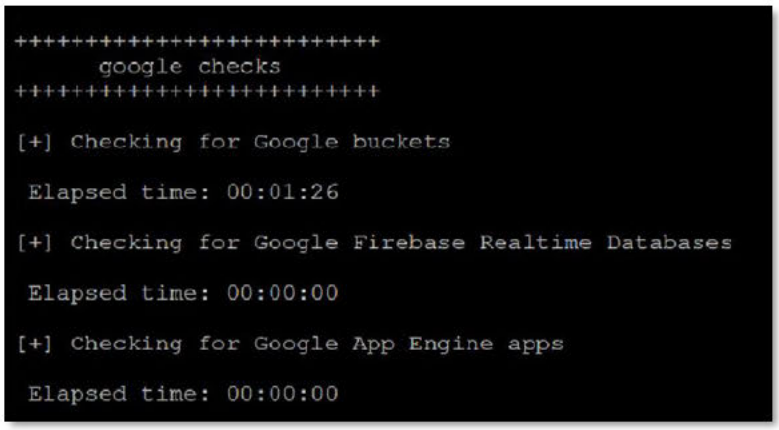

Demonstrating that no exposed Google cloud resources were found via cloud_enum

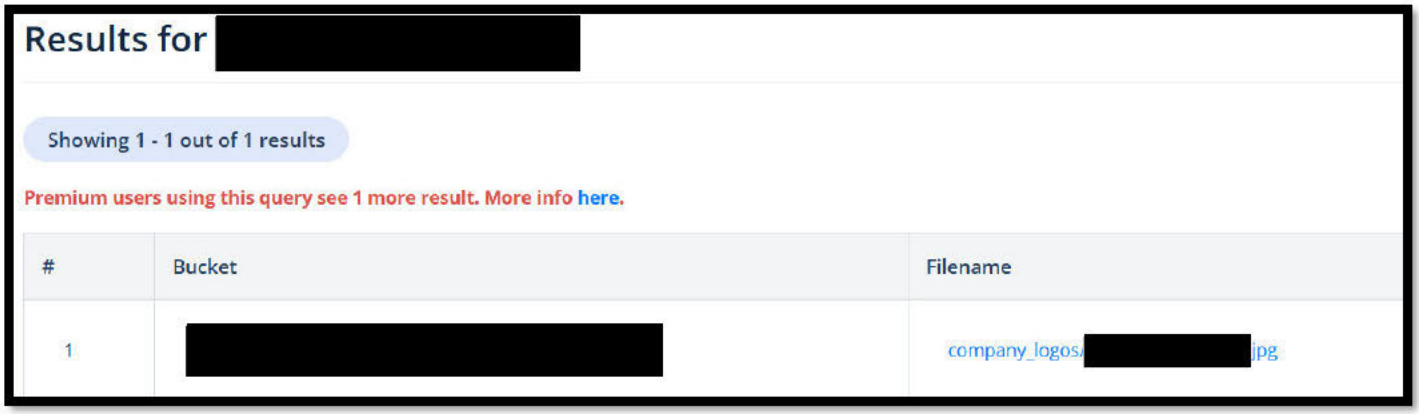

Demonstrating that [redacted] found the [redacted] logo hosted on [redacted]

Microsoft Azure

Once all unauthenticated external enumeration was completed, Abacus leveraged the provided access to Microsoft Azure to further evaluate the security configurations utilizing a combination of well-known open-source tools, custom-built tools, and manual evaluation.

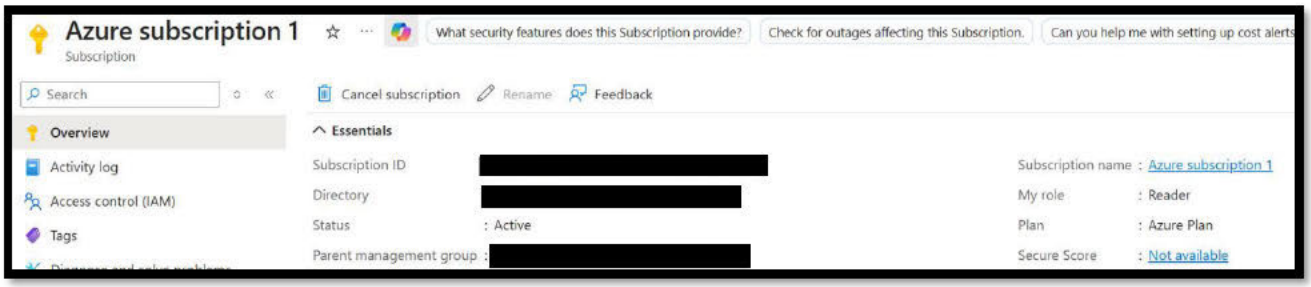

Demonstrating an overview of the subscription that was reviewed during the assessment

Abacus identified several guest user accounts within the Azure environment. These external users were associated with [redacted] own domain or closely affiliated client domains, indicating usage in support of legitimate business collaboration rather than access by unknown third parties. As such, the presence of these guest accounts does not represent a significant security concern; however, periodic review of guest access remains a recommended best practice to ensure continued appropriateness and adherence to the principle of least privilege.

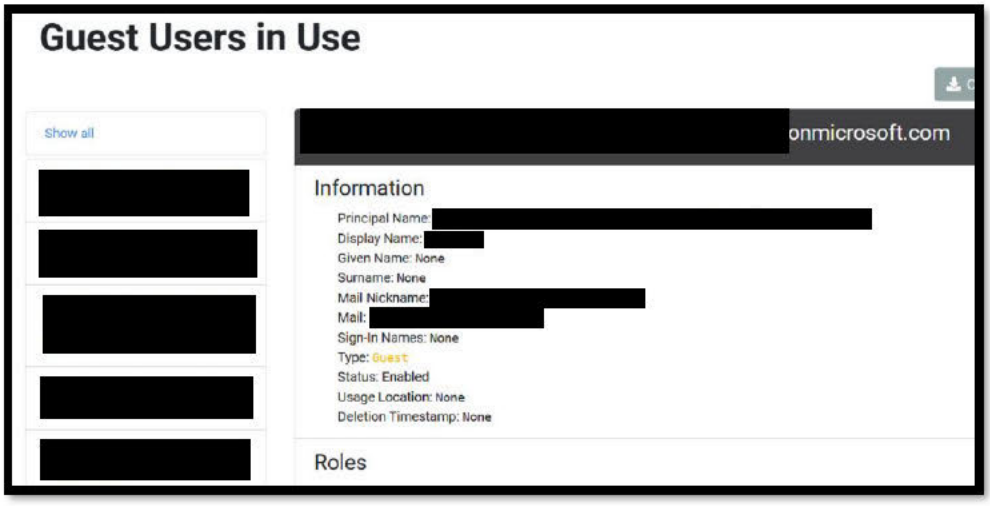

Demonstrating active guest users within the Azure environment

Abacus then utilized API access to aggregate and evaluate substantial configuration information on Azure resources and resource groups by leveraging an open-source tool called ScoutSuite. This scan revealed a very small but well-configured Azure environment.

```
2026-04-23 13:54:01           scout[14862] INFO Launching Scout
2026-04-23 13:54:01           scout[14862] INFO Authenticating to cloud provider
/home/administrator/environments/scoutsuite/lib/python3.12/site-packages/msal/application.py:204: UserWarning:
 Please upgrade msal-extensions. Only msal-extensions 1.2+ can work with msal 1.30+
  warnings.warn(
To sign in, use a web browser to open the page https://login.microsoft.com/device and enter the code
 to authenticate.
2026-04-23 13:54:39           scout[14862] INFO No subscription set, inferring
2026-04-23 13:54:40           scout[14862] INFO Running against subscription
920e5
2026-04-23 13:54:40           scout[14862] INFO Gathering data from APIs
2026-04-23 13:54:40           scout[14862] INFO Fetching resources for the Azure Active Directory service
2026-04-23 13:54:40           scout[14862] INFO Fetching resources for the RBAC service
2026-04-23 13:54:40           scout[14862] INFO Fetching resources for the Security Center service
2026-04-23 13:54:40           scout[14862] INFO Fetching resources for the SQL Database service
2026-04-23 13:54:40           scout[14862] INFO Fetching resources for the Storage Accounts service
2026-04-23 13:54:40           scout[14862] INFO Fetching resources for the Key Vault service
2026-04-23 13:54:40           scout[14862] INFO Fetching resources for the Network service
2026-04-23 13:54:40           scout[14862] INFO Fetching resources for the Virtual Machines service
```

Demonstrating authenticated Azure reconnaissance and vulnerability analysis via ScoutSuite

Network Security Groups were manually reviewed in addition to the automated ScoutSuite review. No insecure inbound rules were identified. Inbound traffic is restricted to internal Azure services, with a default deny rule preventing unsolicited access from external sources. This configuration significantly reduces the exposed attack surface and aligns with Azure security best practices.

Name ↑		Resource Group	Location	Subscription
[redacted]	···	[redacted]	East US	Azure subscription 1
[redacted]	···		East US	Azure subscription 1

Demonstrating the two network security groups

Inbound Security Rules						
65000	AllowVnetInBound	Any	Any	VirtualNetwork	VirtualNetwork	Allow
65001	AllowAzureLoadBalancerInBound	Any	Any	AzureLoadBalancer	Any	Allow
65500	DenyAllInBound	Any	Any	Any	Any	Deny
Outbound Security Rules						
65000	AllowVnetOutBound	Any	Any	VirtualNetwork	VirtualNetwork	Allow
65001	AllowInternetOutBound	Any	Any	Any	Internet	Allow
65500	DenyAllOutBound	Any	Any	Any	Any	Deny

Demonstrating that inbound traffic is restricted to internal Azure resources

Demonstrating a well-configured network setup within Azure

Abacus then reviewed the Network interfaces to validate effective security enforcement. All interfaces are protected by associated Network Security Groups that enforce a default-deny inbound posture. No insecure inbound exposure or unintended network behaviors (such as IP forwarding) were identified. Public IP exposure is mitigated through Network Security Groups enforcing deny-by-default inbound access.

Demonstrating that IP Forwarding is disabled and the public IP is protected by an NSG

Review of virtual machine configurations identified strong baseline security controls. Virtual machines were observed to utilize managed disks with encryption enabled, and no unencrypted or orphaned disks were identified. Disk encryption is implemented using Azure-managed keys, which provides protection for data at rest and aligns with standard Azure security practices for environments of this size and complexity.

While customer-managed disk encryption keys were not observed, the current encryption approach is appropriate given the environment's scope and does not represent a security concern.

Demonstrating that OS and Data Disks are not encrypted with CMK

The review of the Azure role-based access control assignments did not identify a dedicated role or assignment specifically responsible for administering resource locks. While built-in administrative roles are able to manage locks by default, no explicit separation of duties for resource lock management was observed at the time of review.

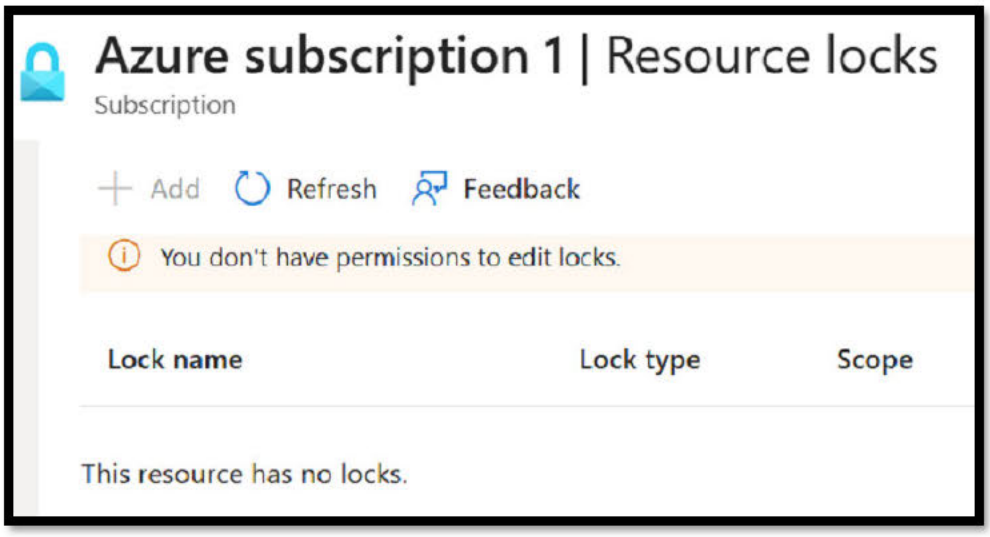

Demonstrating that there are no resource locks

Azure Activity Log Alerts were not observed for security-sensitive configuration changes, such as Network Security Group modifications or policy assignments. While these actions are recorded within Azure Activity Logs, automated alerting to notify administrators of such changes in near-real time was not configured at the time of review.

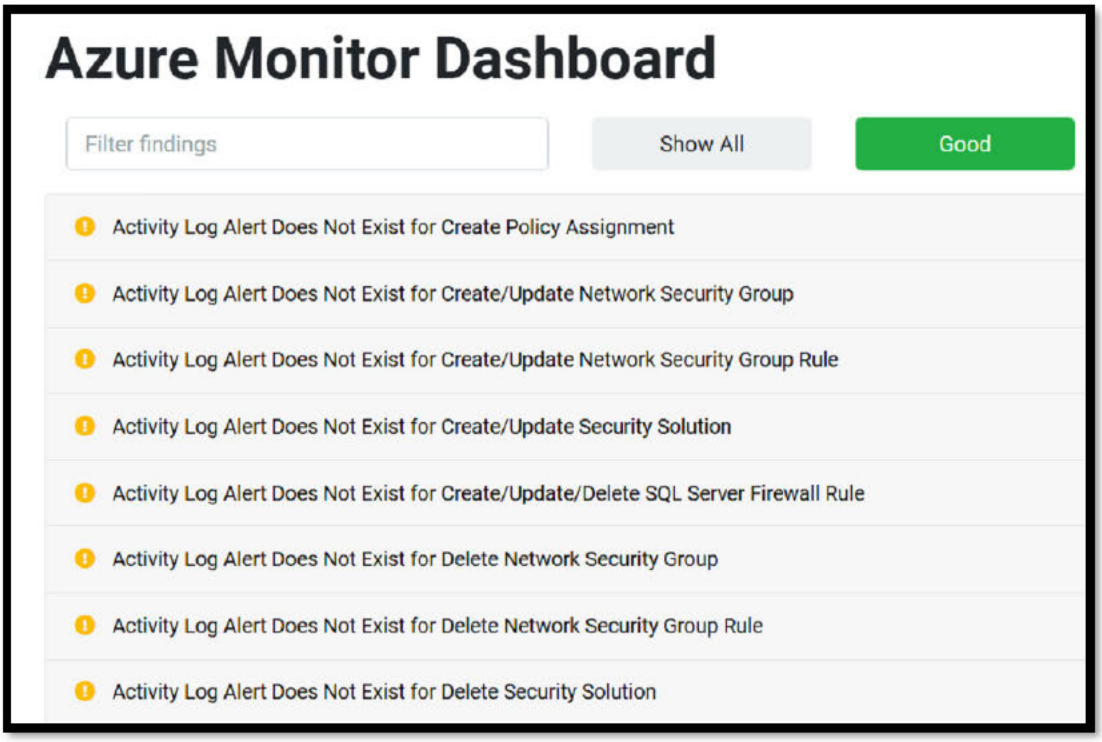

Demonstrating that Activity Log Alerts are not enabled

Attack Surface Analysis

Abacus assessed the external corporate attack surface from the perspective of a malicious actor. Abacus then assessed [redacted] external corporate attack surface by utilizing both active and passive analysis techniques. The fingerprinting process involved collecting details of open TCP/UDP ports, listening services, specific software versions, operating system identifiers, and more. Validation was performed using both automated tools and personal inspection. It is worth noting that while these types of tools and techniques are used in penetration tests, this assessment is not a formal penetration test, as any attempts to exploit identified vulnerabilities would be considered out of scope.

Abacus approached this attack surface analysis from a fully black-box perspective. DNS subdomain enumeration was performed using techniques such as wordlist-based DNS resolution, reverse IP address lookups, MX and SOA record inspection, and Autonomous System (AS) number lookups. This helped provide a baseline of IP addresses and subdomains associated with [redacted] including server infrastructure and external web applications. These techniques allowed Abacus to establish an understanding of the full attack surface of the organization. This information was then corroborated by leveraging a powerful open-source intelligence (OSINT) tool called OWASP Amass.

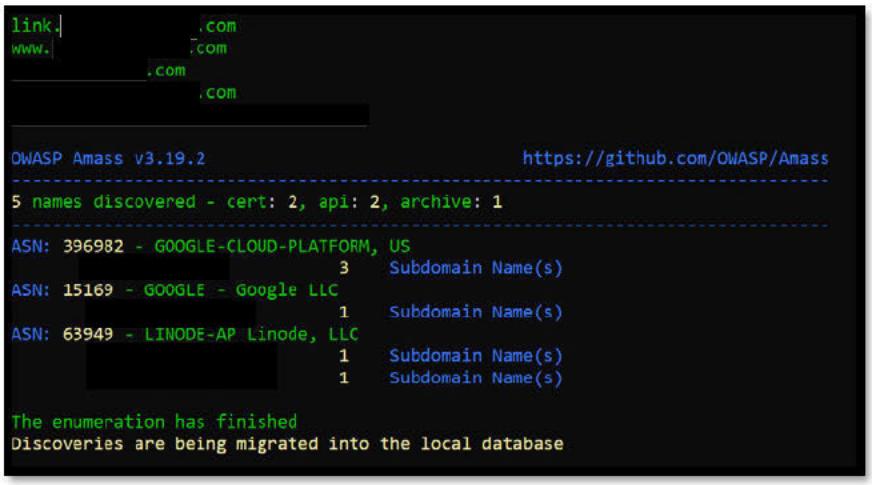

Demonstrating domain and subdomain enumeration

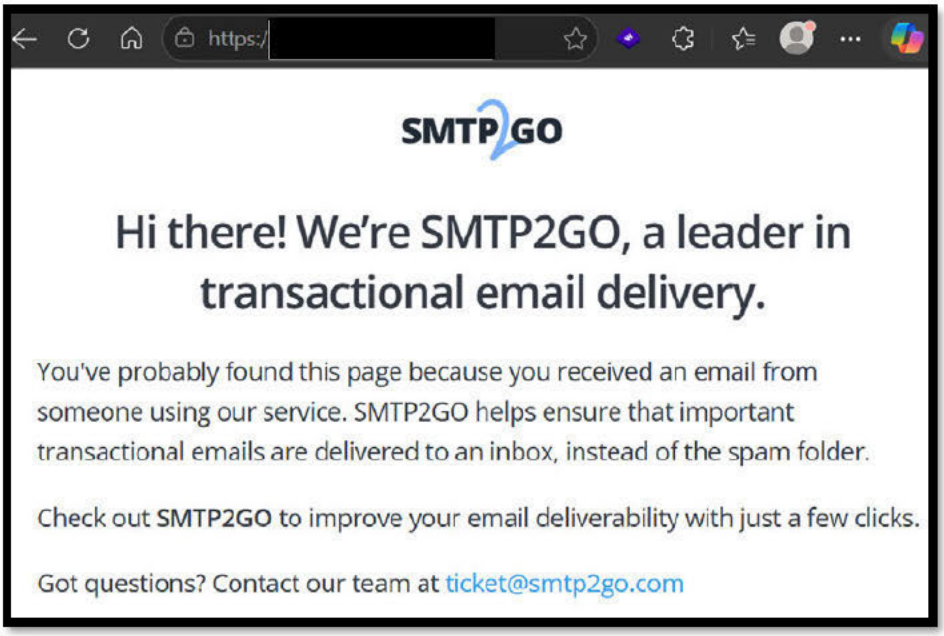

Attempting to navigate to link.[redacted]com, resulting in a landing page for SMTP2GO

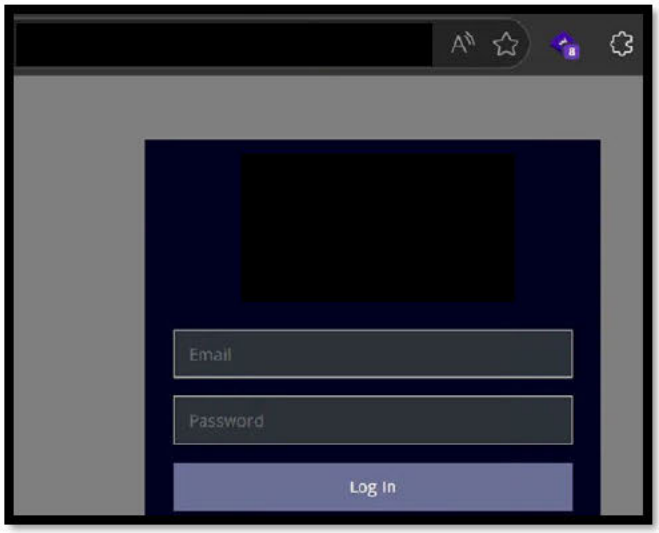

Attempting to navigate to portfoliobuilder.[redacted]com

Virus Total Analysis, AlienVault Open Threat Exchange (OTX), and Criminalip.io were also used to collect information about any potential Indicators of Compromise (IoCs) or relevant OSINT data potentially affecting [redacted] information systems.

IoC analysis via VirusTotal

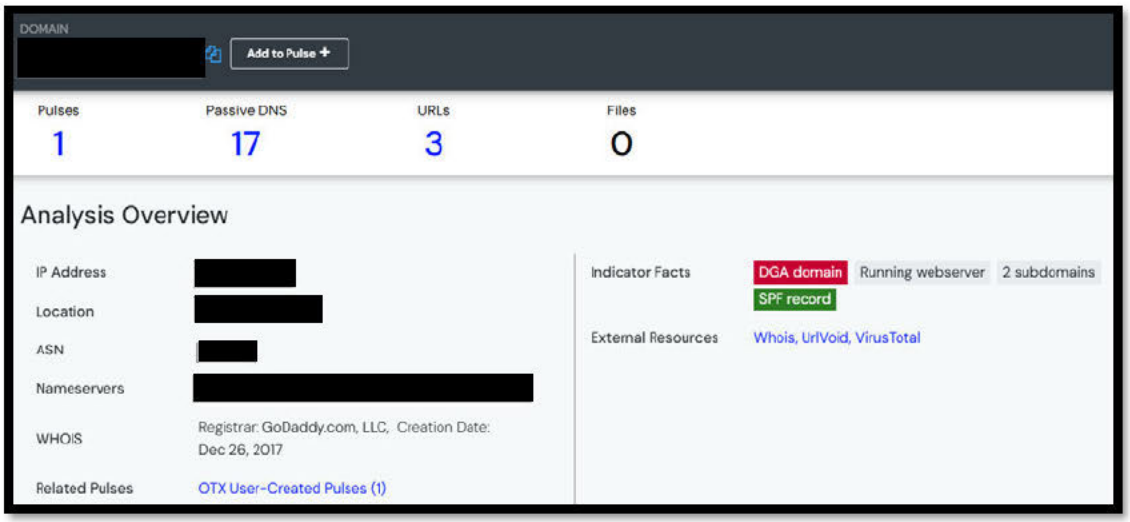

IoC analysis via AlienVault OTX

During a review of the [redacted]*com* domain using AlienVault Open Threat Exchange (OTX), a historical threat intelligence indicator associated with IP address [redacted] was observed. The indicator was created approximately three years prior to the assessment and appears to be related to previously used shared or third-party infrastructure rather than current [redacted] systems. No evidence of active or ongoing malicious activity was identified, and this observation is provided for informational awareness only.

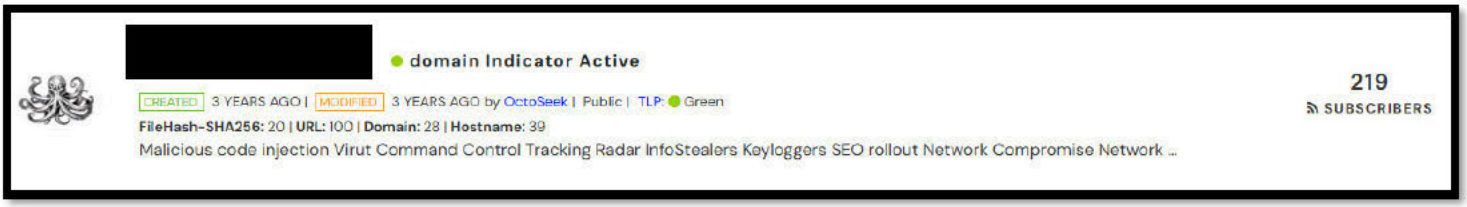

IoC analysis via AlienVault OTX, resulting in the identification of a pulse

Abacus also utilized a publicly available domain impersonation, bitsquatting, and typosquatting monitoring tool called Dnstwist to assess if potentially malicious actors had purchased domains for use in social engineering campaigns against [redacted] employees. Dnstwist is a free tool leveraged by numerous commercial products such as Splunk ESCU, RecordedFuture, SpiderFoot, DigitalShadows, SecurityRisk, SmartFense, ThreatPipes, Palo Alto Cortex XSOAR, Rapid7, Mimecast, Watcher, Intel Owl, PatrOwl, VDA Labs, and Appsecco.

Demonstrating domain impersonation analysis via DNStwist

Mail records for the domain were analyzed for misconfigurations with Sender Policy Framework (SPF), DomainKeys Identified Mail (DKIM), and Domain-based Message Authentication, Reporting, and Conformance (DMARC). SPF works by ensuring that only approved mail servers can send emails on behalf of the organization. DKIM uses digital signatures to verify that email messages have not been tampered with or modified in transit. DMARC expands on these protections by enforcing a policy for unverifiable emails, which include actions like rejection or quarantine. Utilized together, these controls help mitigate email spoofing and help protect against phishing attacks.

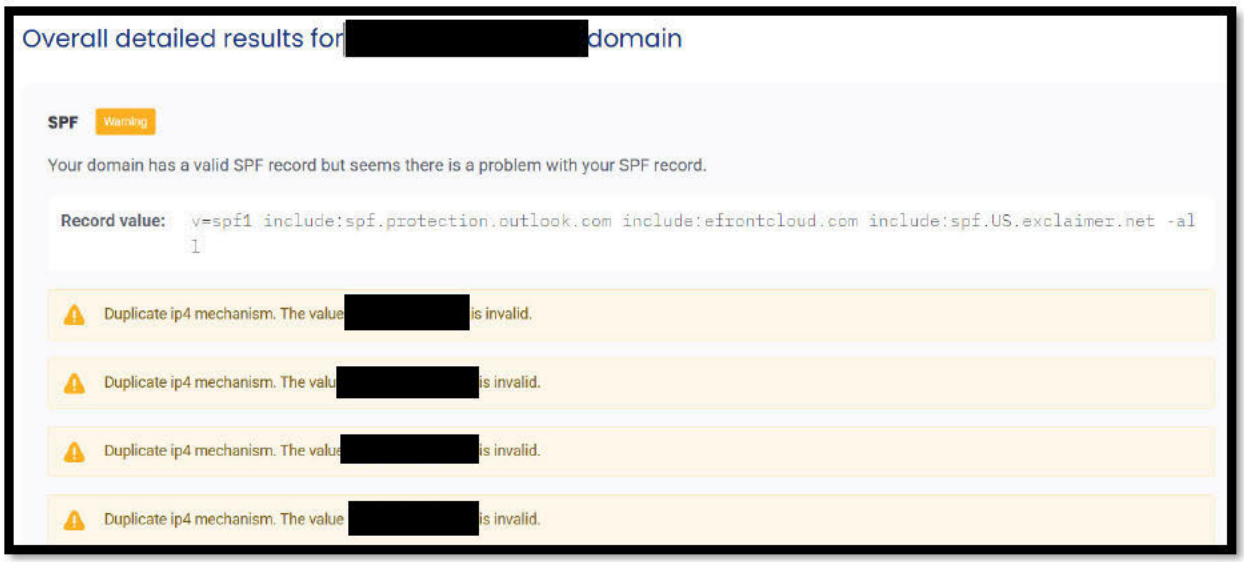

Observed SPF record for [redacted]

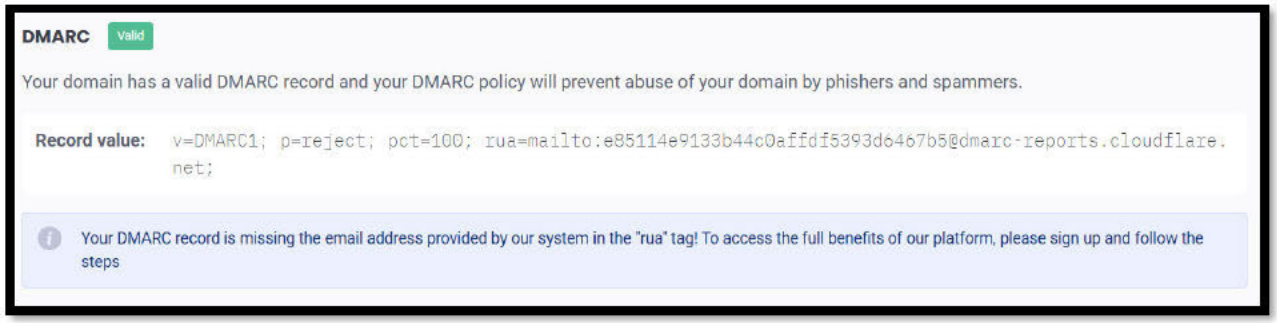

Mail record assessment demonstrating that a DMARC record was enabled

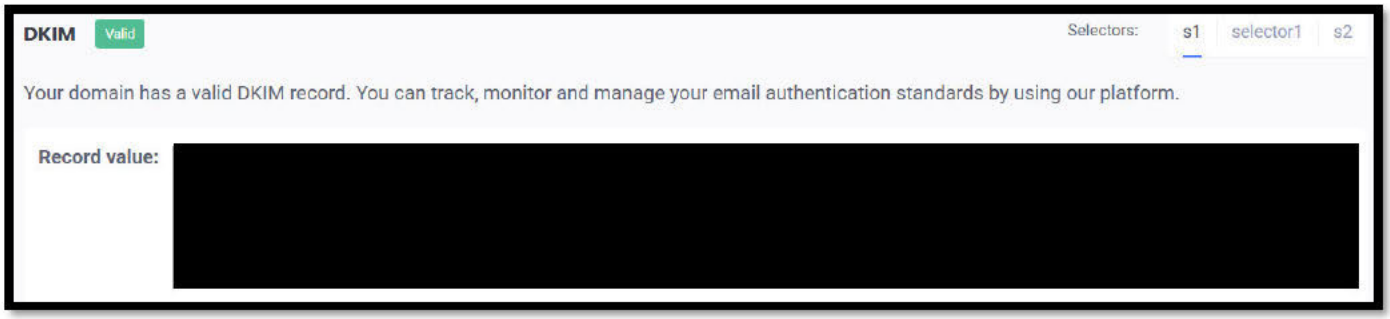

Mail record assessment demonstrating that a DKIM record was enabled

Abacus observed that [redacted]*com* was initially flagged as listed on the Spamhaus ZEN blacklist. Upon further validation, the domain was not found in the Spamhaus database, indicating the alert was likely a false positive.

To further supplement passive analysis and OSINT data collection, Abacus used Google Dorking to search for any exposed log files, company documents, login pages, databases, or paste site dumps that could appear online within search indexes. During reconnaissance exercises, Google Dorking is often used to reveal sensitive metadata about the target company that would assist during active attacks. Abacus identified [redacted] ADV brochure through open-source searches, including Google Dorking techniques. As the ADV brochure is a publicly available regulatory filing, its accessibility does not represent a security concern and is noted for awareness only.

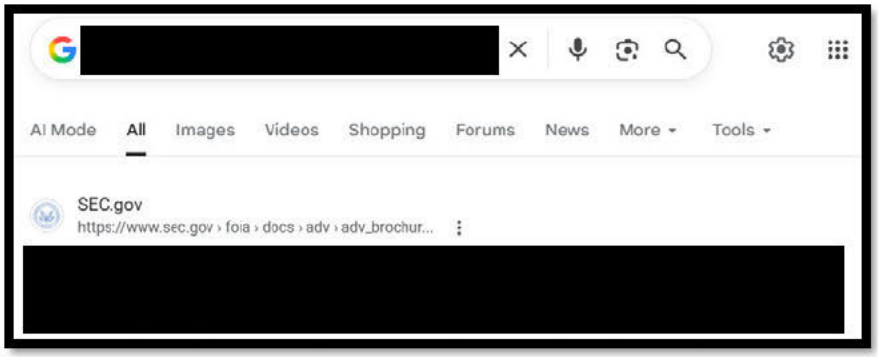

Demonstrating that an adv_brochure was uncovered via Google Dorking

Abacus then confirmed whether the website implemented the recommended security headers. Abacus identified that [redacted] had both a Content Security Policy (CSP) and an HTTP Strict Transport Security (HSTS) policy. This security header (CSP) provides an explicit policy specifying acceptable origins from which various site components may be loaded. This adds an important layer of security that helps to protect against certain classes of attacks, like Cross-Site Scripting (XSS), Clickjacking, and Subresource Integrity violations. In most cases, a well-defined CSP can completely take the place of certain older security headers, such as X-Frame-Options and X-XSS-Protection.

Demonstrating security header analysis

Abacus also leveraged a variety of automated vulnerability scanners across the scope of endpoints. Such vulnerability scanners included, but were not limited to, Nessus, OWASP Zap, Sn1per, and OpenVAS. During this stage of the engagement, Abacus leveraged these vulnerability scanners to corroborate evidence from the manual analysis and to identify any additional vulnerabilities for review and analysis.

Open Port	Service Name	Service Version	Web Fingerprint	Exploit Status
80	http	HTTP	• Radix UI • React • Google Cloud • Next.js • Google Cloud Trace • Google Cloud CDN • Webpack • HTTP/3 • Node.js • Marko • Priority Hints • Open Graph	Not Vulnerable
443	https	HTTPS Google Frontend	• React • Google Cloud • Next.js • Google Cloud Trace • Google Cloud CDN • Webpack • HTTP/3 • Node.js • Marko	Not Vulnerable

[redacted] IP address [redacted]

Demonstrating attack surface analysis

Abacus identified the use of a GraphQL API exposed over HTTPS on [redacted] during a website scan. Review of the GraphQL endpoint revealed multiple confirmed weaknesses related to query complexity and schema exposure, including conditions that allow circular references, alias overloading, field duplication, and directive overloading, all of which could be abused to trigger denial-of-service conditions through resource-intensive queries.

In addition, GraphQL introspection and field suggestion features were found to be enabled, exposing internal schema details that may aid an attacker in enumerating available queries and crafting targeted requests. These conditions indicate that GraphQL security controls, such as query depth limiting, cost analysis, and schema exposure restrictions, are not sufficiently enforced, increasing the risk of service disruption and information disclosure.

Demonstrating the GraphQL risk findings via a website scan

Concurrently with this risk assessment, Abacus also conducted daily Vulnerability Management as a Service (VMaaS) via both agent-based and agentless vulnerability scans for all of [redacted] corporate assets. This program enabled additional insight into each individual workstation and server by establishing a detailed common platform enumeration (CPE) and mapping those platforms to known common vulnerabilities and exposures (CVEs).

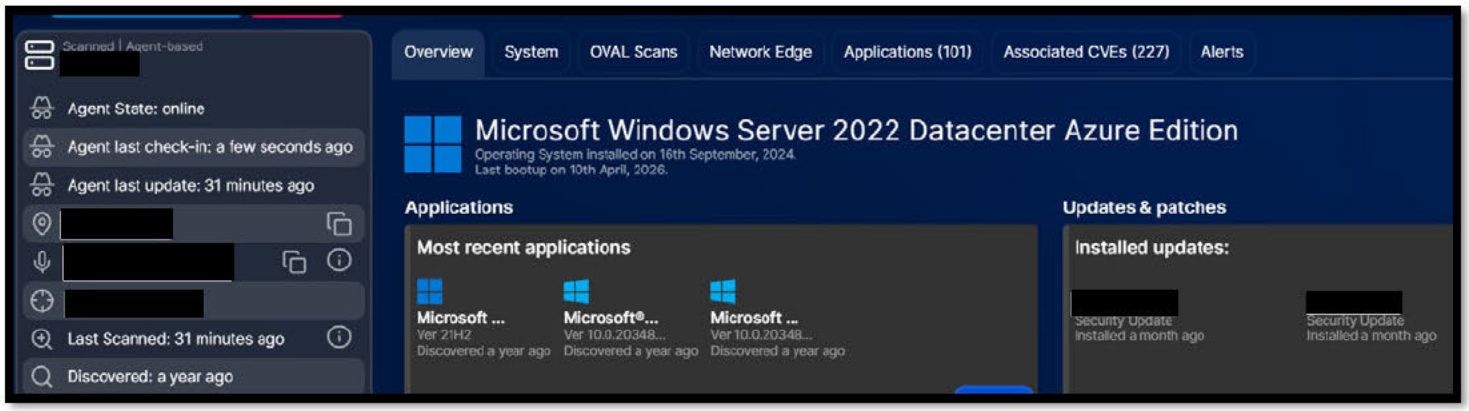

Demonstrating deployed agents on Abacus's VMaaS solution

One unique and insightful feature of Abacus's VMaaS platform is the ability to conduct Open Vulnerability and Assessment Language (OVAL) scans. OVAL is a standard mechanism for benchmarking the security hardening configurations of workstation, server and network systems, originally developed by MITRE and currently maintained by the Center for Internet Security (CIS).

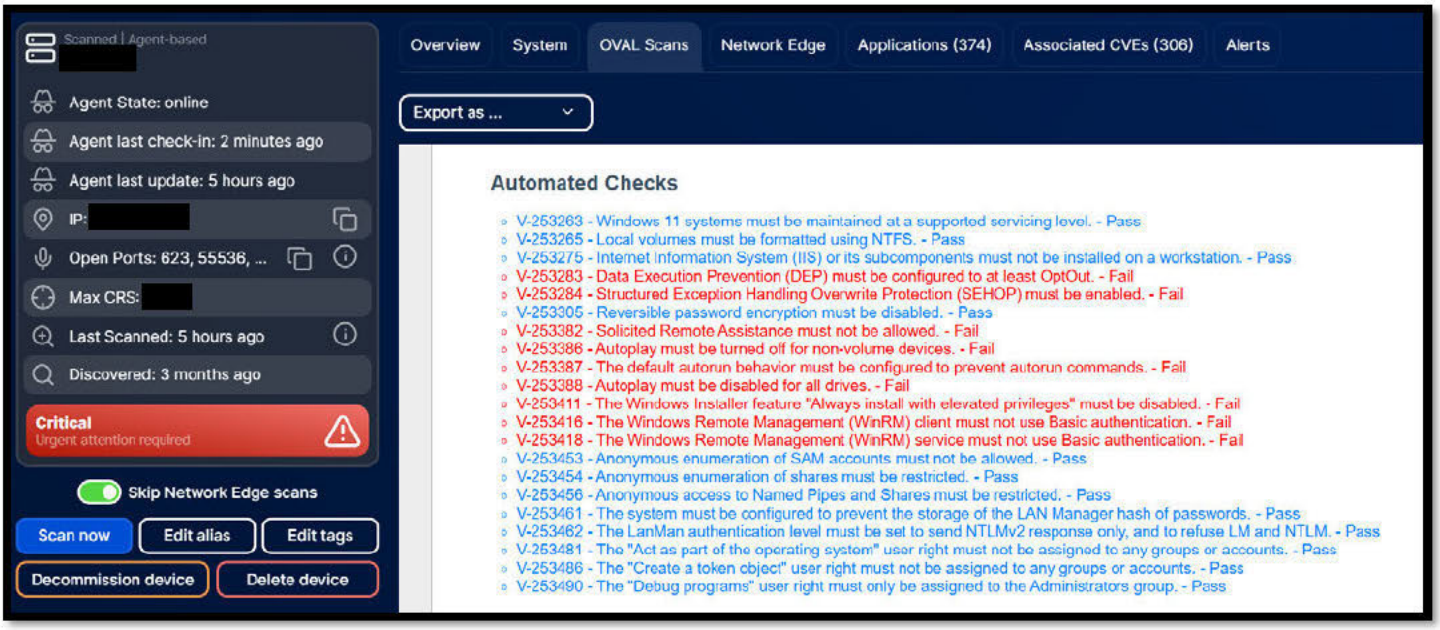

Demonstrating a small portion of OVAL scan results on a single endpoint

This process of OVAL scanning identified several minor misconfigurations in the configuration of endpoints, including multiple weaknesses in Kerberos.

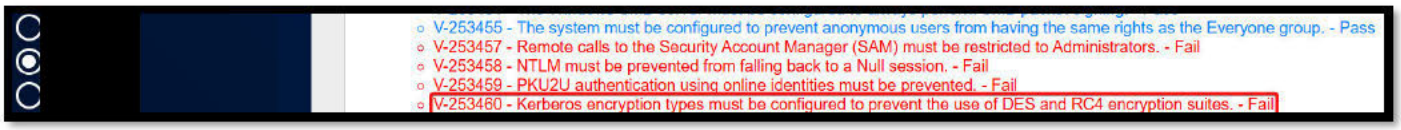

Demonstrating OVAL scanning results (1 of 3)

Autorun commands being allowed to execute may introduce malicious code to a system. Configuring this setting prevents autorun commands from executing, which mitigates a common method of maintaining persistence utilized by malicious actors.

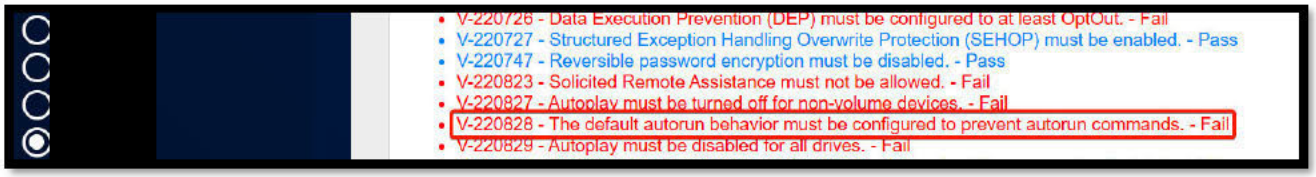

Demonstrating OVAL scanning results (2 of 3)

Abacus discovered that remote assistance was enabled on Windows systems, which allows another user to view or take control of a user's local session.

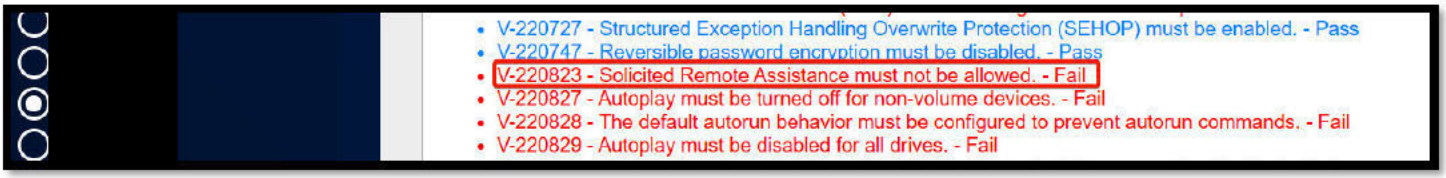

Demonstrating OVAL scanning results (3 of 3)

On-premises Active Directory & Entra ID Review

Abacus conducted a focused review of [redacted] On-premise Active Directory environment using Purple Knight, a non-intrusive AD security assessment tool that scans for common attack paths, privilege exposures, misconfigurations, weak Kerberos settings, and other identity-related risks. The assessment indicated that [redacted] maintains a strong and mature Active Directory security posture, with only a small number of Indicators of Exposure (IOEs), which are configuration weaknesses that could increase security risk, having been identified. A comprehensive list of all risk findings from Purple Knight and the affected endpoints and users will be referenced in Appendix E of this report.

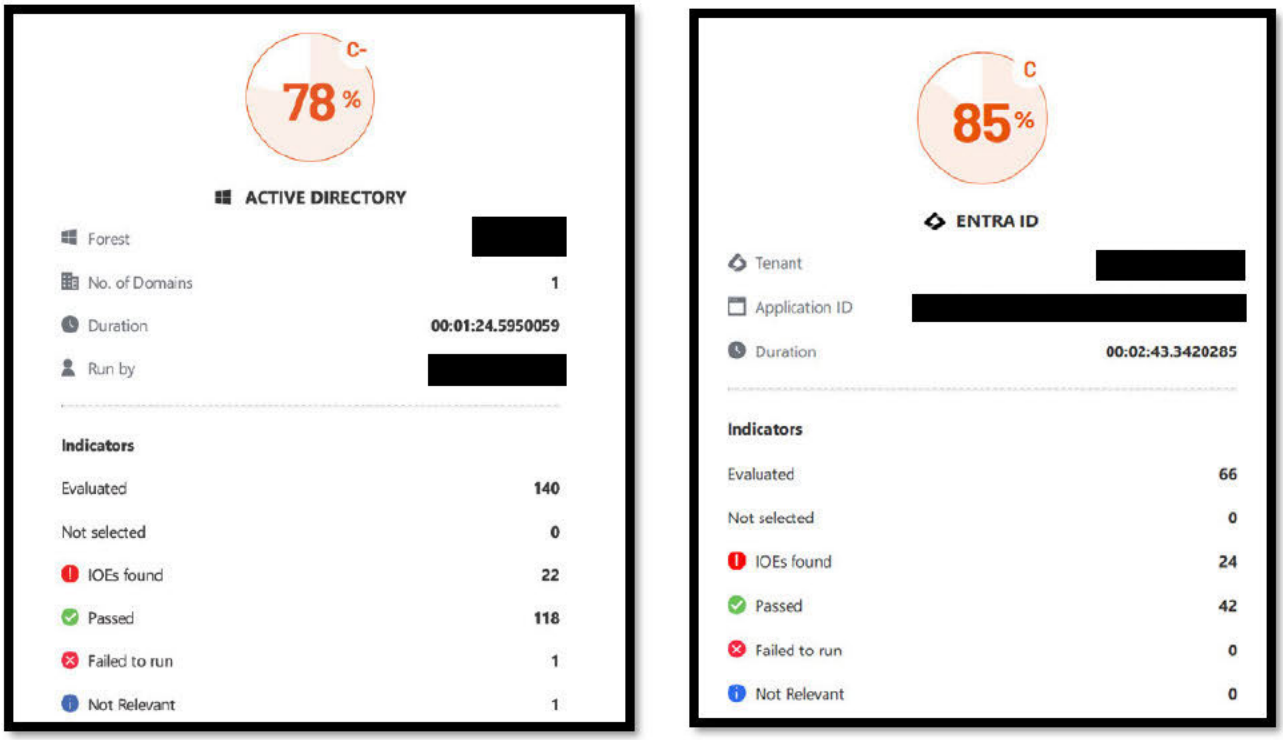

Demonstrating the Active Directory and Entra ID findings via Purple Knight

Purple Knight – Active Directory

Purple Knight identified non-default accounts with directory replication (DCSync) permissions assigned at the domain level. Accounts with DCSync rights can request password replication data directly from domain controllers, effectively allowing them to retrieve password hashes for all users, including domain administrators, without interacting with endpoints.

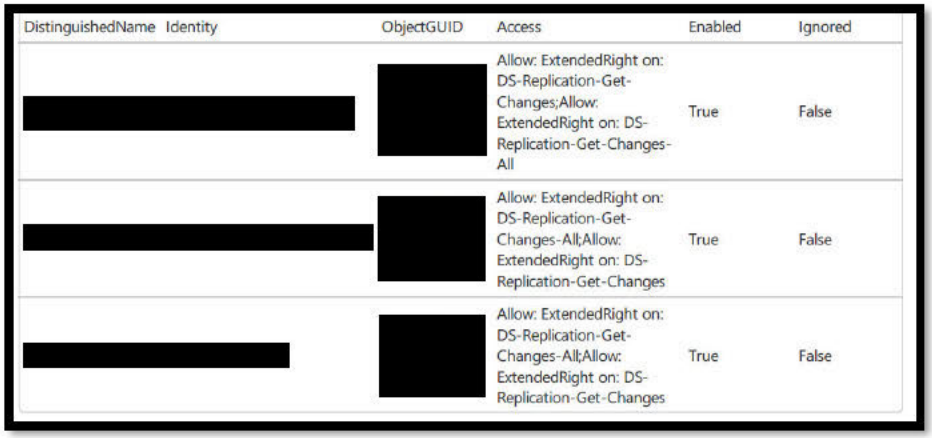

Demonstrating that Non-default principals with DC Sync rights are enabled

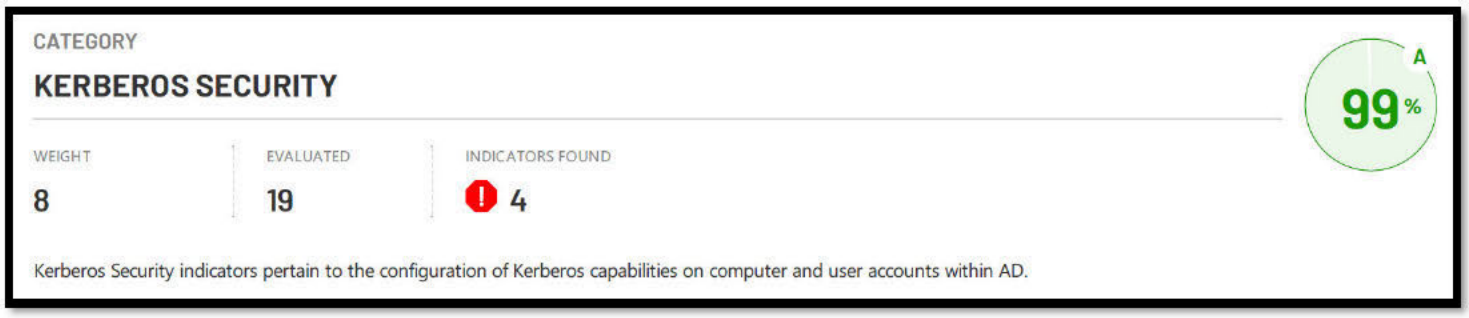

Demonstrating the Kerberos Security category

Multiple Kerberos-related weaknesses were identified, including an outdated KRBTGT password, legacy encryption types (RC4/DES) enabled on domain controllers, risky delegation configurations, and SPNs not supporting AES encryption.

The [redacted] domain's KRBTGT account password has not been rotated since 2019, leaving the Kerberos authentication infrastructure at elevated risk. Because the KRBTGT account is responsible for signing Kerberos tickets, an outdated password significantly increases the impact of credential compromise, allowing attackers, once privileged access is obtained, to forge valid Kerberos tickets and maintain long-term, hard-to-detect access across the domain. Regularly rotating the KRBTGT password is a core Active Directory hardening practice to prevent persistent domain-wide compromise.

Result
Found 1 domains where the KRBTGT account password has not changed in the last 180 days.

DistinguishedName	ObjectGUID	Attribute	LastChanged	DaysPassed	EventTimestamp
[redacted]	[redacted]	unicodePwd	8/21/2019 7:54:28 PM	2423	8/21/2019 7:54:28 PM

Demonstrating that the KRBRTGT account password has not been changed since 2019

Abacus identified that the Domain Controller, [redacted] still permits legacy Kerberos encryption types, including RC4_HMAC_MD5, alongside modern AES options. Continued support for RC4 allows clients or misconfigured systems to negotiate weaker, deprecated encryption during authentication, exposing the domain to downgrade, brute-force, and man-in-the-middle attacks. Modern Kerberos deployments are expected to rely on AES-based encryption, and leaving RC4 enabled increases the risk of credential compromise and weakens overall authentication security.

Result
Found 1 Domain Controllers that support RC4 or DES encryption.

DistinguishedName	ObjectGUID	SupportedEncryptionTypes	EventTimestamp
[redacted]	11dacfb7-1ab6-437e-ba2a-d0e1640d1b77	AES 128; AES 256; RC4_HMAC_MD5	9/17/2024 1:57:52 PM

Demonstrating that the domain controller, [redacted] supports the use of RC4

Purple Knight identified that the built-in Domain Administrator account had not undergone a password change since 2021. Given that this account has a well-known, easily discoverable SID, and even though it should not be used for routine Active Directory administration, its password should be rotated regularly to limit the risk of brute-force compromise.

DistinguishedName	ObjectGUID	SamAccountName	LastChanged	PasswordLastSet	DaysSinceLastSet
[redacted]	[redacted]	Administrator	10/7/2021 3:50:47 PM	10/7/2021 3:50:47 PM	1646

Demonstrating that the built-in domain account password has not been changed since 2021

Purple Knight – Entra ID

Purple Knight identified multiple Microsoft Entra ID configuration weaknesses that reduce the tenant's resilience against identity compromise and privilege escalation. The review identified identity and device lifecycle hygiene gaps, including inactive user accounts and a significant number of stale devices that remain enabled without recent authentication activity. In addition, guest invitation settings were found to be overly permissive, allowing external accounts to be invited without sufficient restriction. Other reported Entra ID items were reviewed and determined to be mitigated, disabled, or not applicable based on the current tenant configuration and are not included as findings below.

Demonstrating users and devices that have been inactive for over 90 days

A significant number of devices (94) remain registered or joined to Entra ID without authentication activity for over 90 days. Stale device objects increase the number of trusted identities in the tenant and complicate device lifecycle management, while potentially enabling unintended access paths if reactivated or misused.

PrincipalName	ObjectID	PrincipalType	OperatingSystem	LastSignIn	Created	Ignored
[redacted]	[redacted]	Device	Windows	9/6/2019 2:48:44 PM	10/8/2018 2:19:35 PM	False
		Device	Windows	8/9/2019 3:18:11 PM	10/9/2018 8:35:27 PM	False
		Device	Windows	2/3/2021 2:18:30 AM	10/10/2018 11:29:22 PM	False
		Device	Windows	3/30/2023 10:27:52 PM	10/16/2018 3:00:40 PM	False
		Device	Windows	4/29/2019 7:08:42 PM	10/19/2018 1:30:52 PM	False

Demonstrating the affected devices that have been inactive for more than 90 days

Guest invitation settings allow users to invite external accounts without sufficient restriction. While guest user access to directory properties is appropriately limited, unrestricted invitation capabilities can lead to unmanaged growth of external identities and make it more difficult to track and govern who has access to organizational resources over time.

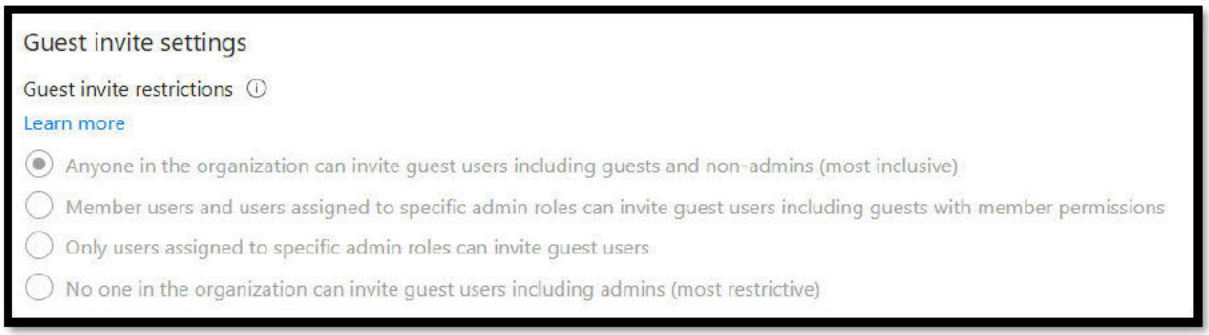

Demonstrating overly permissive guest invite settings

Additionally, Abacus identified that the tenant relies on Microsoft's default banned password protections and does not implement a custom banned password list. While baseline protections are in place, the absence of organization-specific banned password terms represents a missed opportunity to further reduce the risk of weak or predictable password selection.

Demonstrating that the custom banned password protection is not in use

Sentinel One Analysis

Abacus evaluated the configuration policies of [redacted] managed endpoint detection and response (EDR) agents. This analysis included verifying that all in-use endpoints were properly configured with the Sentinel One agent, up to date, and appropriately checking in. Next, Abacus evaluated the detection engine and protection mode settings, which were found to be robust and in line with industry best practices.

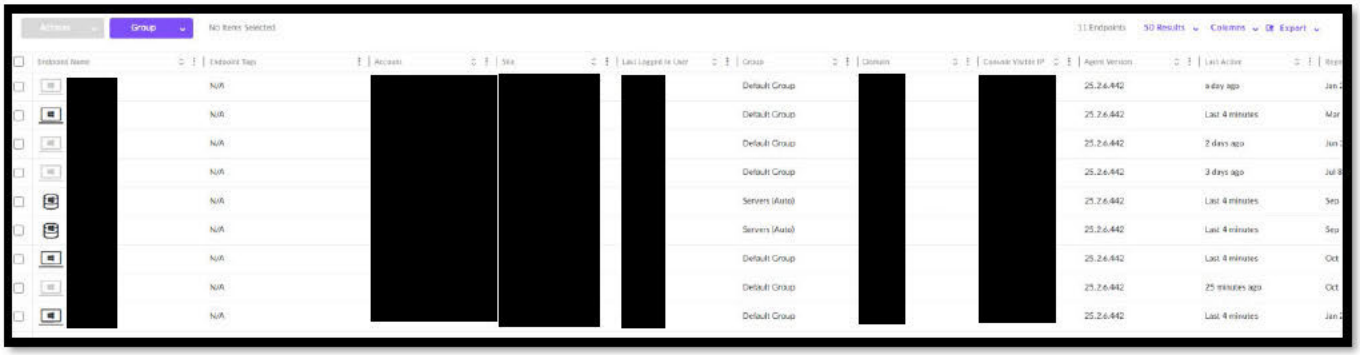

Verifying all actively used endpoints were configured with the SentinelOne EDR Agent

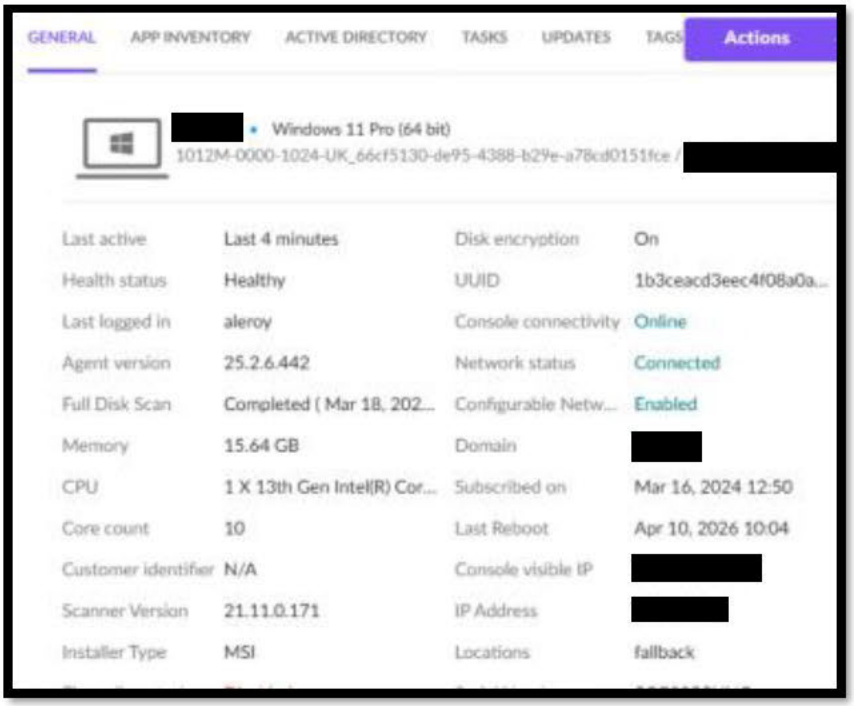

Demonstrating configuration analysis of each endpoint

Verifying Detection Engine and Protection Mode settings

Abacus noted that vulnerability and application scanning were enabled for the endpoints. Considering that [redacted] conducts daily agent-based and agent-less scanning as part of Abacus's Vulnerability Management as a Service (VMaaS) offering, additional scanning via SentinelOne provides increased insight into device-based vulnerabilities.

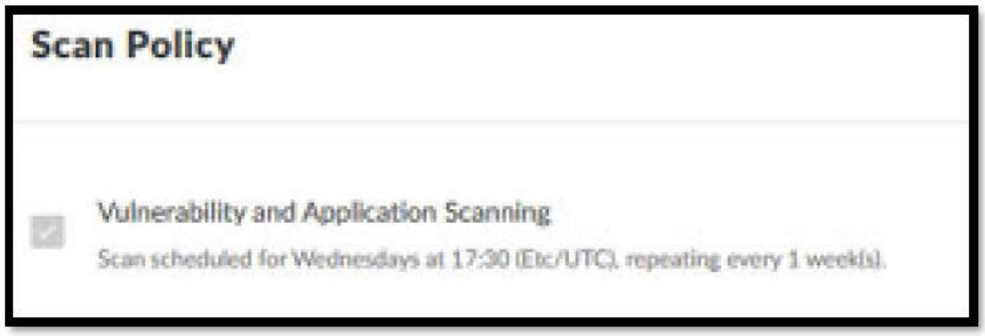

Identifying vulnerability and application scanning enabled within SentinelOne

Interview-Based Analysis

Throughout the engagement, Abacus asked a series of pointed questions to better understand the organization's operations, processes, security controls, and more. The questions, results, and additional interpretations from Abacus can be seen below.

Do you have Cyber Insurance?

1. Response: Yes
2. **Additional Analysis**: Abacus warrants this to be a security best practice.

Do you use a corporate password manager?

1. Response: No, but members are allowed to use their own.
2. **Additional Analysis**: Abacus warrants this to be a moderate security risk. More information can be found in the risk details.

Do you have any technical security controls surrounding the process of sending wires?

1. Response: All wires require 2 approvals by senior management. They require MFA tokens and verbal verification with the bank if they're over a certain amount.
2. **Additional Analysis**: Abacus warrants this as a security best practice. Ensuring RBACs and the separation of authority is key to ensuring that internal processes are secure. Additionally, if an investor requests a change to their wire instructions, treat it as a potential red flag. Always verify the request through a known, trusted communication channel, such as a phone call to a previously confirmed number. Sudden changes in payment details can be a sign of account compromise or fraud.

What regulatory bodies are you beholden too? SEC, FCA, NYDFS, etc.

1. Do you have any presence in the European Union which would fall under the purview of DORA?
2. Response: SEC and no.
3. **Additional Analysis**: Abacus does not warrant the response above to constitute any further security needs.

Do you have customized information security policies or a Written Information Security Program (WISP) in place?

1. Do you have a formalized incident response plan in place?
2. Response: Yes and yes.
3. **Additional Analysis**: Abacus does not warrant the response above to constitute any further security needs.

Do you have a formal security risk management program in place to identify, track, remediate, and close out risks?

1. Response: Yes
2. **Additional Analysis**: Abacus does not warrant the response above to constitute any further security needs.

Do you conduct end-user security awareness training on things like phishing (KnowBe4 or similar)?

1. Response: Yes
2. **Additional Analysis**: Abacus does not warrant the response above to constitute any further security needs.

Are users able to use their own personal devices to access company resources (i.e, email, OneDrive, etc)?

1. Response: Yes
2. **Additional Analysis**: Abacus does not warrant the response above to constitute any further security needs, especially with the review that was conducted regarding this level of access.

Do you develop any applications in-house? If yes, what is the cadence at which the applications get penetration testing?

1. Response: Yes, annually.
2. **Additional Analysis**: Abacus does not warrant the response above to constitute any further security needs.

Risk Details

Print Spooler Service Enabled on Domain Controllers
MITRE ATT&CK Tactic: TA0008 – Lateral Movement
Detected: (QoD 75%) - **Impact:** High – **Ease:** High – **Risk:** High

Summary:
Purple Knight identified that the Windows Print Spooler service is enabled and remotely accessible on the domain controllers. The Print Spooler service is linked to several critical vulnerabilities, most notably PrintNightmare, that allow remote code execution, privilege escalation, and forced authentication. When active on domain controllers, it expands the attack surface of the organization's most sensitive identity systems.

Risk Detection Result(s):

Result
Found 1 DCs that have the Print Spooler service running.

HostName	ObjectGUID	Ignored
[redacted]	11dacfb7-1ab6-437e-ba2a-d0e1640d1b77	False

Risk Mitigation Recommendation(s):

- Disable the Print Spooler service on all domain controllers where printing functionality is not explicitly required.
- If business requirements necessitate Print Spooler usage, restrict service access through network controls and ensure the service is fully patched.
- Validate that no dependent services or workflows rely on Print Spooler functionality on domain controllers prior to disabling.
- Monitor domain controllers for authentication coercion attempts and abnormal service-related activity.
- Periodically review domain controller service configurations to ensure only required services are enabled.

Reference(s):
https://calcomsoftware.com/printspooler-vulnerability/

Affected Endpoint(s):
[redacted]

Non-Default Principals With DCSync Rights
MITRE ATT&CK Tactic: TA0006 – Credential Access
Detected: (QoD 65%) - **Impact:** High – **Ease:** High – **Risk:** High

Summary:
Purple Knight identified non-default accounts with directory replication (DCSync) permissions assigned at the domain level. Accounts with DCSync rights can request password replication data directly from domain controllers, effectively allowing them to retrieve password hashes for all users, including domain administrators, without interacting with endpoints.

Risk Detection Result(s):

Result
Found 3 objects with replication permissions.

DistinguishedName Identity	ObjectGUID	Access	Enabled	Ignored
[redacted]	[redacted]	Allow: ExtendedRight on: DS-Replication-Get-Changes;Allow: ExtendedRight on: DS-Replication-Get-Changes-All	True	False
[redacted]	[redacted]	Allow: ExtendedRight on: DS-Replication-Get-Changes-All;Allow: ExtendedRight on: DS-Replication-Get-Changes	True	False
[redacted]	[redacted]	Allow: ExtendedRight on: DS-Replication-Get-Changes-All;Allow: ExtendedRight on: DS-Replication-Get-Changes	True	False

Risk Mitigation Recommendation(s):

- Review domain permissions and remove Replicating Directory Changes and Replicating Directory Changes All from non-essential accounts.
- Restrict DCSync permissions to default domain controllers and explicitly approved administrative roles.
- Monitor directory permission changes on the domain root object.
- Implement tiered administrative access to prevent lower-tier accounts from holding replication rights.

Reference(s):
https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/manage/understand-security-groups#directory-replication
https://attack.mitre.org/techniques/T1003/006/

Affected Endpoint(s):
[redacted]

Lack of Conditional Access Policy to Block Legacy Authentication
MITRE ATT&CK Tactic: TA0006 – Credential Access
Detected: (QoD 75%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
No conditional access policy was identified to block legacy authentication within [redacted] environment. Legacy authentication is a term that refers to authentication requests made by older Office clients that do not use modern authentication or any client that uses legacy mail protocols such as IMAP, SMTP, or POP3. Legacy authentication methods introduce unnecessary risks to the environment, as many legacy protocols do not support multi-factor authentication (MFA). Whenever possible, it is recommended that these legacy authentication requests be blocked in favor of modern protocols.

Risk Detection Result(s):

Policy name	Created by	State	Creation date
Multifactor authentication and reauthentication for risky sign-ins	MICROSOFT	On	8/4/2025, 6:47:52 PM
Multifactor authentication for admins accessing Microsoft Admin Portal	MICROSOFT	Report-only	12/6/2023, 4:01:54 PM
AzureVPN and Apps - Require Hybrid Joined Machine	USER	Report-only	10/8/2024, 8:34:09 PM
Blocked Locations	USER	On	9/27/2023, 3:38:44 PM
Require Hybrid Azure AD joined device	USER	Report-only	1/16/2025, 10:01:12 AM
Require MFA	USER	On	8/8/2023, 8:24:01 PM

Risk Mitigation Recommendation(s):

- Authenticate to https://entra.microsoft.com and select Protection - Conditional Access from the left hand menu.
- Select "Create new policy from templates" and choose "Block Legacy Authentication".
- Leverage report-only mode to first assess the environment to determine which, if any, exceptions would need to be included before implementing any new Conditional Access policies.

Reference(s):
https://learn.microsoft.com/en-us/entra/identity/conditional-access/block-legacy-authentication

Affected Endpoint(s):
Microsoft Entra ID

Insufficient Credential Management in AD
MITRE ATT&CK Tactic: TA0006 – Credential Access
Detected: (QoD 75%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
Abacus identified multiple password hygiene weaknesses within the on-premises Active Directory environment. These include user and administrative accounts with old passwords, accounts configured with the Password Never Expires or Password Not Required flags, and computer accounts (including an SSO-related computer account) with passwords that have not been rotated for extended periods.

In addition, the built-in Domain Administrator account has not had its password changed in more than 180 days. Extended password lifetimes and non-expiring credentials increase the risk of credential compromise through cracking, reuse, or replication-based attacks. Stale computer account passwords can also be leveraged to establish persistence or facilitate lateral movement within the domain.

Collectively, these conditions indicate gaps in credential lifecycle enforcement and increase the likelihood that compromised credentials could be reused or remain undetected.

Risk Detection Result(s):

Result
Found 4 users with password never expires.

DistinguishedName	SamAccountName	ObjectGUID	PasswordLastSet	ServicePrincipalName
[redacted]	[redacted]	bfc1a387-4ca0-4bf1-a246-a3834ecd7d2f	10/7/2021 11:51:50 AM	
[redacted]	[redacted]	66dd2e12-935b-4294-9c0d-a06f8b82bbe1	3/12/2026 9:59:31 PM	
[redacted]	[redacted]	7d72638f-052b-4653-9e41-fca5447da964	12/18/2024 8:18:21 AM	
[redacted]	[redacted]	ca4af17f-4b25-483c-9d64-655afe0941b8	9/17/2024 10:10:32 AM	

Risk Mitigation Recommendation(s):

- Enforce regular password rotation for user and administrative accounts, including the built-in Domain Administrator
- Remove Password Never Expires and Password Not Required flags except where explicitly justified
- Rotate computer account passwords exceeding recommended age thresholds
- Periodically review privileged and service accounts to ensure credential policies align with security standards

Reference(s):
https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/appendix-d--securing-built-in-administrator-accounts-in-active-directory
https://learn.microsoft.com/en-us/troubleshoot/windows-server/active-directory/password-change-processing-conflict-resolution-function

Affected Endpoint(s):
On-premises Active Directory

Insufficient Endpoint Hardening Controls
MITRE ATT&CK Tactic: TA0001 – Initial Access
Detected: (QoD 75%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
A review of the VMaaS dashboard identified that many recommended Windows security configuration controls, such as enabling strong Kerberos encryption, disabling autorun, restricting anonymous enumeration of shares, Windows Installer with elevated permissions and machine activity limit, are not currently implemented.

The absence of these baseline hardening controls increases the risk of unauthorized access, credential compromise, malware propagation, and data exposure across Windows systems. While no active exploitation was observed, implementing these controls would improve the overall security baseline and reduce the attack surface of the environment.

Risk Detection Result(s):

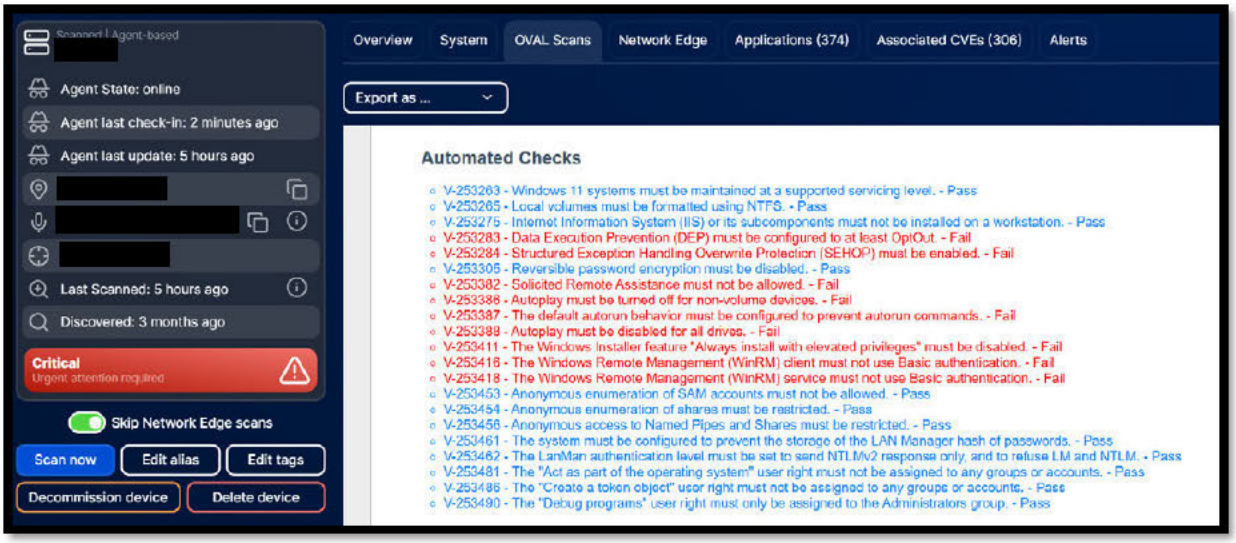

Risk Mitigation Recommendation(s):

- Implement baseline Windows hardening controls via GPO or automation. Some examples include:
- Configure Kerberos encryption.
- Ensure default autorun behavior is configured to prevent autorun commands.
- Ensure solicited remote assistance is not allowed.
- Ensure Windows installer feature "Allows install with elevated privileges" is disabled.
- Ensure anonymous enumeration of shares is restricted.
- Ensure the LanMan authentication level is using NTLMv2 response only and refusing LM and NTLM.

Reference(s):
https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-10/security/threat-protection/security-policy-settings/network-security-configure-encryption-types-allowed-for-kerberos
https://learn.microsoft.com/en-us/windows/win32/shell/autoplay-reg
https://learn.microsoft.com/en-us/windows/client-management/mdm/policy-csp-remoteassistance
https://learn.microsoft.com/en-us/windows/win32/msi/alwaysinstallelevated
https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-10/security/threat-protection/security-policy-settings/network-access-do-not-allow-anonymous-enumeration-of-sam-accounts-and-shares
https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-10/security/threat-protection/security-policy-settings/network-security-lan-manager-authentication-level

Affected Endpoint(s):
Windows Devices

Lack of Kerberos Security Hardening
MITRE ATT&CK Tactic: TA0006 – Credential Access
Detected: (QoD 75%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
Multiple Kerberos-related weaknesses were identified, including an outdated KRBTGT password, legacy encryption types (RC4/DES) enabled on domain controllers, risky delegation configurations, SPNs not supporting AES encryption, and improper write access to resource-based constrained delegation (RBCD) on sensitive accounts.

Kerberos relies on strong encryption and strict delegation boundaries. Legacy encryption enables offline cracking attacks, outdated KRBTGT credentials increase Golden Ticket risk, and misconfigured delegation or RBCD allows attackers to impersonate users or services across the domain.

Risk Detection Result(s):

Result
Found 1 domains where the KRBTGT account password has not changed in the last 180 days.

DistinguishedName	ObjectGUID	Attribute	LastChanged	DaysPassed	EventTimestamp
[redacted]	[redacted]	[redacted]	8/21/2019 7:54:28 PM	2423	8/21/2019 7:54:28 PM

Risk Mitigation Recommendation(s):

- Reset the KRBTGT account password using a two-phase rotation to invalidate existing Kerberos tickets.
- Disable RC4 and DES encryption types on domain controllers and user/service accounts.
- Ensure all users and services with SPNs support AES Kerberos encryption.
- Review and restrict Kerberos delegation configurations to only systems that strictly require it.
- Remove unauthorized write permissions to Resource-Based Constrained Delegation (RBCD), especially on privileged or Tier 0 accounts.
- Regularly audit delegation and SPN configurations.

Reference(s):
https://learn.microsoft.com/en-us/windows-server/security/kerberos/password-change-process
https://learn.microsoft.com/en-us/windows/security/threat-protection/security-policy-settings/network-security-configure-encryption-types-allowed-for-kerberos
https://learn.microsoft.com/en-us/windows-server/security/kerberos/kerberos-delegation-overview

Affected Endpoint(s):
[redacted]

Lack of Conditional Access Policy to Lock Down Global Admin Access
MITRE ATT&CK Tactic: TA0004 – Privilege Escalation
Detected: (QoD 75%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
No conditional access policy was identified to lock down global admin access within [redacted] environment. A lock down global admin access policy refers to a set of security measures designed to restrict and control access to global administrator accounts within an organization.

Risk Detection Result(s):

Policy name	Created by	State	Creation date
Multifactor authentication and reauthentication for risky sign-ins	MICROSOFT	On	8/4/2025, 6:47:52 PM
Multifactor authentication for admins accessing Microsoft Admin Portal	MICROSOFT	Report-only	12/6/2023, 4:01:54 PM
AzureVPN and Apps - Require Hybrid Joined Machine	USER	Report-only	10/8/2024, 8:34:09 PM
Blocked Locations	USER	On	9/27/2023, 3:38:44 PM
Require Hybrid Azure AD joined device	USER	Report-only	1/16/2025, 10:01:12 AM
Require MFA	USER	On	8/8/2023, 8:24:01 PM

Risk Mitigation Recommendation(s):

- Conditional Access: Implementing rules that require specific conditions to be met before access is granted, such as multi-factor authentication (MFA), location-based restrictions, or device compliance checks.
- Least Privilege Principle: Ensuring that global admin accounts have only the necessary permissions to perform their tasks, minimizing the risk of misuse.
- Access Reviews: Regularly reviewing and auditing global admin access to ensure that only authorized personnel have the necessary permissions.
- Emergency Access Accounts: Maintaining separate, highly secure accounts that are only used in emergencies to prevent accidental lockouts.

Reference(s):
https://learn.microsoft.com/en-us/entra/identity/conditional-access/policy-block-by-location
https://learn.microsoft.com/en-us/answers/questions/1003487/all-admins-locked-out-due-to-faulty-conditional-ac
https://answers.microsoft.com/en-us/msoffice/forum/all/hello-is-there-a-way-to-prevent-other-global-admin/e83050e2-4157-4a60-8b21-1315eb752d34

Affected Endpoint(s):
Microsoft Entra ID

Credentials Not Rotated for User, Admin and Privileged Accounts
MITRE ATT&CK Tactic: TA0001 – Initial Access
Detected: (QoD 75%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
Purple Knight identified multiple user, administrative, and privileged accounts with passwords older than 180 days, as well as accounts configured with the password never expires flag. These conditions weaken credential hygiene by allowing passwords to remain unchanged for long periods, making them more susceptible to reuse, historical breach exposure, or password attacks. Traditional Active Directory defaults, such as the 42-day maximum password age, are designed to encourage regular password rotation and reduce credential compromise risk.

Regular rotation of all user and privileged passwords significantly reduces exposure by preventing long-term use of static credentials and aligns with widely accepted best practices for safeguarding identity systems.

Risk Detection Result(s):

DistinguishedName	ObjectGUID	SamAccountName	PasswordLastSet	DaysSinceLastSet	Ignored
[redacted]	1143021d-f89f-4ea0-a462-35e5ce79abb4	[redacted]	4/17/2020 1:03:18 AM	2184	False
[redacted]	c4436e66-3260-42e6-8643-1df673323074	[redacted]	4/8/2024 4:05:56 PM	732	False
[redacted]	dabd7e1a-298b-4b07-9f73-8a7f92eafb3c	[redacted]	1/31/2024 7:56:19 PM	799	False

Risk Mitigation Recommendation(s):

- Conduct recurring password audits to identify old or non-expiring credentials and enforce timely rotation, prioritizing privileged and admin accounts and applying fine-grained password policies where needed.

Reference(s):
https://rsw.io/keeping-active-directory-secure-and-tidy-with-powershell-identifying-stale-accounts/

Affected Endpoint(s):
[redacted]

Privileged Group Membership Governance in AD
MITRE ATT&CK Tactic: TA0004 – Privilege Escalation
Detected: (QoD 75%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
Abacus identified weaknesses in the governance of privileged Active Directory groups. The [redacted] contains members, and there is evidence of temporary (ephemeral) administrative access as well as recent changes to privileged group membership.

While administrative role changes can be operationally necessary, insufficient controls or monitoring around privileged group assignments reduce visibility into who holds elevated access and for how long. Membership in highly privileged groups such as [redacted] grants the ability to make irreversible changes to the directory schema, increasing the potential impact of misuse or compromise.

These conditions highlight the need for stronger controls around privileged role assignment, approval, monitoring, and regular review.

Risk Detection Result(s):

Result
Found 2 members in the Schema Admins group

DistinguishedName	ObjectGUID	GroupDistinguishedName	ForestName
[redacted]	bfc1a387-4ca0-4bf1-a246-a3834ecd7d2f	[redacted]	
[redacted]	e4be5f4e-87c9-4eca-b195-835ca7940db3	[redacted]	

Risk Mitigation Recommendation(s):

- Ensure [redacted] membership is strictly limited and time-bound
- Implement clear approval and documentation requirements for privileged group changes
- Monitor and regularly review privileged group membership and recent changes
- Remove standing or temporary administrative access that is no longer required

Reference(s):
https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/manage/understand-security-group[redacted]
https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/plan/security-best-practices/best-practices-for-securing-active-directory

Affected Endpoint(s):
Administrator, [redacted]

Lack of Data Loss Prevention Controls
MITRE ATT&CK Tactic: TA0010 – Exfiltration
Detected: (QoD 75%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
Data classification controls refer to policies that identify, label, and handle data based on predefined classification categories, enabling organizations to apply appropriate protections to sensitive information such as Personally Identifiable Information (PII) or financial data. These controls are often used in conjunction with Data Loss Prevention (DLP) mechanisms to help prevent unauthorized transmission or exfiltration of sensitive data.

Without adequate DLP controls, sensitive information could be inadvertently or intentionally exposed by a malicious actor or insider, potentially contributing to fraud, data misuse, or reputational impact. During the assessment, an existing DLP policy was identified to be in simulation mode and associated with sensitive M365 Copilot interactions.

Although data classification and DLP enforcement may be partially handled through other third-party platforms, enabling additional DLP policies within this environment, specifically targeting the external transmission of sensitive data types such as PII or financial information, could provide improved visibility and defense-in-depth. These enhancements are considered beneficial but not critical, given existing controls already in place.

Risk Detection Result(s):

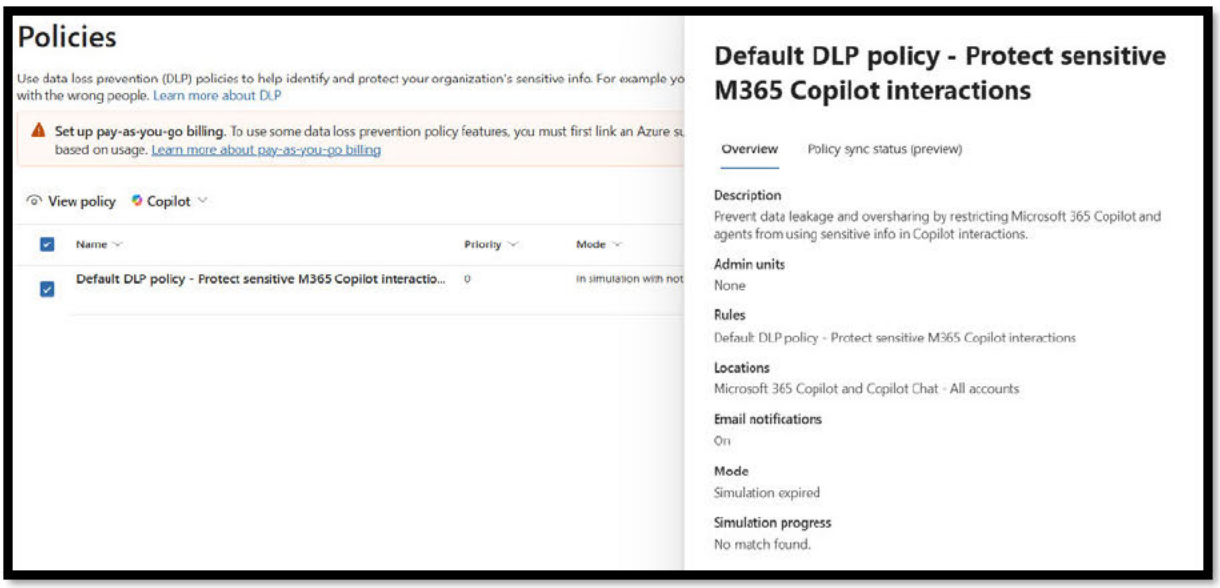

Risk Mitigation Recommendation(s):

- Implement data loss prevention (DLP) policies related specifically to sensitive information being sent externally, such as social security numbers and bank account information.

Reference(s):
https://learn.microsoft.com/en-us/purview/dlp-learn-about-dlp

Affected Endpoint(s):
Microsoft Purview

Insufficient Use of Password Manager
MITRE ATT&CK Tactic: TA0010 – Exfiltration
Detected: (QoD 50%) - **Impact:** Moderate – **Ease**: Moderate – **Risk**: Medium

Summary:
It was discovered in the interview questionnaire that the organization does not use a standard corporate password manager and instead allows employees to use individual password managers. This approach creates inconsistent password practices, limited visibility into weak or reused passwords, and no reliable way to enforce MFA, sharing controls, or emergency access.

It also increases the risk of credential loss when employees leave, as passwords may be stored in personal vaults that the organization cannot access. From a security perspective, this fragmented model expands the attack surface and complicates incident response and forensic investigations. Overall, the absence of an approved enterprise password manager elevates the likelihood of account compromise, data exposure, and audit or compliance findings.

Risk Mitigation Recommendation(s):

- Define and document role-based requirements for password manager usage, ensuring that employees who access corporate systems, applications, or shared credentials are required to use the approved corporate password manager.
- Update internal policies to clearly distinguish between roles that require secure credential storage and those that do not, based on business and operational need.
- Enforce the use of the corporate password manager for in-scope personnel and prohibit insecure password storage practices such as spreadsheets or plain text files.
- Periodically review and reassess password manager adoption to ensure alignment with evolving roles and access requirements.

Reference(s):
https://pages.nist.gov/800-63-3/sp800-63b.html
https://www.keepersecurity.com/blog/2024/09/26/the-risks-of-storing-your-passwords-in-google-sheets/
https://www.keepersecurity.com/resources/spreadsheets-and-sticky-notes-are-insecure.html

Affected Endpoint/System(s):
Organization-wide

Insecure Microsoft SharePoint Guest Access Policy
MITRE ATT&CK Tactic: TA0001 – Initial Access
Detected: (QoD 75%) - **Impact:** Low – **Ease:** Low – **Risk:** Low

Summary:
Abacus identified a security misconfiguration within Microsoft SharePoint. This configuration allows external guests to share items they do not own, which presents a security risk.

- Unauthorized sharing of sensitive information could result in confidential data being shared with unauthorized personnel.
- This setting may contravene various data protection regulations, potentially exposing the organization to legal and financial penalties.
- By enabling external guests to share items freely, the organization inadvertently increases its vulnerability to cyber threats, providing potential attackers with additional avenues to exploit.
- The ability for external users to share items they don't own undermines the organization's control over its data assets, potentially leading to version control issues and data integrity problems.

Risk Detection Result(s):

Risk Mitigation Recommendation(s):

- Consider adjusting the guest access policy within the Microsoft SharePoint admin center (admin.teams.microsoft.com) to allow external communication only to existing guests.
  - This setting can be found by searching "Organizational Settings" then selecting "External Access".

- It is possible that some external guests may be considered internal personnel, and this may be complicated by the difference in primary domain, which has been previously recommended to be addressed.

Reference(s):
https://adminsharepoint.microsoft.com
https://learn.microsoft.com/en-us/microsoft-365/solutions/collaborate-in-site?view=o365-worldwide
https://www.hubsite365.com/en-ww/crm-pages/policies-to-safely-enable-external-sharing-guest-access-in-microsoft-365.htm
https://learn.microsoft.com/en-us/sharepoint/turn-external-sharing-on-or-off

Affected Endpoint(s):
Microsoft SharePoint

Insecure Microsoft Teams External Access Policy
MITRE ATT&CK Tactic: TA0001 – Initial Access
Detected: (QoD 75%) - **Impact:** Moderate – **Ease:** Moderate – **Risk:** Low

Summary:
Abacus detected an insecure setting within Microsoft Teams, which enables users to communicate with all external domains. This could be directly utilized as part of a greater social engineering attack leveraging Microsoft Teams as a direct communication medium.

Malicious actors are increasingly leveraging enterprise communications platforms such as Microsoft Teams for sophisticated phishing campaigns, as it inherently bypasses most anti-phishing security controls associated with email systems.

Risk Detection Result(s):

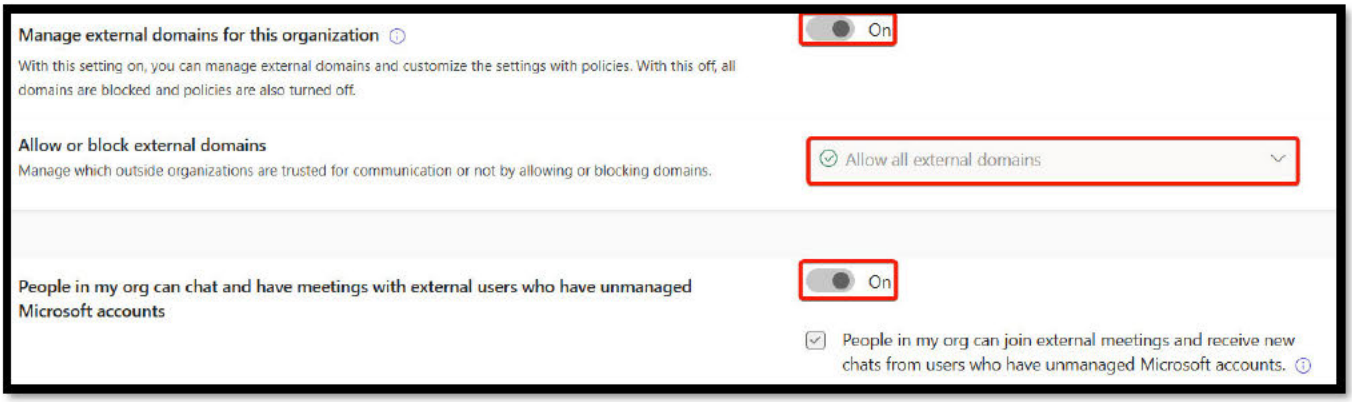

Risk Mitigation Recommendation(s):

- Consider adjusting the external access policy within the Microsoft Teams admin center (admin.teams.microsoft.com) to allow external communication only to specific domains.
  - This setting can be found by searching "Organizational Settings" then selecting "External Access".

Reference(s):
https://admin.teams.microsoft.com
https://learn.microsoft.com/en-us/microsoftteams/communicate-with-users-from-other-organizations
https://labs.jumpsec.com/advisory-idor-in-microsoft-teams-allows-for-external-tenants-to-introduce-malware/

Affected Endpoint(s):
Microsoft Teams

GraphQL Query and Schema Exposure
MITRE ATT&CK Tactic: TA0040 – Impact
Detected: (QoD 75%) - **Impact:** Low – **Ease:** Low – **Risk:** Low

Summary:
Abacus identified the use of a GraphQL API exposed over HTTPS (TCP 443) on [redacted]*com*. Review of the GraphQL endpoint confirmed several GraphQL specific conditions related to query complexity handling and schema exposure, including the ability to submit circular references, alias overloading, field duplication, and directive overloading queries, as well as enabled schema introspection and field suggestion features.

Exploitation of these conditions would require deliberate submission of resource-intensive or intentionally malformed GraphQL queries and does not represent a traditional vulnerability such as authentication bypass, direct data exposure, or remote code execution. No evidence of abuse, excessive query execution, or service instability was observed during testing. Accordingly, these issues are considered low-impact hardening opportunities, rather than indicators of active risk, and may be addressed to reduce unnecessary exposure and improve resilience against misuse.

Risk Detection Result(s):

Risk Mitigation Recommendation(s):
- Implement GraphQL query depth and complexity limits to reduce the potential impact of resource intensive queries
- Restrict or disable GraphQL introspection and field suggestions in production environments where not strictly required
- Apply existing rate limiting or edge protections consistently to GraphQL endpoints
- Periodically review GraphQL security settings as part of application lifecycle management

Reference(s):
https://graphql.org/learn/best-practices/
https://cheatsheetseries.owasp.org/cheatsheets/GraphQL_Cheat_Sheet.html

Affected Endpoint(s):
[redacted]com

Misconfigured SPF Mail Record
MITRE ATT&CK Tactic: TA0001 – Initial Access
Informational

Summary:

[redacted] SPF records were observed to generate warnings from automated validation tools indicating the presence of duplicate or overlapping IPv4 and IPv6 entries. Due to tooling limitations, the specific IP ranges were not visible during review; however, these warnings are typically associated with redundant or improperly formatted netblocks rather than active misconfiguration.

These conditions do not indicate that unauthorized parties are permitted to send email on behalf of the domain, nor do they suggest an immediate spoofing risk. However, unresolved SPF warnings can reduce the effectiveness and clarity of SPF enforcement and may lead to inconsistent evaluation by recipient mail servers. Reviewing and cleaning up duplicate or overlapping SPF entries may help improve long-term email policy maintainability and deliverability.

Risk Detection Result(s):

Risk Mitigation Recommendation(s):

- Review the SPF record for duplicate or overlapping IPv4/IPv6 entries that may be generating validation warnings.
- Remove redundant or unnecessary SPF mechanisms where feasible to improve clarity and maintainability of the record.
- Confirm that any IP ranges included in the SPF policy are still required for legitimate mail-sending services.
- Consider periodically validating SPF configuration using automated tooling to ensure continued alignment with current email infrastructure.

Reference(s):

https://docs.microsoft.com/en-us/exchange/mail-flow-best-practices/how-to-set-up-a-multifunction-device-or-application-to-send-email-using-microsoft-365-or-office-365
https://docs.microsoft.com/en-us/microsoft-365/security/office-365-security/set-up-spf-in-office-365-to-help-prevent-spoofing?view=o365-worldwide

Affected Endpoint(s):

[redacted]com

Previously Breached Credentials Identified
MITRE ATT&CK Tactic: TA0006 – Credential Access
Informational

Summary:
As part of the assessment, eight employee email addresses were identified as having appeared in previously reported third-party data breaches associated with external services, including Betterment, Canva, and DemandScience. These exposures appear to be limited to the referenced third-party platforms and were identified through publicly available breach aggregation datasets.

The presence of these email addresses does not indicate compromise of the [redacted] environment and is presented for awareness purposes only. Such findings are common and primarily serve as a reminder to maintain good credential hygiene, particularly around password reuse across external services.

Risk Detection Result(s):

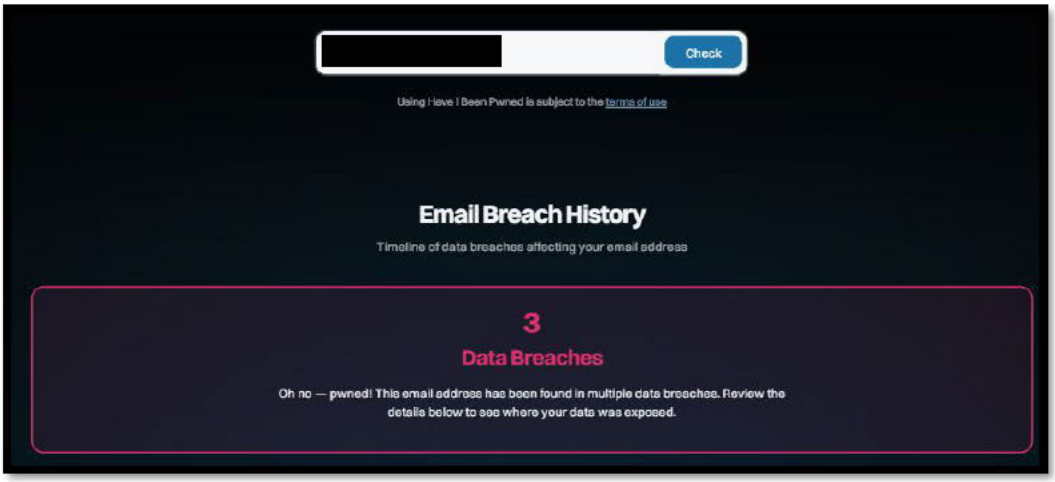

Risk Mitigation Recommendation(s):

- Ensure that identified accounts are not reusing previously breached credentials within the company environment and if so, immediately force password changes.
- Implement an organization-wide password policy that encourages strong passwords and passphrases and discourages password reuse.
- Consider implementing domain breached credential monitoring services for the organization.
- Consider utilizing a corporate password manager for randomly generated secure passwords.

Reference(s):
https://www.proxynova.com/tools/comb/
https://haveibeenpwned.com/

Affected Endpoint(s):

[redacted]

Inadequate Governance of Guest Accounts
MITRE ATT&CK Tactic: TA0001 – Initial Access
Informational

Summary:
Abacus identified that guest invitation settings are broadly permissive, allowing users to invite external guest accounts without restriction. While guest user directory access is appropriately limited and aligns with Microsoft's recommended defaults, unrestricted guest invitations can lead to unmanaged growth of external identities over time.

Additionally, many guest accounts appear to represent internal employees rather than third-party users, reducing overall risk. As a result, this finding does not indicate active exposure but highlights an opportunity to strengthen identity governance by limiting who can initiate guest invitations.

Risk Detection Result(s):

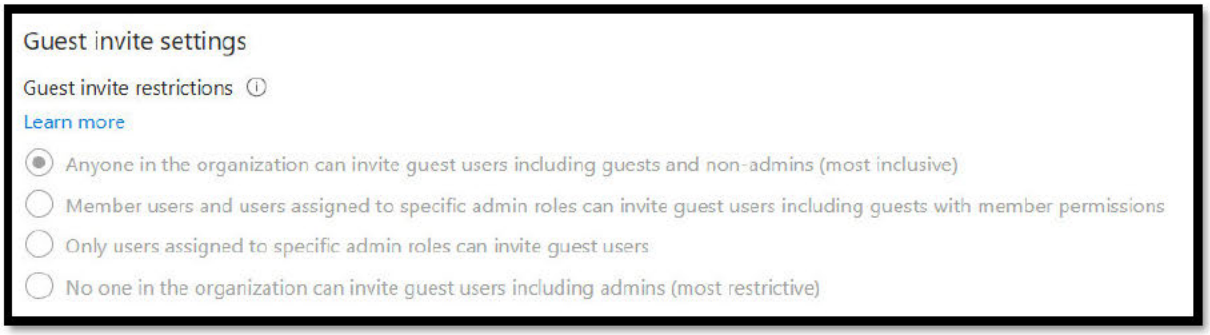

Risk Mitigation Recommendation(s):

- Restrict guest invitations to approved roles (e.g., Guest Inviter) or a defined security group
- Periodically review guest accounts to ensure they continue to have a valid business purpose

Reference(s):
https://learn.microsoft.com/en-us/entra/external-id/external-collaboration-settings-configure

Affected Endpoint(s):
Microsoft Entra ID

Activity Log Alerts Not Set Within Azure
MITRE ATT&CK Tactic: TA0005 – Defense Evasion
Informational

Summary:
During the assessment, Abacus observed that Azure Activity Log Alerts for certain security-relevant configuration changes (such as Network Security Group and policy modifications) were not implemented. While all evaluated actions are logged within Azure's native activity logging, the absence of automated alerts may limit visibility into configuration changes as they occur.

In the context of a small Azure environment, this does not represent an immediate security risk. However, implementing activity log alerting can provide additional assurance and improve security monitoring as the environment evolves or scales.

Risk Detection Result(s):

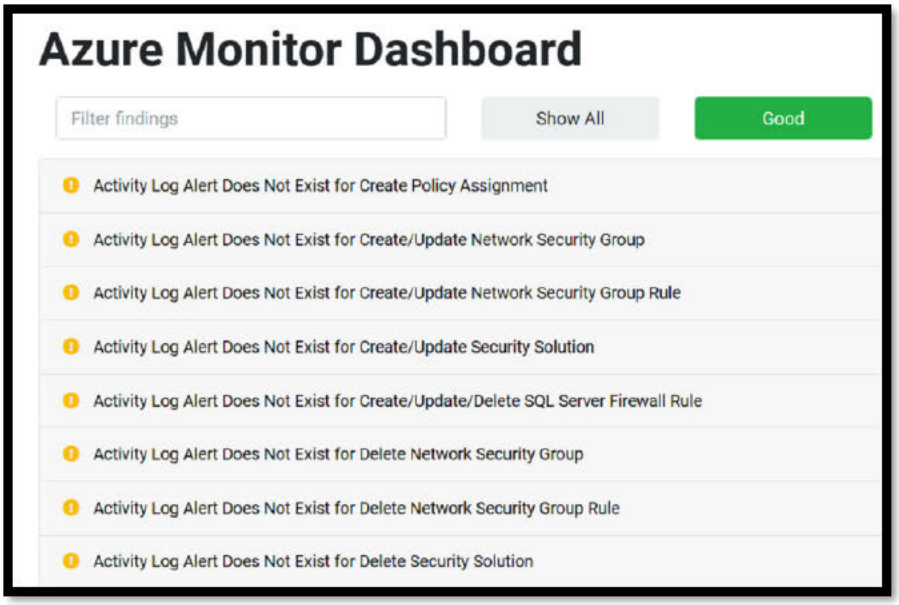

Risk Mitigation Recommendation(s):

- Configure Azure Activity Log Alerts to monitor security-sensitive administrative actions, such as the creation, modification, or deletion of Network Security Groups, NSG rules, and Azure Policy assignments. These alerts can notify administrators when changes occur and help detect unintended or unauthorized configuration changes.
- Associate Activity Log Alerts with an Azure Monitor Action Group (for example, email or webhook notifications) to ensure timely awareness of security-relevant events without requiring continuous manual log review.
- As the environment matures, consider aligning alert coverage with Microsoft's Azure Monitor Baseline Alerts (AMBA) recommendations to establish consistent monitoring for high-impact configuration changes across the subscription.

Reference(s):
https://learn.microsoft.com/en-us/azure/azure-monitor/platform/activity-log?tabs=azure-cli%2Clog-analytics%2Ccsv-azure-cli

Affected Endpoint(s):
Microsoft Azure

Conclusion

While Abacus makes a best effort to rank the severity of vulnerabilities, it is not possible for any assessment to guarantee that low-risk issues cannot be chained together and significantly impact [redacted] systems. All risk ratings in this report are a combination of industry-standard ratings and considerations for unique business risk as identified by Abacus. Special consideration should also be given to the fact that multiple low-risk issues may compound into higher-risk threats.

As a result of this cybersecurity risk assessment, the following was discovered:

- 2 High Risks
- 9 Moderate Risks
- 3 Low Risks
- 4 Informational Risks

Abacus completed the cybersecurity risk assessment, having identified two high findings, nine moderate findings, three low findings, and four informational findings.

Abacus observed that [redacted] overall security posture is generally robust, with many foundational controls in place across identity, infrastructure, and cloud services. The highest-impact findings were limited in number and primarily centered on on-premises Active Directory attack surface reduction, including the presence of the Print Spooler service on domain controllers and non-default principals with DCSync rights. These issues represent well-understood escalation paths that, if chained with credential compromise, could significantly increase the impact of an attack and should be prioritized for remediation. Additional Moderate findings, such as gaps in credential lifecycle management, privileged group governance, Conditional Access coverage for legacy authentication and Global Administrator access, and endpoint hardening controls, reflect opportunities to strengthen consistency and reduce blast radius rather than indications of active compromise.

Lower-severity findings identified across collaboration platforms, guest access governance, application exposure, and email configuration further highlight areas where incremental improvements can enhance overall resilience and monitoring maturity. Addressing these items in a structured, risk-based manner, while prioritizing the most impactful identity and domain security controls, will meaningfully strengthen [redacted] defensive posture and reduce the likelihood that common attack paths could be successfully leveraged in the future.

It is important to note that the risk mitigation recommendations contained in this report reflect industry best practices and Abacus' limited knowledge of [redacted] Abacus does not warrant the compatibility or operational effectiveness of the risk mitigation recommendations provided in this report, as Abacus does not have any understanding of the nuances and caveats of [redacted] network environment. Abacus highly recommends that [redacted] technology department assesses the compatibility and operational effectiveness of the risk mitigation recommendations provided in this report and uses a change management process to validate, plan and implement remediation changes in a UAT environment first.

Appendix A - MITRE ATT&CK Tactic Definitions

TA0043 – Reconnaissance
Reconnaissance consists of techniques that involve adversaries actively or passively gathering information that can be used to support targeting. Such information may include details of the victim organization, infrastructure, or staff/personnel. This information can be leveraged by the adversary to aid in other phases of the adversary lifecycle, such as using gathered information to plan and execute Initial Access, to scope and prioritize post-compromise objectives, or to drive and lead further Reconnaissance efforts.

TA0042 – Resource Development
Resource Development consists of techniques that involve adversaries creating, purchasing, or compromising/stealing resources that can be used to support targeting. Such resources include infrastructure, accounts, or capabilities. These resources can be leveraged by the adversary to aid in other phases of the adversary lifecycle, such as using purchased domains to support Command and Control, email accounts for phishing as a part of Initial Access or stealing code signing certificates to help with Defense Evasion.

TA0001 – Initial Access
Initial Access consists of techniques that use various entry vectors to gain their initial foothold within a network. Techniques used to gain a foothold include targeted spear phishing and exploiting weaknesses on public-facing web servers. Footholds gained through initial access may allow for continued access, like valid accounts and use of external remote services, or may be limited-use due to changing passwords.

TA0002 – Execution
Execution consists of techniques that result in adversary-controlled code running on a local or remote system. Techniques that run malicious code are often paired with techniques from all other tactics to achieve broader goals, like exploring a network or stealing data. For example, an adversary might use a remote access tool to run a PowerShell script that does Remote System Discovery.

TA0003 – Persistence
Persistence consists of techniques that adversaries use to keep access to systems across restarts, changed credentials, and other interruptions that could cut off their access. Techniques used for persistence include any access, action, or configuration changes that let them maintain their foothold on systems, such as replacing or hijacking legitimate code or adding startup code.

TA0004 – Privilege Escalation
Privilege Escalation consists of techniques that adversaries use to gain higher-level permissions on a system or network. Adversaries can often enter and explore a network with unprivileged access but require elevated permissions to follow through on their objectives. Common approaches are to take advantage of system weaknesses, misconfigurations, and vulnerabilities.

TA0005 – Defense Evasion
Defense Evasion consists of techniques that adversaries use to avoid detection throughout their compromise. Techniques used for defense evasion include uninstalling/disabling security software or obfuscating/encrypting data and scripts. Adversaries also leverage and abuse trusted processes to hide and masquerade their malware. Other tactics' techniques are cross listed here when those techniques include the added benefit of subverting defenses.

TA0006 – Credential Access
Credential Access consists of techniques for stealing credentials like account names and passwords. Techniques used to get credentials include keylogging or credential dumping. Using legitimate credentials can give adversaries access to systems, make them harder to detect, and provide the opportunity to create more accounts to help achieve their goals.

TA0007 – Discovery
Discovery consists of techniques an adversary may use to gain knowledge about the system and internal network. These techniques help adversaries observe the environment and orient themselves before deciding how to act. They also allow adversaries to explore what they can control and what's around their entry point in order to discover how it could benefit their current objective. Native operating system tools are often used toward this post-compromise information-gathering objective.

TA0008 – Lateral Movement
Lateral Movement consists of techniques that adversaries use to enter and control remote systems on a network. Following through on their primary objective often requires exploring the network to find their target and subsequently gaining access to it. Reaching their objective often involves pivoting through multiple systems and accounts to gain. Adversaries might install their own

remote access tools to accomplish Lateral Movement or use legitimate credentials with native network and operating system tools, which may be stealthier.

TA0009 – Collection
Collection consists of techniques adversaries may use to gather information and the sources information is collected from that are relevant to following through on the adversary's objectives. Frequently, the next goal after collecting data is to steal (exfiltrate) the data. Common target sources include various drive types, browsers, audio, video, and email. Common collection methods include capturing screenshots and keyboard input.

TA0011 – Command and Control
Command and Control consists of techniques that adversaries may use to communicate with systems under their control within a victim network. Adversaries commonly attempt to mimic normal, expected traffic to avoid detection. There are many ways an adversary can establish command and control with various levels of stealth depending on the victim's network structure and defenses.

TA0010 – Exfiltration
Exfiltration consists of techniques that adversaries may use to steal data from your network. Once they've collected data, adversaries often package it to avoid detection while removing it. This can include compression and encryption. Techniques for getting data out of a target network typically include transferring it over their command-and-control channel or an alternate channel and may also include putting size limits on the transmission.

TA0040 – Impact
Impact consists of techniques that adversaries use to disrupt availability or compromise integrity by manipulating business and operational processes. Techniques used for impact can include destroying or tampering with data. In some cases, business processes can look fine, but may have been altered to benefit the adversaries' goals. These techniques might be used by adversaries to follow through on their end goal or to provide cover for a confidentiality breach.

Appendix B - Microsoft 365 Security Resources

Given that a significant amount of [redacted] systems is within the Microsoft ecosystem, measuring the efficacy of compliance with company policies would most likely be best accomplished with Microsoft Purview, which is bundled with Business Premium licenses. Microsoft Purview has the ability to ingest and evaluate the configurations of Office 365 services, Defender services, Intune, Secure Score, and Azure systems.

It would be ideal if the Microsoft 365 Compliance CIS assessment dashboards were reviewed on at least a quarterly cadence to ensure a consistent burn-down of security control gaps while also ensuring that new gaps do not appear from unintended regressions. CIS-specific references for security control implementation recommendations and security hardening are also provided.

Specific Resources:

Microsoft 365 Purview Compliance Manager helps automatically evaluate and manage your organization's compliance requirements. There are multiple compliance checklist templates that are provided, including templates for CIS, which can also be modified and customized.
Microsoft Compliance Configuration Analyzer for Purview is a preview PowerShell-based utility that will fetch your organization's current configurations and validate them against Microsoft 365 recommended best practices. MCCA can help you quickly see which improvement actions in Purview apply to your current Microsoft 365 environment. Each action identified by MCCA will give you recommendations for implementation, with direct links to Purview and the applicable solution to start taking corrective action.
CIS Microsoft Azure Foundations Benchmark provides a step-by-step checklist for securing Azure.
CIS Hardened Images on Microsoft Azure are Azure certified and preconfigured to the security recommendations of the CIS Benchmarks. They are available on both Azure and Azure Government.
Azure Blueprint for CIS Microsoft Azure Foundations Benchmark helps customers deploy a core set of policies for any Azure-based architecture that must implement CIS Azure Foundations Benchmark recommendations.
Azure Policy recommendation mapping provides details on policy definitions included within the above Blueprint and how these policy definitions map to the compliance domains and controls in CIS Microsoft Azure Foundations Benchmark. When assigned to an architecture, resources are evaluated by Azure Policy for non-compliance with assigned policy definitions.
CIS Controls Cloud Companion Guide provides guidance on applying security best practices in CIS Controls Version 7 to cloud environments.
CIS Microsoft 365 Foundations Benchmark provides prescriptive guidance for establishing a secure baseline configuration for Microsoft 365.
Microsoft 365 security roadmap Minimize the potential of a data breach or compromised account by following this roadmap.
Windows security baselines Follow these guidelines for the effective use of security baselines in your organization.

Appendix C - Risks Connected to AI Usage

Before offering guidance for the use of AI, it is important to communicate the risks present in the use of these tools. While this list is not exhaustive, it is generally considerate of the higher risks present in the current landscape of AI usage.

Unauthorized or Unintentional Data Exposure
AI usage can expose firms to having confidential data leaked to outside parties or unauthorized internal parties if appropriate protections are not in place. Without de-authorization, your data may be used to train AI models and lead to potential leakage outside your organization. Further, prompt injection techniques may be used to forcibly expose your data connected to these AI systems.

Regulatory, Legal, and Compliance Risks
The use of AI is growing in regulatory oversight. The EU AI Act is strictly enforced for all manners of AI usage and processing and carries penalties akin to those under the GDPR. In the same vein, personal data used in AI systems may need to be communicated to individuals as part of required privacy notices. Leveraging AI without prior authorization may also expose firms to violating existing contracts, which require confidentiality of shared data.

Hallucinations
AI is an imperfect tool and reliant on a set of growing, but limited data references. As such, asking highly specific questions to which there is no appropriate reference backing may lead the AI solutions to generate arbitrary and incorrect answers.

Insecure Code Development
In 2025, OWASP identified inappropriate trust in AI generated code (colloquially called "vibe coding") as introducing significant security risk in newly developed applications and programs. For internally used programs the risk may be lower, but more public-facing or highly sensitive applications developed using this method are inherently very dangerous.

Undiscovered Risks
Just as there are newly discovered security exploits made every day, so too are risks stemming from AI systems. It is important to always remain informed and understand the risk posture of the solution you are seeking to implement and weigh those findings against the risk appetite for your organization.

Appendix D - Countering AI-centric Risks

It should be made clear that in any system, operation, or action, there is no such thing as no risk - there are inherent risks (risks present from the outset) and actual/residual risks (risks remaining following steps taken to mitigate the risk). That said, there are a series of recommended practices that your organization should adopt and confirm as adequate *before* implementing any new AI system.

Technical Controls
Technical controls will range in ease of implementation depending on the degree to which you are mitigating risks, and the level of licensing used for the AI tool. The following are broad recommendations for technical policies to implement and are vendor-agnostic.

- Disable the allowance of the tool to train on your data
- Minimize retention periods of data or disable “memories” outright
- Enable SSO/SAML authentication where possible
- Enable audit logging within the solution - ideally connect them to your Security Information and Event Management (SIEM) solution
- Disable or highly restrict any plugins, connectors or integrations with other systems (e.g., your Microsoft environment, your cloud storage solution)
- Enforce domain restrictions to only authorize your organization’s domain within the tenant
- Set automatic log-off for idle sessions, ideally within 15 minutes (this is different from the automatic device log-off)
- Disable external sharing of AI conversations outside of the organization
- Restrict user creation of API keys
- Follow the principles of least privilege and role-based access controls when assigning licenses and permissions
- Under no circumstances should AI tools be authorized without organization-sponsored licensing. Personal licenses (both free and advanced licenses) within an organization should be prohibited in all cases

It is important to perform due diligence and review documentation to ensure your organization has the appropriate licensing to enact the listed controls.

A note on Data Loss Prevention (DLP) tooling
DLP tools like Microsoft Purview and Egnyte are becoming increasingly popular in the IT security landscape. These tools help prevent data exposure by automatically labeling data and setting access controls to prevent unauthorized exposure to both external and internal parties.
If you are seeking to leverage an AI tool to process, review and understand confidential or sensitive data you have, even if it is internally facing data, it is of utmost importance that you integrate a DLP solution with your authorized AI tool.

Administrative Controls
While the technical controls are rather straightforward in approach, the administrative controls are where most organizations falter. The administrative controls are not simply policies to allow or disallow the use of AI, but should be comprehensive to properly identify and address the risks of AI usage.
Before seeking to implement an AI solution, it is important to set very clear criteria and inputs for your overall operations. These include:

- A defined data classification schema to understand what data you have in tow and the sensitivity of such data
- Risk management considerations to identify, document, and address risks that may arise from the use of AI systems - with these considerations being proactive and ongoing
- Defined governance structure to determine who owns and is responsible for enforcement of all security controls related to AI
- Defining approved use cases for AI solutions and what data is authorized for those use cases
- Due diligence on the interested solutions, both before onboarding and on an ongoing basis to identify incidents and security failures for the tool's infrastructure
    - Note: Standard due diligence is usually related to the tool provider and not an audit of the tool itself. Individual configurations of tools should be considered with respect to the previously listed technical controls.

- Identifying regulatory, legal, and compliance risks tied to the implementation of an AI solution
- Monitoring for potential incidents stemming from an organization's use of AI, as well as non-compliance with organizational policies
    - Note: The monitoring of security incidents related to AI may be considered separate to broad security monitoring provided by a security service provider (e.g., a SOC). Mechanisms to monitor and respond to AI-centric incidents may fall more on internal teams and introduce new risks.

- Creating or amending Secure Software Development Lifecycle (SSDLC) practices to reflect the accepted uses of AI for coding/programming
- Training staff in acceptable practices and ensuring they understand their responsibilities related to an AI system

This list may seem quite long but can realistically be incorporated into existing Security Policies like a WISP. The more difficult part organizations experience is ensuring enforcement of these policies and procedures since those deficiencies may come up during an audit.

Furthermore, these recommendations are not to be considered "one size fits all" and are broad strokes for policy coverage. Procedural elements and the depth of policies should be specific to each organization, their operations, and the degree to which they can be enforced.

Appendix E - Purple Knight Findings

Purple Knight - Active Directory Findings

Non-default principals with DC Sync rights on the domain	[redacted]
Unprivileged users can add computer accounts to the domain	[redacted]
Built-in domain administrator account with old password (180 days)	Administrator
Ephemeral admins	Administrators
Schema Admins group is not empty IOE Found 64	[redacted] Administrator
Users with old passwords	[redacted]
Admins with old passwords	[redacted] N-able Service Account
Risky Pre-Windows 2000 Compatible Access Group membership IOE Found 91	Pre-Windows 200 Compatible Access
Changes to privileged group membership in the last 7 days	Administrators
Protected Users group not in use	Administrator, [redacted] N-able
Users with Password Never Expires flag set	[redacted]
Privileged accounts with a password that never expires	[redacted] N-Able Service Account
User accounts with password not required	[redacted]
Computers with password last set over 90 days ago	[redacted]
Print spooler service is enabled on a DC	[redacted].[redacted]local
Kerberos [redacted] account with old password	[redacted]
Users with SPN defined	[redacted]
Primary users with SPN not supporting AES encryption on Kerberos	[redacted]
RC4 or DES encryption type are supported by DCs	[redacted].[redacted]local
SSO computer account with password last set over 90 days ago	[redacted]

RC4 or DES encryption is supported by SSO computer account	[redacted]

Purple Knight – Entra ID Findings

Certificate-based authentication persistence exposure	[redacted]
[redacted] units are not being used	
Guests with permission to invite other guests	
Users or devices inactive for over 90 days	[redacted]
[redacted] authentication is allowed	
Users with weak or no MFA	
Custom banned password protection not in use	
Conditional Access policy that disables admin token persistence	
Guest users are not restricted	
Entra tenant is susceptible to Hidden Consent Grant attack	[redacted]
Guest accounts inactive over 30 days	[redacted]
Permanently active privileged role assignment	[redacted]
Privileged accounts with mailbox	[redacted]
Report suspicious MFA activity disabled	
Conditional Access policy that does not require MFA when sign-in risk has been identified	

List of Risky users (medium or high level)	[redacted]
Guest invites not accepted in last 30 days	[redacted]
Non-privileged users cannot complete MFA	[redacted]
Unresolved Entra ID privileged role members	Object not found
AD privileged users that are synched to Entra ID	[redacted]