LMNTRIX
BE THE HUNTER | NOT THE PREY

IDENTITY **PROTECTION** IN A DIGITAL **WORLD**

A Comprehensive Guide to Safeguarding Digital Identities and Active Directory

WHITEPAPER 2023

LMNTRIX USA
333City Blvd West
18th Floor, Suite
1805 Orange,
CA 92868
+1.888.958.4555

LMNTRIX UK
200 Brook Drive,
Green Park
Reading, RG2 6UB
+44.808.164.9442

LMNTRIX AUSTRALIA
Level 32, 101 Miller
Street, North Sydney
NSW 2060
+61.288.805.198

lmntrix.com

LMNTRIX SINGAPORE
60 Kaki Bukit Place
#05-19
Eunos TechPark
+65 3159 0639

LMNTRIX HONG KONG
14F, Manning House,
38-48 Queen's Road
Central , Hong Kong
+852.580.885.33

LMNTRIX INDIA
VR Bengaluru, Level 5
ITPL Main Rd,
Devasandra Industrial
Estate Bengaluru,
Karnataka v560048
Email: sales@lmntrix.com
+91-22-49712788

EXECUTIVE **SUMMARY**

Organizations manage access to digital resources using identity-based authentication and authorization processes, where every user, device, and application that accesses resources has its unique digital identity. The explosion in the use of intelligent devices and the move to cloud-based services has dramatically increased the number of digital identities that organizations manage.

Using digital identities to manage and control access to resources has made these a principal target for attackers. Compromising just one identity will allow an attacker to access systems and gain a foothold in a manner that traditional perimeter-based security controls cannot detect. While observing a low-skilled attacker using a stolen password through unusual behavior patterns is possible, sophisticated attackers will have the skills to imitate typical user actions to avoid detection by traditional controls.

Identity protection solutions minimize the risks of identity-related attacks using specialist processes to detect and respond, integrating with Identity and Access Management processes and other security controls to provide comprehensive security coverage. Identity protection revolves around measures to confirm identities and implement secure authentication and session management processes to protect against identity-related attacks and misuse of compromised credentials.

Identity protection solutions achieve protective results by correlating identification data with indicators of attack and compromise from traditional security controls such as endpoint and network monitoring, behavioral and intelligence-led analysis with other logging and telemetry data. This integrated approach gives security analysts enhanced visibility across an attacker's kill chain. Combining identity threat detection information into cohesive threat data offers faster, automatic threat detection and prioritization, supports faster intelligence-led investigation, and allows automated remediation processes.

The critical challenge for identity protection solutions is the handling and analyzing vast quantities of identity-related data and correlating this against other system information to spot the signs of system compromise to detect attacks as early as possible within the attack lifecycle to minimize impact.

This paper looks at how identity management and protection solutions operate and how LMNTRIX leverages this technology to deliver enhanced security protection to our clients.

CONTENTS

INTRODUCTION TO IDENTITY PROTECTION

INTRODUCTION TO DIGITAL IDENTITIES

A verified and trusted digital identity is essential for completing any type of transaction within an information system and interacting across the boundaries of interconnected systems. This applies to users, devices, and applications, all requiring a means to identify themselves with each other on an as-required basis legitimately.

Verifiable and trustworthy digital identities are also essential for implementing Zero Trust architectures to counter the increasingly sophisticated and persistent cybersecurity threats that organizations face daily. A Zero Trust philosophy has no assumption of authenticated trust within system boundaries, requiring identities to be verified and trust earned using the approach of assuming transactions are suspect until proven legitimate using robust authentication techniques irrespective of whether the transaction originates inside or outside system boundaries. This approach makes Zero Trust ideal for securing connections from anywhere cloud-based architectures, and the Zero Trust policy is under adoption across US government departments as the best practice for countering the increasing risks in the current threat landscape. We discuss Zero Trust in more detail later in this white paper if you want to learn more about this access management approach to securing systems.

One of the leading security challenges for managing digital identities for users is the proliferation of different identities for distinct systems creating fragmented identities that are difficult to reconcile and maintain and often incompatible across diverse systems. For users handling multiple identities, the temptation to use easily remembered weak credentials or reuse the same credentials across various systems can undermine the integrity of the authentication processes. Service providers also prioritize service availability over security to make it easier for users who have forgotten their credentials to gain access using techniques such as password resets using unencrypted email, creating opportunities for malicious exploitation.

Diverse and incompatible techniques and technologies for verifying, authenticating, and managing digital identities create opportunities for attackers to exploit weaknesses and vulnerabilities to steal legitimate identity information or trick a system into validating an illegitimate identity to gain unauthorized access.

Another regulatory challenge is the divergence of personal data protection legislation worldwide which can impact the operational deployment of digital identity verification and authentication processes across national boundaries.

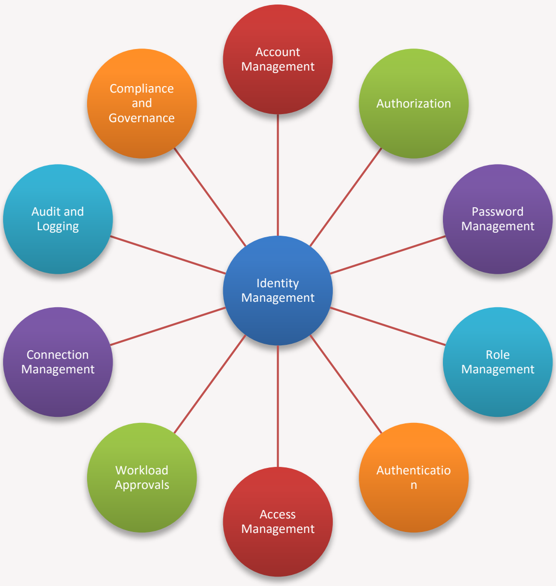

Figure 1 – Identity Management Complexity

Digital identities are the technological means of implementing trust-based methods for authorizing the provision of goods and services between parties, covering everything from providing users and applications with access to data and services to approving the transfer of physical goods, funds, and other tangible products between organizations and individuals.

Physical identification credentials such as a passport were designed originally to enable face-to-face transactions such as movement across borders to proof of entitlement to service. The credentials provide the means to prove a person is who they claim to be using attributes such as a description or photograph. The digital age requires an alternate means of proving identity that replaces face-to-face transactions with electronic interactions. This remote authentication process requires electronic attributes such as biometric information or password data to replace visual confirmation techniques for user identity.

Digital identity techniques have since evolved to include other types of entities such as devices, applications, and services in addition to users to allow systems to verify and authenticate the identity of all entities that provide and consume data in an information system.

A digital identity is now classed as the unique representation of an entity engaged in an online transaction in the context of a specific digital service. The digital identity does not necessarily uniquely identify the entity in all contexts or indeed reveal the real-life identity of the entity.

The challenge for organizations is ensuring digital identification attributes' confidentiality, integrity, and availability to protect their systems and entities. Any compromise of the confidentiality of a digital identity would allow an attacker to steal and reuse the identity attributes for malicious purposes. Any compromise of the integrity or availability of a digital identity would inhibit authentication and deny services to the affected entity.

The proliferation of interconnected devices exacerbates this challenge through the widespread adoption of intelligent machines with the Internet of Things and their role in the Industry 4.0 revolution. There are currently tens of billions of connected devices worldwide, each with a unique digital identity used to authenticate access to systems and services. The number and complexity of these interconnected entities performing digital transactions will inevitably increase going forward. This scaling will only increase the criticality of identity authentication processes to support trusted and secure transactions between entities through these devices.

The Better Cloud surveys of SaaS adoption revealed that a typical organization now relies on around 130 different cloud applications, each with its own digital identity-based authentication and authorization process. This is a significant increase from the average of 12 applications just six years earlier.

Organizations using digital identities to manage access controls must balance the demands of an efficient and scalable solution with security, governance, and data privacy constraints. As these organizations embrace digital transformation and adopt hybrid and remote work practices, the challenges will only increase in ensuring access to resources, including information, applications, and services, are available to the right people without compromising security.

Digital access rights define the need to know the entitlement of the entity requesting access. They should operate on the least privilege principle to limit rights to only the resources the entity legitimately requires when performing its authorized functions. This places a requirement on the organization to establish a comprehensive, consistent, and effective classification policy across all its systems to manage access. Information security integrity directly depends on implementing effective identity and access management controls.

A risk-based approach to managing digital identity access rights can ensure highly scaled systems remain manageable. Examples of high risks include:

- Privileged entity accounts with elevated access permissions
- Entity accounts requesting unusual or abnormal accesses
- Entity accounts with active policy violations
- Entity accounts with pending remediation actions
- Orphaned entity accounts
- Entity accounts with aged credentials

Over 80% of data breaches involve the compromise of entity credentials, which allow an attacker to bypass simple event-driven security controls by masquerading as a legitimate user. Sophisticated attackers can use compromised access credentials using techniques that make it difficult for security controls to differentiate malicious actions from legitimate user activities.Recent changes to work practices, including greater use of remote and hybrid working along with the proliferation of connected devices, have increased the attack surface available for advanced persistent threats and made it more challenging to discern abnormal user actions using traditional behavioral monitoring and analysis techniques.

The problem with standard password-based access controls is that most security incidents can be traced back to using weak, easy-to-guess passwords or compromised services that hardcode passwords or store them with little or no protection. Hundreds of millions of stolen login details with plain text passwords are circulating the deep and dark Web. A glance at a sample will reveal just how many users still use "password" or "12345678" despite years of warnings by security professionals.

OVERVIEW OF IDENTITY PROTECTION

Identity protection solutions protect the digital identities of all entities within an organization's systems. This includes all users, devices, applications, and services that need to authenticate their identity to gain access to information assets that the organization is responsible for protecting, whether they reside in on-premises or cloud-based infrastructure, irrespective of if it is owned, operated, and managed by the organization. Identity protection differs in philosophy from some of the more typical security controls that it includes protecting the digital identities of users of the system who operate outside of the system as consumers, suppliers, customers, or indirectly connected entities along the supply chain and the compromise of such third-party digital identities within an organization's infrastructure can typically allow an attacker to exploit that compromised identity outside of the organization.

An identity protection solution aims to minimize the risk of an identity-related security breach and provide the processes to detect and respond to any such breach. However, to be effective, it must integrate with the organization's Identity and Access Management (IAM) processes and other security controls to provide comprehensive security coverage.

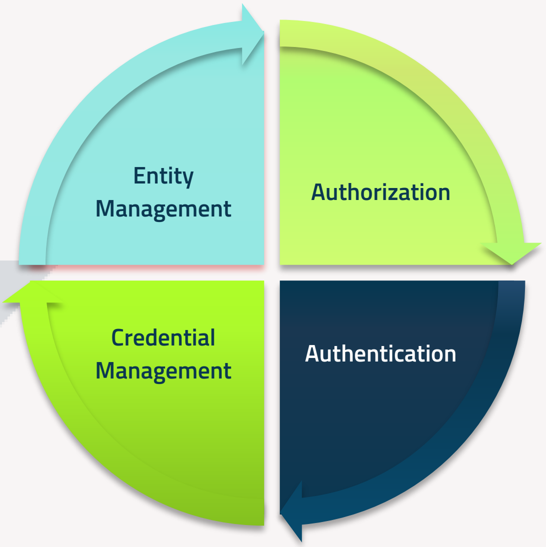

Figure 2 – Identity Management Complexity

Standalone identity protection solutions can overwhelm security analysis and SOC staff with credential-based alerts that require correlation with behavioral analytics and event-based data creating a significant investigation overhead with the potential to increase the risk of alert fatigue. A lack of integration will also limit detection and response process automation with a corresponding increase in process times.

IAM processes support digital identity and associated access rights management to implement access controls for systems and services to implement risk-based controls using the least privilege philosophy as part of an integrated security solution. They also facilitate productivity efficiencies by implementing single sign-on (SSO) and multifactor authentication (MFA) processes. However, IAM processes cannot detect the exploitation of compromised credentials.

Implementing effective identity protection requires a systematic approach beginning with a risk-based evaluation of the systems, entities, and information assets within the scope of the identity protection solution. The goal is to define what elements require protection from what credible threats, quantifying the likelihood of threats materializing and qualifying the impact of occurrence to assess risks. Those risks that exceed the risk appetite for the business can then be prioritized for applying risk mitigation measures to reduce the levels to an acceptable level through technological or procedural controls.

Security controls include implementing robust authentication processes, including MFA techniques, or using biometrics to replace memorable data where possible. Other measures include enhancing and expanding data encryption across systems and using secure communications channels to protect information at rest and in transit.

IAM configuration should follow best practices, including least privilege principles, non-persistent role-based privilege allocation, and segregation of permissions for multi-role and elevated permissions accounts. An example of the former is that local administrator privileges needed to configure peripherals such as printers should only be enabled on an as-required basis. An example of the latter is that no single account should be permitted to configure a business-critical process and initiate its operation, even if a single user performs the two steps.

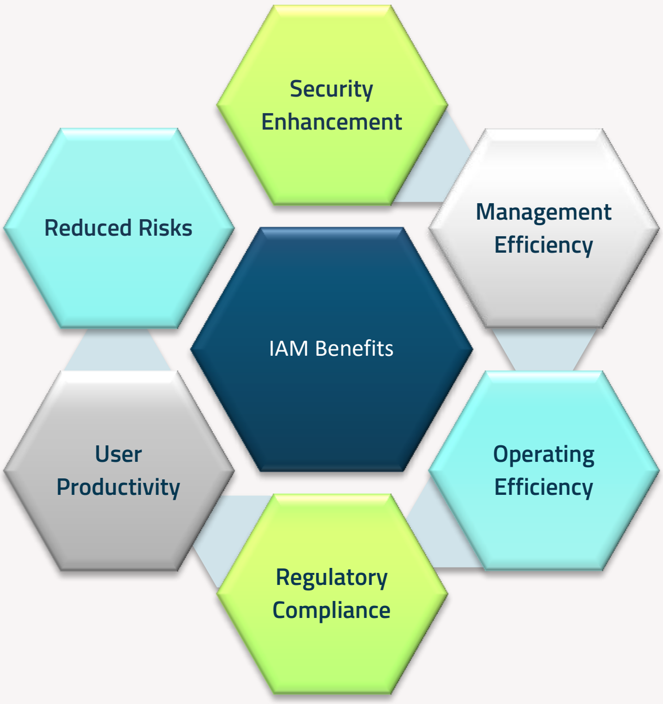

Figure 3 - Identity and Access Management Benefits

Information asset protection should be regularly audited against baseline information security standards and relevant regulations and legislation to ensure compliance is established and maintained over the life of the information.

Threat detection and response services should integrate the IAM solution with extended detection and response to provide comprehensive protection and link with incident response and recovery processes in an end-to-end solution.

Users should undergo appropriate training before being granted permissions and periodic awareness training to maintain currency and respond to changes in the systems or the threat landscape they reside within.

IDENTITY PROTECTION BEST PRACTICES

Identity protection best practices revolve around measures to confirm identities and implement secure authentication and session management processes to protect against identity-related attacks. These best practices include the following:

- Disabling or changing all default credentials before a service or application goes live, particularly for accounts and entities with privileged access rights.
- Implementing a comprehensive password policy with automated checks to enforce length and complexity rules and prevent the use of weak access credentials such as known compromised, easily guessable, reused, or commonly used passwords.

- The implementation of multifactor authentication to prevent automated credential stuffing, brute force attacks, and the reuse of stolen credentials.
- Implementation of automated processes to log and manage failed login attempts with secure auto-remediation options for non-malicious failures and security alerting for detected credential stuffing, brute force, or other account authorization process attacks.
- Hardening all registration, credential recovery, and API pathways against account enumeration attacks using the same messages for all outcomes.
- Implementation of secure server-side session management processes that automatically generate all new session ID with random, high entropy values to prevent brute force attacks that shares and stores these session IDs securely.
- Implementation of session management processes that automatically invalidate session IDs after a set maximum period, after logout, and following defined periods of zero account activity.

Verify Trust
Least Privilege Principles
Central Management and Control
Complete Visibility

Figure 4 - Identity Protection Best Practices

IDENTITY PROTECTION RISK

The reliance on digital identities for managing access to resources has made these a principal target for attackers. The latest Verizon data breach reports estimate that more than 80 percent of breaches now involve identity-based attack techniques.

Compromising a single account for a corporate entity can allow an attacker to gain access to large areas of infrastructure, bypassing security controls to search out opportunities to elevate privileges and move laterally to gain greater permission until they achieve sufficient access to complete their attack. The greater the range of digital identities, the increased probability of exploitable weaknesses allowing compromise. Sophisticated attack techniques allow compromise of credentials by exploiting weaknesses and vulnerabilities in technology, personnel, and processes rather than the traditional approach of brute force password cracking techniques.

Identity protection solutions correlate identification data with indicators of attack and compromise from traditional security controls such as endpoint and network monitoring, behavioral and intelligence-led analysis with other logging and telemetry data. This integrated approach gives security analysts enhanced visibility to get a comprehensive view of the attacker's kill chain. Combining identity threat detection information into cohesive threat data offers faster, automatic threat detection and prioritization, supports faster intelligence-led investigation, and allows automated remediation processes.

Identity protection solutions offer organizations a range of benefits when implemented as part of an integrated security solution. The key benefits realized from real-time, continuous visibility of authentication data and processes include the following:

- Support the timely detection and response to advanced identity-based threats, including supply chain attacks and ransomware infection, with real-time automated monitoring and alerting processes integrated with an extended detection and response capability.
- Enhance detection of exploitation of credentials, including lateral movement and privilege escalation, with end-to-end visibility across system touchpoints and boundaries and enforcement of authentication policies.
- Enhance the monitoring of privileged accounts to counter high-risk threats from exploiting compromised credentials with context-based incident prioritization to minimize response times.
- Improve the robustness of identity and credential repository protection controls with continuous monitoring and IAM integration for automating auditing and governance processes.
- Provide security teams with enhanced continuous end-to-end visibility of identity credential processes in hybrid working environments and across cloud-based infrastructure.
- Enhance traditional AD security by enforcing conditional access policies that manage identity-based risks based on entity behavior, context, and risk across system boundaries using coherent organization-wide technology- and system-agnostic procedures.
- Support successfully implementing a comprehensive Zero Trust policy across systems and boundaries.

INTRODUCTION TO ACTIVE DIRECTORY

OVERVIEW OF ACTIVE DIRECTORY

Active Directory (AD) services underpin systems by providing a data store-based repository for object-based access policies as part of objects' logical and hierarchical records. AD manages the authentication of entities requesting access to an object indexed through its records and handles the access authorization process. AD domain services extend the AD processes across networks by providing a domain controller to manage the access credentials and implement permissions-based authentication and authorization processes.

The critical security benefit of AD technology is the centralized access rights and user account management that allows configuration, distribution, and enforcement of authentication and authorization policies in network-based environments. AD also enables users to authenticate once to access multiple resources using SSO processes where business productivity requirements outweigh robust security needs such as a Zero Trust policy. AD also supports central data storage to simplify information sharing and collaborative work processes and support backup and recovery processes.

The issue with AD is that any compromise of access credentials can undermine the complete AD infrastructure allowing an attacker to escalate the attack using data compromise, privilege escalation, and lateral movement techniques. AD makes critical data useful for attackers, such as identity records, credentials, and configuration data, easier to find. AD also uses token-based technology to manage authorization keys to control resource access across systems, creating a significant security risk when tokens are compromised.

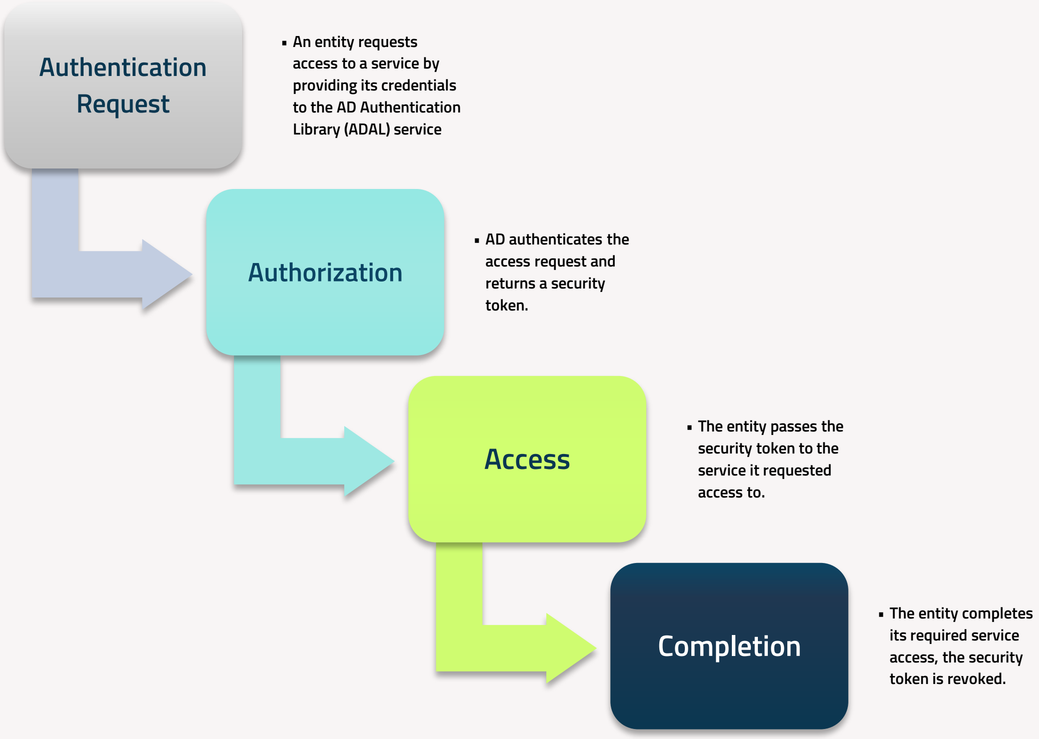

- An entity requests access to a service by providing its credentials to the AD Authentication Library (ADAL) service
- AD authenticates the access request and returns a security token.
- The entity passes the security token to the service it requested access to.
- The entity completes its required service access, the security token is revoked.

Figure 5 – AD Token-Based Authentication

AD uses a hierarchical structure for managing information in complex systems. At the lowest level is the concept of an organizational unit (OU) grouping shared role entities. These OUs then form an AD domain that contains the records for entities and objects in an administrative group, such as a business department or geographic location. AD domains can be structured in a hierarchical modular structure known as a tree to allow multiple administrative domains to be managed more easily separately. AD trees can then be grouped into an AD forest, where the boundary around all the constituent trees represents a security boundary. This approach relies on trust relationships between all trees within the AD forest. Large complex organizations can then implement multiple AD forests with a standard global catalog of all objects to manage infrastructure with internal security boundaries to manage information asset protection. Digital identity and permission configuration can then be shared across AD forests using an AD federation service to manage token provision to provide a federated IAM capability. AD federation services also enable SSO functionality to operate across security boundaries, including internet access points and between enterprise systems. AD processes issues and manage digital certificates.

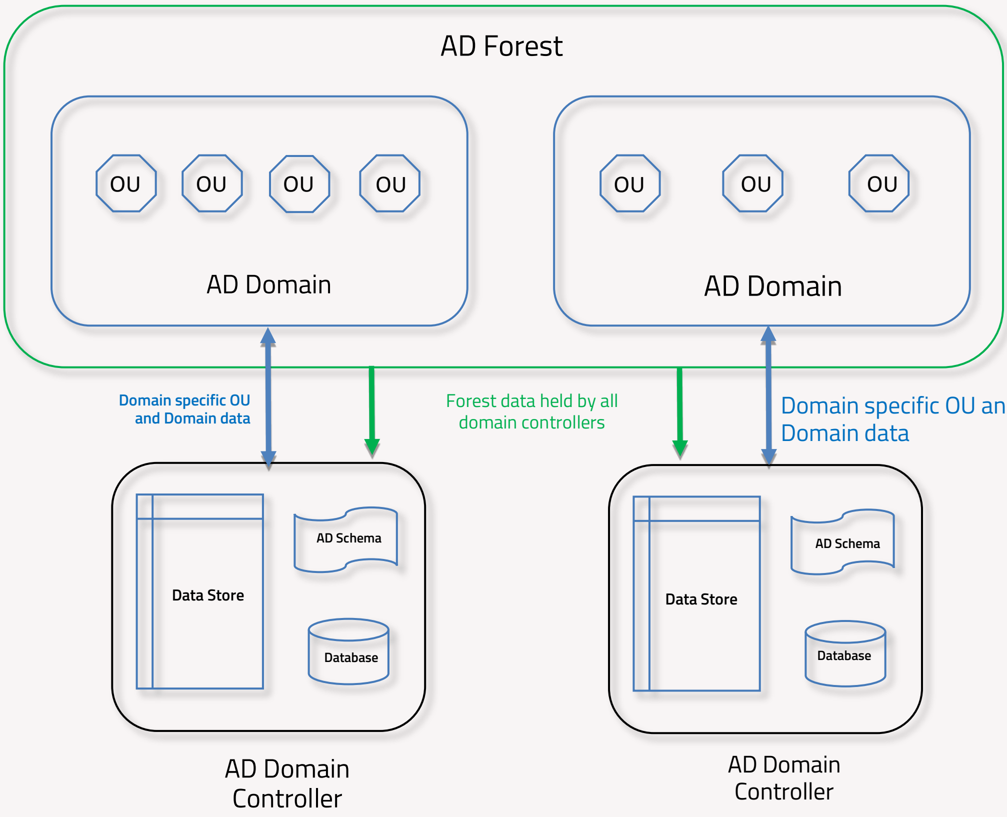

Figure 6 - AD Hierarchical Structure

The critical operating philosophy of AD solutions is the centralized issue and management of security tokens to allow entities to authenticate once for multiple authorizations for access to different resources.

The key benefit of the AD solution is the centralized permissions management to control information access across entities and roles within an organization from any connected device. The solution includes failover redundancy to manage equipment failures transparently for users and replication to support recovery processes.

The latest generation of cloud-based AD solutions moves to a managed Identity-as-a-Service (IaaS) approach that can span different cloud services and integrate with on-premises systems using pass-through authentication techniques

AD RISK REPORTING

AD solutions offer automated risk reporting and scoring to support business risk identification and management processes. This includes scoring the organization's identity security posture and comparison against industry benchmarks to help highlight weaknesses compared to industry best practices.

Identity-based risks cover user accounts and authentication sign-in events with an inherently high risk. Examples of high-risk user accounts are those with excessive privileges or weak authentication controls. Examples of high-risk sign-in events are those where an access request was successfully granted without using MFA, along with access by entities using legacy authentication processes assessed as being weaker than standard processes.

User account risk reporting is based on the probability that the account has been compromised using information such as abnormal user behavior, user actions that correlate with known attack tactics, techniques, and procedures (TTP), or intelligence, including discovering the account credentials within leaked breach data.

Sign-in risk reporting is based on the probability that the authorized identity owner did not initiate the authentication request. Typical drivers for such risks include requests made using an abnormal or compromised security token or originating from a suspicious endpoint device, IP address, geographic location, or other intelligence-led parameters. Other examples include multiple authentication requests from different places within a timeframe that precludes these being legitimate requests from a single user.

One of the most significant risks to AD solutions is that users employ weak or reused passwords that allow an attacker to compromise an account by using a previously compromised password, using brute force or intuitive deduction techniques to direct compromise passwords, or cracking the associated password hash to reverse engineer the password. Organizations enabling remote access without enforcing robust MFA techniques and without the enforcement of strong passwords to prioritize continuity of business operations over security significantly increase the likelihood of account compromise. Strong password policies and MFA can substantially reduce the risk of account compromise. At the same time, conditional access techniques that support behavioral analysis can improve the detection rate and minimize response times for remediating attacks.

Remote access vulnerabilities for endpoints that are misconfigured to enable remote desktop protocols represent a discoverable and exploitable ingress path for attackers. Weak credentials make brute force attacks on the remote access function a practical attack vector for gaining a foothold within a system from which further attacks can be launched. Remote desktop protocols should be turned off by default and only enabled in response to a specific need with robust security controls, including any MFA application.

Failure to proactively apply least privilege principles and role separation practices to AD account creation can result in the proliferation of domain administration accounts, including orphaned and unmanaged service accounts with excess access rights that can be compromised by an attacker with local access rights to gain full domain-wide administrator access to compromised systems by hijacking such an account. Techniques like rotating administrator passwords after use and frequent security auditing can reduce exploitation risks and support detection and response processes.

Including domain users in a local administration group will allow an attacker who has compromised a domain user account to gain privileged access to the relevant device and perform lateral movement across the system. Local administrative access to a compromised endpoint can allow attackers to elevate privileges and change network configuration settings to obtain full domain access and compromise security settings. Any organizational need for a domain user to have local administration privileges should be assessed using a risk-based approach, and if risk levels are deemed acceptable, then granted on a non-persistent temporary basis for the minimum time needed using least privilege principles.

A final weakness of AD solutions is that attackers can easily extract valuable information from the stored data records. This includes data that can allow the deduction of identity information and relationships and dependencies between entities and accounts that can be exploited to elevate privileges and move laterally across systems using complex attack paths that cannot be readily perceivable by security analysts without access to the same relational information

AD RISK-BASED ACCESS MANAGEMENT

AD solutions allow the application of risk-based access control policies to enhance protection in the event that a high-risk account is detected, or a sign-in event occurs. These conditional access policies allow the system to manage high-risk identity-related events with appropriate controls automatically.

Multiple risk-based policies can be defined to manage risks at a granular level to consider different user groups, roles, criticalities, and locations so access controls can be tuned to balance productivity against security across different risk levels. This approach also allows the system to handle changes of circumstances, such as a user temporarily changing work location without incorrectly denying access to that user but identifying the geographic variation as a change in risk level.

Examples of good risk-based policy practices include an explicit definition of known permitted access parameters such as device identifiers, geographic location, IP address, and sign-on times as low risk, along with credible possible parameters as medium risk. All other variations of these parameters would then be assigned as high risk. For example, a specific user group may only access systems from one company premise representing the low-risk access request. However, it may be credible for this user group to access systems infrequently from other specific company premises within the same country, representing a medium risk. Any attempt for this user group to sign on from any different location would be deemed high-risk and subject to conditional access controls.

Examples of additional controls before allowing access include enforcing MFA as part of the authentication process or requiring the user to change a compromised password. Other options include blocking access where the risk is deemed unacceptable or allowing access with no additional controls in cases where the risk is considered acceptable.

A key benefit of conditional risk-based access policies is the ability to automate risk mitigation using user auto-remediation processes for cases where enforcing MFA or changing a compromised password provides risk reduction to an acceptable level. Only those risks at the highest levels will either result in loss of access or require administrative intervention to resolve.

Risk-based policies can also be applied to non-user entities such as application or service workloads. The critical difference so these entities is they rely on sharing stored access credentials for authorization and cannot support MFA or perform auto-remediation steps. Detection of high-risk events can either block access to the workload or alert the security team to the need for further investigation.

Typical examples of high-risk workload access events include unusual sign-in request parameters, behavior correlating with known attack TTP, use of known compromised credentials, and detection of malicious actions.

While conditional access policies offer robust protection against identity-based threats, they can be misconfigured or maliciously exploited to block user access by falsely identifying all authentication requests as high risk. Good practice advice is that systems should include an emergency access account not covered by the conditional access policies to ensure access can be established under any circumstances. Such an account will require an alternate protection strategy to ensure it cannot be compromised under credible circumstances.

IMPLEMENTING AD IDENTITY PROTECTION

ESTABLISHING BASELINES

The first step in implementing AD protection is establishing a baseline of which entities require access to which resources following security policies governing aspects that affect AD configuration, including privilege allocation and segregation of duties. Entities typically include users and applications in this context, while resources include information repositories, data stores, and services.

The goal of the baseline is to ensure that all resources that require protection from credible threats have sufficient AD controls to reduce risks to below the maximum acceptable level. The effectiveness of this process is dependent on the rigor of information asset discovery processes to identify all resources requiring protection and risk assessment processes to establish the risks for each information asset.

THREAT DISCOVERY

An essential initial step for implementing AD solutions is ensuring that the system is free from dormant, unrecognized, and undetected threats that can compromise the integrity of the identity protection deployment process and remain hidden until activated later. All suspicious, abnormal, and unexpected behavior should be investigated and remediated until a sufficiently high level of certainty of a threat-free environment is established.

Another aspect of inherent threat discovery is the review of the results of the baseline AD access requirements to identify high-risk entities, such as user accounts or applications that require significantly high access privileges, and implement additional measures to reduce risk levels, such as splitting single high privilege accounts into two or more role-based lower privilege accounts or adding other technological and procedural controls to manage risks actively.

DEFINE ACCESS REQUIREMENTS

Authentication and authorization processes can be managed using a variety of technological and procedural controls, such as enforced password strength requirements or implementing MFA processes. The challenge is devising adequate controls strong enough to prevent compromise but user-friendly enough not to inhibit productivity or encourage users to bypass controls.

Users forced to create complex passwords that are impossible to remember to comply with complicated rules will look for a workaround that sidesteps controls, such as using easily guessed variations such as “Password1!” or “Qwertyuiop!” why tweaking a simple password to meet the requirements rather than following the spirit of the need for strong passwords and being secure. Enforcing good practice achieves better results through checks that prevent users from selecting passwords similar to any commonly used passwords rather than forcing the inclusion of a special symbol that typically results in most users simply adding a number and an exclamation mark to the end of their favorite password.

MFA protects against password compromise if implemented correctly. Effective MFA requires independence between the factors, so a single compromise cannot affect all factors. For example, if the two factors are a password and characters from a memorable word, both can be compromised using a single attack. Similarly, if one factor is a password and the other is a code sent to the device used to request authentication, then theft of that device can compromise both factors.

True independence combines unrelated factors such as something you know, something you have, something you are, and somewhere you are. For example, authentication requests sent from a known device in a known location with a password and biometric data, such as a fingerprint reader, provide robust MFA that is difficult to compromise.

CONDITIONAL ACCESS POLICIES

The next step is defining and implementing conditional access policies by establishing business-specific definitions of low, medium, and high risks and determining the remediation methods necessary for each risk level to reduce risks below the business risk appetite.

Self-remediation techniques such as requesting MFA for medium-risk authentication requests offer a good balance of threat prevention against user operability to increase security posture with minimal impact on business operations. However, an organization may feel that the risk of the self-remediation process being compromised by advanced persistent threats exploiting vulnerabilities in the process or the underlying technology may make this unsuitable for high-risk authentication requests. This may require manual intervention by system administrators or other authorized personnel to intervene in the authorization process to verify the validity of the authentication request before approval.

Good practice for businesses looking to establish new conditional access policies without impacting business operations is to implement report-only policies that do not impede authentication requests initially but alert system management resources to medium and high-risk access requests to allow post-access investigation. Once the conditional access policies' effectiveness and integrity are proven, remediation processes can be activated.

CONTINUOUS ACCESS EVALUATION POLICIES

The next step is defining and implementing continuous access evaluation policies to determine the criteria by which access to authorized entities should be revoked to protect systems in response to suspicious or malicious actions or when access should be extended in response to business performance needs and system resilience requirements. Access is managed through security tokens that, by default, have a fixed lifespan but can be expired or expended on demand in response to automated policy decisions

SECURITY INTEGRATION

The AD solution will generate logging and alerting information that should be integrated into the organization's managed security solution to offer the security analysts enhanced threat visibility of attack kill chains and allow automated threat detection and response.

The correlation of AD authentication request and authorization process data with logging and telemetry data from network and endpoint monitoring processes offers more comprehensive attack indicators for faster threat detection and alert prioritization and improved remediation results.

OPERATIONAL DEPLOYMENT

Once an AD solution is deployed, a process of continuous monitoring and improvement can follow to ensure the solution delivers the required results in terms of security effectiveness and system useability.

INTRODUCTION TO ACTIVE DIRECTORY

OVERVIEW OF ACTIVE DIRECTORY

Identity security of AD-based infrastructure encompasses managing and governance of identity-based access processes by protecting against threats and security systems and detecting and remediating attacks. In modern information systems, this protection covers on-premises and in-cloud resources with access requests from entities internal and external to the infrastructure boundaries in a sophisticated and hostile threat landscape.

ACTIVE DIRECTORY AUDITING

OVERVIEW

AD security audits are an assessment of the Active Directory Environment that evaluates the AD security levels using a defined methodology. Best practice audit methodologies should use risk-based evaluation of threats using a maturity framework to produce repeatable and comparable metrics that allow a view of security posture relative to industry peers and measure posture change over time.

The AD security audit aims to highlight risks that require management as part of organizational security practices and provide evidence of continuous improvement to support governance processes. Four core security indicators provide a risk-based assessment of AD security posture

STATE OBJECTS

Stale objects are inactive entities with defined permissions present within the AD structure that are available for activation by an attacker hijacking the entity or, for example, a former employee using their old user account for unauthorized purposes. Stale entities can result from the presence of obsolete technology or the deployment of upgrades creating new duplicate instances of existing entities. They can also exist due to a failure to remove redundant accounts or roles becoming redundant within the organization. Inactivity can be measured using logged entity activity data and remediated by de-provisioning unnecessary entities. Stale objects also include entities configured to allow access without a password which represents a significant risk or the presence of known vulnerabilities that have not been patched or otherwise mitigated to prevent exploitation.

PRIVILEGED ACCOUNTS

Privileged accounts are entities with excessive assigned rights and permissions or inadequate controls to minimize use that are configured. For example, privileged accounts should not allow delegation to prevent impersonation by attackers. It also detects non-privileged accounts that can perform privilege operations due to misconfigurations, such as standard users able to modify AD Group Policy Objects (GPO) due to inherited rights. Other checks include ensuring privileged accounts cannot be used on insecure endpoints with a high risk that credentials can be stolen or using obfuscation rather than encryption to protect stored passwords.

TRUST RELATIONSHIPS

Trust relationship misconfiguration can create significant security vulnerabilities in AD solutions. Forest and domain trust relationships can exist as one-way or two-way relationships, and their requirements can change over time, leading to excessive trust relationships accruing and reducing the overall security posture. Redundant or superseded trust relationships should be removed to minimize risks associated with the lateral movement of attackers across AD domains to elevate privileges and expand the attack surface.

ANOMALIES

Anomalies are entities or AD structures with suspicious or anomalous configurations that offer opportunities for attackers to compromise accounts and hide their activities. Critical issues include weak password policies, including an inadequate password policy for local administrator accounts, and the use of weak, vulnerable, or depreciated security protocols or encryption technology to transfer authentication requests. Other risk indicators include evidence of standard accounts granted temporary privileged rights, erroneous or missing AD backups, and an AD audit policy inadequate to support threat detection and response processes.

CONTROL PATH ANALYSIS

An important element of an AD security audit is a control paths analysis that assesses exploitable permissions issues that can allow an attacker to take control of a domain. Critical issues include accounts within the control path of a domain located in different foreign domains or excessive numbers of entities with indirect access to a domain that an attack on the domain can exploit.

AUDIT REPORTING AND FREQUENCY

The output of the AD security audit should be a prioritized list of risks requiring remediation with advice and guidance for risk management options ranging from implementing new technological controls, introducing operating procedures, or using a cost-benefit analysis to justify risk acceptance.

The goal of remediation guidance and recommended best practices is to improve the AD infrastructure performance and security posture.

An AD security audit should be undertaken as the first step in implementing an identity security solution. Then, as a minimum, repeated periodically at an interval that reflects the criticality of the system being protected and the intensity of the threat landscape. Additional AS security audits should also be performed following any significant events, including functional changes to the system under protection, substantial changes to organizational business practices and security policies, significant changes to the threat landscape, or in response to a major security incident where identity compromise was a factor in the attack chain.

IDENTITY THREAT DETECTION & RESPONSE

Identity threat detection and response (ITDR) services add another layer to security defenses to protect against attacks leveraging the compromise of digital identities and authentication processes. Attack vectors can range from exploitable vulnerabilities in IAM systems allowing the discovery and theft of authorization data to social engineering techniques to trick users into revealing access credentials.

ITDR encompasses tools and processes to protect identity systems that combine threat intelligence with best practices to implement detection mechanisms—key to threat detection of monitoring and investigating changes to configurations, authorizations, and behaviors.

ITDR good practices include ensuring the IAM infrastructure implements a single authoritative access directory that is controlled using configuration management and change control processes. The IAM must be fully supported and maintained, including robust security update management processes, complying with the latest standards and industry best practices.

The IAM should offer single sign-on access management with continuous assessment of user context attributes, including user and entity behavior analytics (UEBA) functions, account takeover (ATO) fraud detection functions, and support identity governance and administration processes.

Effective ITDR requires integration into the organization's managed security solution, typically achieved by collecting identity and access log data and feeding it into an extended threat detection and response service. Identity, access, and AD detection rules can then be configured to trigger security alerts monitored and analyzed as part of the standard security incident management processes. A list of the recommended rules for monitoring the security status of AD solutions and supporting threat detection processes is provided as an appendix to this white paper to help select an identity security solution provider.

Identity security assessment and auditing can provide visibility of AD misconfigurations, suspicious configuration or credential changes, and unauthorized access events to reduce the risk of identity-related attacks.

ACTIVE DIRECTORY DECOYS

DISRUPTING THE IDENTITY ATTACK LIFECYCLE

An effective technique for detecting the more advanced threats a business faces is using proactive protection measures using decoy credentials and services to act as tripwires for alerting the security team. Decoys provide a highly effective method of detecting attacks at the earliest stages of an attack before any significant damage or compromise.

A key element of typical attack paths is the search for high-value access credentials for systems and services or privileged accounts. Creating false user identities and permissions offers the attacker a tempting target, and these decoy credentials recorded in the AD directory will never, under normal circumstances, be accessed by a legitimate entity. Therefore, any access to these fake entry points to either an unintentional act or reconnaissance by an attacker. A validated alert is automatically triggered when the attacker has these false credentials. Generating a validated security alert for access to a fake certificate or service provides security analysts with early warning of an active identity attack path. Creating personas linked to false user credentials enhances authenticity and believability and encourages attacker interaction to maximize engagement.

This proactive approach using deception techniques enables the detection of an attacker's presence within a system as they undertake passive surveillance and reconnaissance or difficult-to-detect active privilege escalation and lateral movement actions.

Deception techniques also provide security analysts with valuable intelligence on attacker TTP, behavior, and intentions to support remediation and recovery actions and inform the broader security community. The security team can also employ the information gained to hamper any ongoing attack using misdirection and misinformation.

INITIAL COMPROMISE

During the initial attack phase, an attacker uses compromised identity information to establish a foothold to gain permanent remote access to a system to extend their attack. Decoy technology provides alerts for any attempts by an attacker to compromise systems or communicate with command and control servers. Misinformation allows the defenders to move and contain the attacker within a decoy environment.

INTERNAL RECONNAISSANCE

Once an attacker has established a foothold, internal system reconnaissance allows the attacker to discover critical systems and sensitive information needed to further the attack, such as system configuration settings, credential repositories, AD information, and other intelligence beneficial for the attack. Decoy applications, services, and data can be employed to trigger alerts to detect their presence and track activities.

PRIVILEGE ESCALATION

Identity-based attacks typically require a sequence of privilege escalation actions unless they were lucky enough to compromise a high-value administrator account to gain a foothold. Privilege escalation exploits the information uncovered during the internal reconnaissance phase. Decoy credentials can be employed to trigger alerts and allow defenders to move the attacker to a decoy environment.

LATERAL MOVEMENT

Identity-based attacks typically involve movement across domains and systems to gain access to valuable information assets and extend the attack to achieve the most expansive reach within systems. Transversing across boundaries also allows attackers to move across organizational boundaries in supply chain attacks. Decoy networks can be employed to trigger alerts and allow defenders to contain the attacker within a decoy environment.

RECONNAISSANCE

IDENTITY AND CREDENTIAL BREACH MONITORING

Identity breach monitoring is the process of discovering and tracking stolen identity-related information, including account credentials, on the Internet. Data stolen from individuals and organizations, including usernames and passwords, are collated and traded on deep and dark Web marketplaces beyond normal Internet users' reach. There is a growing business space where sophisticated attackers steal identity-related information and then sell this to less capable criminal actors for fraudulent activities. Other sensitive information traded includes network diagrams, firewall configurations, and AD visualizations obtained during reconnaissance that other attackers can use to undertake their own differently motivated and actioned attacks.

Identity breach monitoring will shine a light on the deep and dark Web using intelligence-led processes. The critical benefit of identity breach monitoring is that it provides evidence that an organization's systems have been breached in situations where a sophisticated attacker has compromised the system undetected using compromised credentials and left no visible indicators of compromise on the system. Detailed analysis of the information uncovered on the deep and dark Web can indicate when the data exfiltration is likely to have occurred and the extent of the breach, providing valuable insight for security analysts investigating the attack to halt and remediate the attacker's kill chain. Information gleaned from the stolen data and its presence in a criminal marketplace may also allow analysts to identify who performed the attack and their TTP to provide intelligence helpful in preventing further breaches.

Identity breach monitoring services should cover the entire World Wide Web, the surface Web indexed by regular search engines, and the hidden deep and dark Web with its content that cannot be found using traditional search engines and is not accessible using standard browsers. Specialist security service providers have special tools and applications to discover and access deep and dark Web content. This includes private websites, hidden criminal chat rooms, peer-to-peer networks, Internet relay chat (IRC) channels, social media platforms, and black markets. Scanning also requires advanced search techniques, including automatic OCR for images and other non-textual formats that can hide information from text-based searches.

There are four levels to breach monitoring:

- Level 1 covers scanning surface Web and social media sites, with breach monitoring performed using policy-driven crawling across the Internet in the search for unintentional data leakages.
- Level 2 covers scanning the more accessible areas of the deep Web for data spills, crawling the entire open FTP Internet Spectrum, Databases, Sharing, and Storage sites, and employing crawling and sniffing techniques to map the whole Bit Torrent network.
- Level 3 covers scanning the semi-restricted areas of the deep Web, including crawling binary sites, leakage forums monitoring, checks for private and identifying information and contracts monitoring and employing automatic entity extraction and pattern recognition algorithms.
- Level 4 covers scanning the dark Web, black marketplaces, and underground networks, including Onion router (Tor) network searches, Onion sites, escrow and marketplaces, forums, and the Invisible Internet Project (I2P) fully encrypted private network layer.

Effective identity and credential breach monitoring should employ wide-reaching sources and types using automation, machine learning, and crawlers for continuous data accumulation combined with ingesting and indexing non-crawlable data. Analyst-driven augmentation techniques should source from the latest underground communities, new channels, authenticated forums, and chat platforms. Services also will require continuous product adaptation to new technical formats, such as the changes to the Tor network.

IDENTITY AND CREDENTIAL BREACH MANAGEMENT

Responding to identity and credential data quickly and effectively is crucial for minimizing the breach's impact. This impact can be tangible regarding disruption of business operations, financial loss, and remediation costs. However, there can also be intangible losses in terms of damage to business reputation, the impact of theft of personal information of employees or customers, and the disclosure of sensitive organizational information into the public domain.

Business disaster recovery and business continuity processes should include a well-planned and comprehensive step-by-step strategy to ensure the effectiveness of post-breach remediation in the event of any identity and credential-related breach. Where the organization operates in a regulated business or manages personal identifiable information, additional steps may be needed to notify the applicable regulatory bodies within the required timeframes, demonstrate corrective actions, and offer a transparent investigation of the root cause.

The first step in managing an identity-related data breach is to regain control of the affected system. This step can be more challenging for identity-related attacks due to the potential for communications with affected systems to be compromised, either preventing legitimate access or allowing a sophisticated attacker to monitor and disrupt recovery actions. Establishing secure communications is critical for an effective incident investigation and response.

The next step is to fully understand the type of breach, its scope, its severity, and its business impact to ensure the correct response and remediation is undertaken to resolve the incident quickly and effectively.

The next step is to establish if any credentials or identity information has been compromised and take measures to prevent the exploitation of compromised credentials by taking positive actions such as locking accounts, forcing the change of passwords, or limiting access until the integrity of each identity is established. Once complete system control has been regained, measures can be taken to remediate or block possible persistence techniques that the attacker may have employed and detect and remediate any new access exploits that may have been deployed.

The next step is to inform all affected parties on the nature of the attack, its impact on digital identities, the actions the organization is taking, and any actions the affected users should take. Where personal identifiable information is compromised, measures must be taken to minimize the risk to involved parties. Where credentials to third-party systems are compromised, the affected users and the system owners must be kept fully informed. It is crucial to have processes to manage communications, including formal statements and public announcement procedures, to ensure clear, consistent communications that minimize reputational impact.

The next step is to implement the relevant post-breach incident response plans as part of the security response process to contain and halt ongoing attacks, reverse the impact of the attacker’s actions, and resolve the underlying weaknesses and vulnerabilities that the attacker was able to exploit in the conduct of the attack.

The final step in the immediate response is a security review process with lessons learned output to ensure that the nature of the attack and its impact on the business is fully understood and that such an attack cannot reoccur within practical business constraints of costs and resource requirements.

Then the business should undertake periodic data exposure monitoring to establish any new, additional, or distribution of stolen credentials, assessing the impact of any change and taking remedial actions as necessary in the case of recent disclosures.

IDENTITY PROTECTION THREATS

TYPICAL IDENTITY-RELATED RISK

The following represent the most common identity-related risks currently faced by organizations.

EXCESSIVE PERMISSIONS

This risk comes from the accruement of excessive permissions in a dynamically changing organization where permissions are added to users in response to changes of roles or additional duties but where there is no clear process to remove permissions that are no longer needed. Managing this risk requires robust methods to control permissions management supported by regular audits to identify discrepancies against the permissions policy and processes to resolve identified issues.

STALE ACCOUNTS

This risk comes from the failure to correctly offboard users or decommission devices such that the access credentials with associated permissions remain assigned within a system and available for exploitation at a later day. This risk can materialize from an attacker discovering and misusing the stale account or from a variant of an insider attack where a disgruntled ex-employee uses their old access credentials to log on to the system. Managing this risk requires robust processes to manage access credentials supported with regular audits to identify discrepancies with techniques to resolve identified issues.

MISCONFIGURATION

Access control configuration can be complex, and errors are often not readily apparent or have an identifiable impact on system operations. This issue is particularly problematic for access configurations performed by deploying applications or services where administrator visibility of changes may be limited. A simple misconfiguration can make private data publicly visible or shared more widely than necessary. Configuration risks can be managed by employing an identity and access management solution that can automatically detect accidental and malicious misconfigurations supported with processes to resolve identified issues.

UNMANAGED VULNERABILITIES

Weaknesses and vulnerabilities in applications, services, and systems are regularly uncovered, including in technology and processes that handle access management and authentication and authorization actions. Once a vulnerability becomes known, the provider of the affected product will swiftly move to resolve the issue, typically with the release of a security patch or application upgrade. In some circumstances, inherently vulnerable technology may be replaced with a more secure alternative. However, there is a window of opportunity between discovering a vulnerability and affected user organizations implementing mitigation measures that an attacker can exploit. The slower an organization is to install patches or replace technology, the more time and more significant the risk an attacker may leverage the vulnerability to attack systems. Managing vulnerability risks can be achieved by adopting robust patching policies and actively mitigating known vulnerabilities as soon as they are discovered.

Typical identity-related threats to organizations revolve around attackers gaining a foothold using compromised credentials and extending their reach within the affected systems using privilege escalation and lateral movement techniques.

Attackers typically use administrative permissions through an on-premises compromise to gain access to the organization’s trusted Security Assertion Markup Language (SAML) token-signing certificate. SAML is the open standard for exchanging authorization and authentication data between identity and service providers. Access to the certificate allows an attacker to forge SAML tokens that impersonate an existing privileged user account. These tokens can be used against any on-premises or cloud-based resources configured to trust the organization’s security tokens. This technique allows attackers to add their access credentials to any existing application service principal to grant themselves access to the APIs.

COMMON CREDENTIAL THEFT TECHNIQUES

Credential theft is one of the simplest methods for attackers to bypass perimeter security controls and gain undetected system access. While vulnerabilities such as weak credentials allow remote attackers to brute force access, more subtle techniques are available for more sophisticated attackers looking to compromise systems with robust security controls.

SOCIAL ENGINEERING ATTACKS

Social engineering requires very little technology and knowledge to implement, relying on the attacker's ability to convince a target user to perform an action or divulge sensitive information so they can gain access to credentials. It may be as simple as phoning a user pretending to be a system administrator and asking them to reveal their password to resolve some invented problem.

Phishing emails and smishing text messages represent a low-skill social engineering attack vector that requires the attack to have very little knowledge about the target users. Instead, it relies on contacting many users and hoping that just one will fall victim and either respond to the message or click on a link.

Spear-phishing and whale-phishing attacks refine phishing techniques where the attacker uses knowledge about a specific individual to send them a targeted message. This method has a greater success rate but takes much longer to produce, requiring the attacker to have some level of skill to perform correctly.

Common social engineering techniques against large organizations use links to shared resources that mimic applications such as OneDrive or SharePoint to steal Microsoft 365 access credentials from a victim to use in an attack. For a user with poor security awareness, logging into what appears to be a Microsoft resource seems less risky than providing a password to a more sensitive business application.

MALWARE ATTACKS

Malware attacks rely on malicious software downloaded onto a user's device to steal credentials entered into the device. This can use techniques like keylogger software to record key presses to capture helpful information such as usernames and passwords entered into the affected device. Other malware may imitate specific applications and services to capture the credentials used to access these in a more targeted attack.

Malware can also employ other techniques, such as screen captures, to defeat authentication processes that do not rely on users typing credentials, such as using pull-down menus to enter characters from a memorable word.

An example of a successful malware attack was the use of a phishing technique that resulted in the download of trojan malware that included a keylogger function by employees at a healthcare organization. Five employees fell victim to the phishing attack, and the result was access credentials were obtained that allowed the attacker to steal nearly eighty million sensitive medical records.

AUTOMATED ATTACKS

Repetitive attack methods such as brute force guessing passwords, dictionary attacks, credential stuffing, password spraying, and other trial and error techniques can be automated to interrogate system authorization processes to discover valid credentials. This allows attackers with minimal skills to uncover credentials as long as they can access a computer and find a target system vulnerable to such attack methods.

While requiring little skill, these attack types have yielded significant results. These include a brute force attack on the Northern Irish parliament that compromised email services, a credential stuffing attack on the Canadian Revenue Agency (CRA) that compromised 48,500 accounts, and a password spraying attack on Citrix that allowed access to information held in a shared network resource.

VULNERABILITY EXPLOITATION

Applications and endpoints with known vulnerabilities are common targets for attackers looking to gain a foothold in a system from where they can steal credentials within the system boundaries and extend their attack.

Competent attackers will have access to tools that can scan Internet-facing applications and services to uncover unpatched vulnerabilities with known exploitation techniques. This approach is a common technique to start a kill chain due to the large number of organizations that do not have robust patching policies that ensure security updates are actioned within reasonable timeframes.

The attack on Facebook in 2018 is an example where Internet-facing systems were left with multiple unpatched vulnerabilities for over a year. The attack that leveraged these vulnerabilities impacted over fifty million user accounts, compromising personal information and resulting in a $18 million fine for violating European data protection regulations.

COMMON CREDENTIAL THEFT TECHNIQUES

The SolarWinds attack is an excellent example of a sophisticated identity-related threat used to attack many high-value target organizations with capable security defenses and good security postures.

INTRODUCTION TO SOLARWINDS

SolarWinds Inc. is an American IT services company that produces management software that helps large organizations manage their IT infrastructure and networks. One of their products is the Orion Network Management System (NMS), a centralized performance monitoring and network administration solution. It is designed to oversee and maintain networks, including managing configuration settings and rolling out software updates and critical security patches. The product requires privileged access permissions to access and modify the systems on which it is installed to perform these actions.

In December 2020, it was reported that the SolarWinds Orion product had been infected with malicious code. It is believed that the infection took place in March 2020. However, the initial breach of the SolarWinds systems that led to this breach would have been significantly earlier than this date due to the attack's complex nature. This resulted in what is believed to be one of the most damaging supply chain cyber-attacks. Several labels describe the SolarWinds attack, including SUNBURST, SUPERNOVA, SUNSPOT, TEARDROP, and RAINDROP.

The state-sponsored actors that successfully attacked the SolarWinds systems and inserted malicious code into the Orion product used the SUNBURST process. This process created a vulnerability the attackers could exploit once the infected Orion software was installed on customer systems. The SUNBURST vulnerability was then used to download malware code labeled SUPERNOVA onto infected systems. The process for injecting the SUNBURST backdoor during the Orion Platform build process is marked SUNSPOT. The SUNBURST vulnerability uses TEARDROP and RAINDROP malware loaders as part of its operation.

MALWARE DEPLOYMENT

The attack started when sophisticated state-backed actors gained access to the SolarWinds systems. They achieved this by allowing their presence to remain undetected while enabling them to access the updated servers. This enabled them to add malicious code into the software that the Orion product downloaded into the client's systems as part of the update process. The clever part was that this malicious code remained inactive while the updated software was verified on development systems. It only became active once it was deployed into a production environment.

Once operational, the malicious code would then communicate back to the attacker's systems to download additional resources to give the attackers control over the infected systems. Communications were routed through IP addresses located in the victim's own country to reduce the chance of detection. This was implemented as a multi-layer attack using various attack vectors to maximize the probability of success when attacking robustly protected networks. The key to success was exploiting successful infections without detection, spreading the infection, and extracting as much information as possible.

The malicious code automatically gained access to a process with some of the highest privileges and the most significant reach in the affected networks using the software patching mechanism as the attack vector.

The malicious code was inserted in the SolarWinds.Orion.Core.BusinessLayer.dll plugin component of the SolarWinds Orion software. The code contained a backdoor function that used the Hypertext Transfer Protocol (HTTP) to communicate with attacker-controlled servers. This code was sufficiently obfuscated to avoid detection. After an initial dormancy period and verifying that the code was on a live network by checking for connectivity with the servers, the code became active. It accessed the remote servers to download commands to fulfill various functions. These included identifying anti-virus and forensic tools, disabling system services, profiling the infected system, transferring data, and executing files. These functions were designed to mimic legitimate operations performed by the SolarWinds Orion software to avoid detection. Of particular use was the Orion Improvement Program (OIP) protocol that collects evaluation, performance, and usage data from users. It does this to enable SolarWinds to monitor software performance and aid fault diagnostic processes. The SUNBURST malware collected data and reported it to its servers by emulating this protocol without raising suspicions.

Another technique employed was to hijack a legitimate process and replace the functionality of the standard operating system process with a malware function, executing the malware and then replacing the malware code with the original legitimate code. This prevented any forensic analysis from identifying that the malware code had been run, with log files only recording the legitimate code's execution.

IDENTITY

A critical element of the initial phase of the SolarWinds attack was that the attackers were able to bypass an outdated MFA by stealing a Web cookie. The use of known weak MFA technology gave a false sense of security to the access control processes that the attackers could identify and exploit. This vulnerability was believed to be found as the initial stages of the attack focused on the target organization's identity infrastructure.

The subsequent attack kill chain also included stealing passwords using Kerberoasting to elevate privileges, stealing SAML certificates to enable identity authentication by cloud services, and creating new accounts on the AD server. These attacks on the AD environment enabled the attackers to move laterally from the on-premises environment into a Microsoft Azure cloud environment to extend the reach of the attack.

IMPACT

The breach was uncovered when the US cybersecurity company FireEye identified that its systems had been breached and traced the attack vector back to the SolarWinds Orion product. They dubbed the malware SUNBURST and produced a set of signatures to enable organizations to scan their systems to determine if they had been infected. This set off a chain of events where tens of thousands of Orion product users were alerted to the breach.

It was estimated that around 18,000 out of the 33,000 SolarWinds Orion users had installed the infected software, representing approximately 6% of the total SolarWinds customer base.

The majority of infected users were non-governmental organizations. However, this customer base included US organizations handling some of the most sensitive national security information.

Overall, it is estimated that more than 250 US federal agencies were affected by the security breach. It also affected worldwide governmental organizations and critical corporations, including Microsoft, Cisco, Intel, Visa, and AT&T.

REMEDIATION

The SolarWinds attack has highlighted the vulnerability of supply chains when a link in that chain is compromised. The reliance and levels of trust that the SolarWinds customer base placed on the products that SolarWinds provided allowed the attackers to bypass security controls. It is not uncommon for procurement processes to focus on suppliers' financial history and stability rather than cybersecurity measures. The attack provides a valuable wake-up call that corporate risk management processes should include supply chain security. Often supplier risks focus on the failure to deliver on time and quality. They also need to look at the impact of security incidents that originate in suppliers.

One solution is to adopt a zero-trust philosophy. That is, applications and devices inside and outside an organization's IT systems must prove their legitimacy before being allowed to perform any function. This adds a layer of bureaucracy to processes but would prevent a software update containing malicious code from being allowed to be executed before verifying its legitimacy. There are limits to how practical a zero-trust philosophy would be in the face of highly sophisticated malware that can hide until after the completion of verification processes.

IDENTITY THREAT LEVELS

Statistics from the 2023 Trends in Securing Digital Identities report produced by the Identity Defined Security Alliance (IDSA) show that 90% of respondent organizations across the USA experienced at least one identity-related breach in the past year. Of these, 68% reported that they suffered a direct business impact due to the breach.

The report found that the number and type of digital identities businesses manage are increasing, driven by the growth in remote and hybrid work practices, expanding digital relationships with contractors and third parties, and the explosion of machine identities. This significantly increased the attack surface for identity-related attacks and increased the probability of attacks leading to breaches.

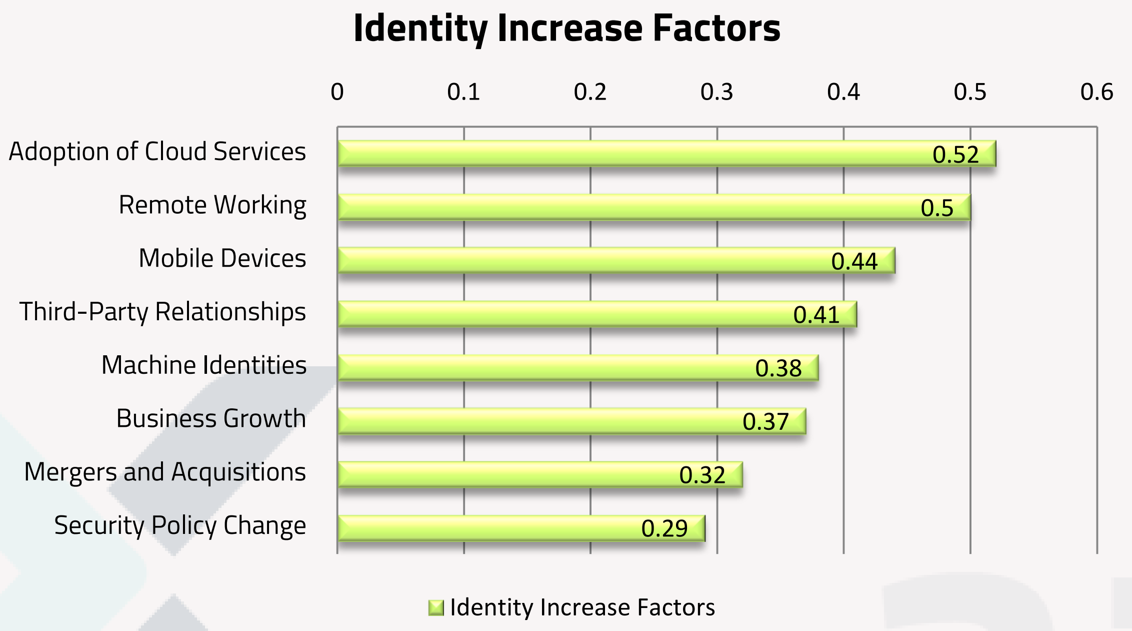

Figure 7 - Factors Driving the Increase in Digital Identities

Regarding identity and credential-related attacks, the greatest released threat was from phishing attacks ranging from broad-based organizational-wide attacks to more sophisticated and targeted spear-phishing and whale-phishing campaigns.

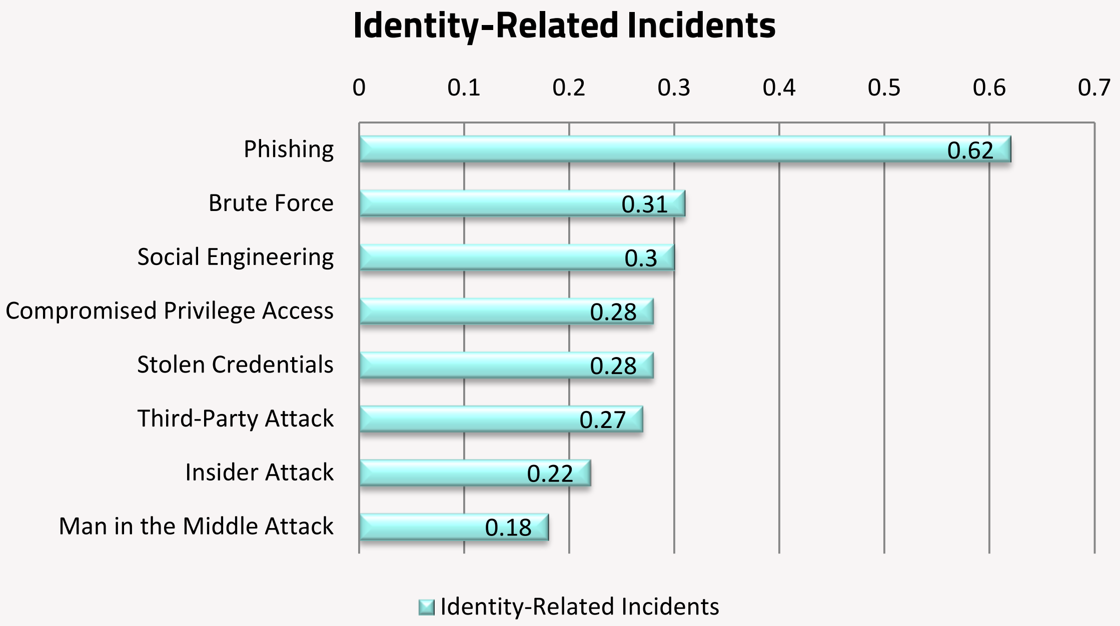

Figure 8 - Recorded Identity-Related Incidents

System users' poor security awareness and practices were recorded as identity-related incidents' most significant underlying causes. This includes succumbing to a phishing attack by clicking on an email link and reusing passwords across multiple accounts, including work and personal accounts. Using unauthorized and weakly secured personal devices to connect to business systems was also a factor in these incidents.

These factors are undoubtedly behind the importance that businesses place on managing and securing digital identities, with 96% of respondents rating it in their top ten security priorities.

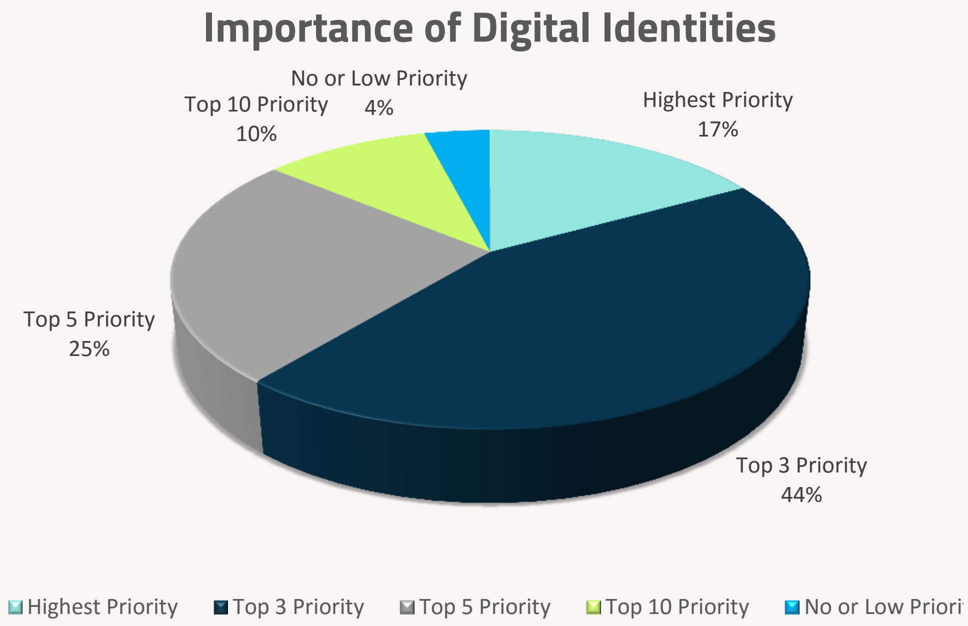

Figure 9 - The Perceived Importance of Digital Identity Security

Regarding the impact of identity-related incidents, 68% of respondents reported that the attack had a direct business impact. Of these, 39% incurred a financial hit due to the cost of recovering from the breach, and 33% reported a productivity impact due to the incident distracting from core business. Significantly 17% of respondents reported an adverse business impact due to lawsuits and other legal action, and 25% reported a perceivable effect on the business reputation due to the breach.

Managing and securing digital identities is also seen as a significant challenge for businesses due to technological and commercial barriers and resource availability challenges having the most significant impact. One of the fundamental underlying causes of barriers is a lack of proactive investment from the management team to invest and support identity security. This is primarily seen to be due to the leadership team’s lack of understanding of identity and security risks. Interestingly 29% of respondents reported that their leadership team only engaged with identity security support following an incident. Unless these barriers are overcome, businesses will expect to see an increase in identity-related breaches in the near future as attacker sophistication increases.

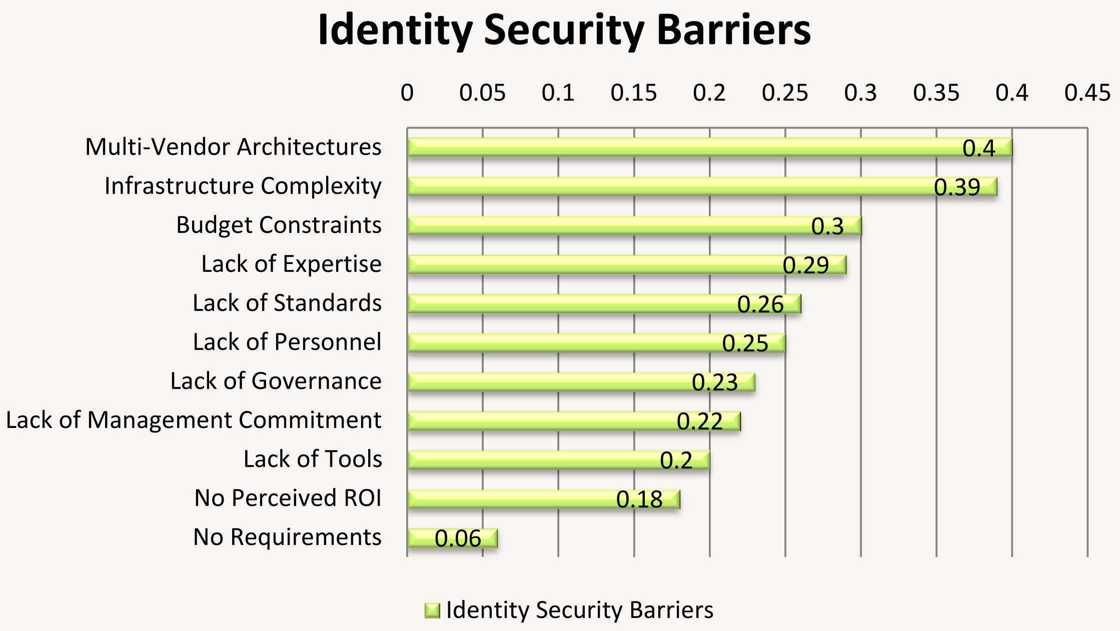

Figure 10 - Barriers to Effective Identity Security

An interesting outcome of the survey was the respondents reporting on the measures they believe would have prevented or minimized the impact of the security-related breaches suffered by their business. The results provide a helpful cross-reference against industry best practices, demonstrating their practical value as preventative measures.

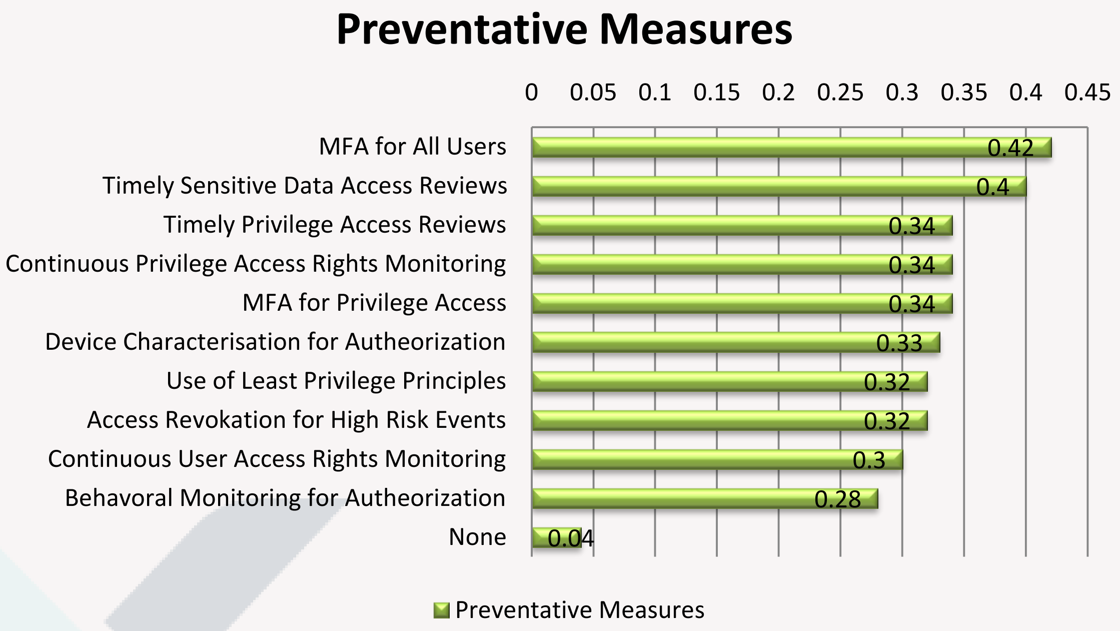

Figure 11 - Factors that Could Minimize Incident-Related Security Incident Impact

IDENTITY PROTECTION SOLUTION

EFFECTIVE IDENTITY PROTECTION

Identity protection secures your digital identity and authentication information and alerts you of threats. Employing an identity protection solution will help an organization manage identity-based threats with a service that will detect, investigate, and remediate identity-based risks quickly, efficiently, and effectively. This includes auto-remediation techniques, risk-based conditional access mechanisms, and integration with existing security solutions, such as an advanced extended detection and response (XDR) solution or a more fundamental security information and event management (SIEM) tool for investigation and correlation.

Managing and maintaining an AD environment can become complex and challenging as organizations grow and evolve. The constantly changing security enhancements and configuration options result in enterprises failing to retain their AD environments correctly.

Identity protection solutions should assist organizations in enhancing their processes and configurations and security and monitoring controls necessary to secure an AD environment effectively. This should include the ability to deliver guidance and recommendations as part of assessments that reflect real-world tested techniques and expertise with a track record of successfully eradicating attackers from client networks and helping to contain and remediate threats.

IDENTITY PROTECTION APPROACH

Identity protection solutions should be capable of conducting security assessments remotely with support from key client stakeholders. Collaborative knowledge-sharing processes, including online workshops, will allow the review of existing on-premises and cloud systems and network architectures. Using software scripts to collect information from AD servers can also allow the identity protection solution's security analysts to identify misconfigurations and potential attack paths within the AD infrastructure.

IDENTITY PROTECTION ASSESSMENT FOCUS

The identity protection security assessment provided by the identity protection solution should include consideration of the following critical elements:

- Domain overall risk level
- Stale objects
- Privileged accounts
- Trusts
- Anomalies
- Forest architecture
- Operational processes
- Monitoring and Response
- Group policy controls and enforcement
- Permission delegation
- Service accounts and service principal names (SPNs)
- Remote access controls and hardening
- Endpoint configuration and hardening
- Integration with Microsoft Azure and Microsoft Office 365

The goal of the security assessment is to gain insight into the AD environment and identify exploitable AD configuration issues and identify the reduction of AD-privileged access through the development and enforcement of stricter policies. The benefit to businesses is the ability to better detect and respond to AD environment threats using a combination of best-practice defensive and offensive approaches.

IDENTITY PROTECTION PRODUCTS

Identity protection solutions should deliver a detailed report of the identity threat assessment and recommendations, including an overview of the current AD security issues and misconfigurations.

Clients should also receive advice and guidance for improving identity security posture that encompasses the best practices for securing the AD environment and managing privileged user access and accounts in line with operational processes. This should also include detailed recommendations for locking down the AD environment and resolving all issues and misconfigurations identified by security assessments.

ZERO TRUST SOLUTIONS

OVERVIEW

One security management technique for minimizing identity-related risk levels that is gaining prominence is the adoption of a Zero Trust philosophy. This adoption is partly driven by the US Government's decision to move towards greater use of Zero Trust Architectures to help combat the increasing cybersecurity threats from progressively sophisticated and persistent threat actors. This move includes the issue of a Federal Zero Trust Strategy by the US Office of Management and Budget (OMB) to initiate the process of protecting US Government systems implementation of rigorous access and monitoring controls for all users, devices, and systems, irrespective of their location, in line with trust no one attitude.

ZERO TRUST PRINCIPLES

The Zero Trust philosophy assumes that all users and devices are untrusted, significantly altering how security policies assign and enforce roles and responsibilities. The principle of least privilege ensures authorized users are only allowed access to those services and specific data necessary for performing their duties. This minimizes security incidents due to unauthorized or accidental access to services or data but must include logical and physical access restrictions to be effective.

Division of duties policies prevent a single user from performing a complete end-to-end process but instead impose rules that permit a user to execute part of the process and then require an independent user to perform the following function. This approach will prevent a single compromised account from completing potentially damaging operations. Demonstration of true independence is necessary; for example, separate users should not access a service from the same device.

Dual operator policies require actions by two separate users before an activity is permitted to provide independent verification that the action is authorized. Financial transactions or granting privileged access will typically use this approach. A demonstration of true independence is again necessary for this to be effective.

ZERO TRUST ARCHITECTURES

Zero Trust Architectures are systems that assume you can trust no one and deny access to every user and device until proven trustworthy with granted access to services and data limited to the minimum necessary resources.

This approach is a significant shift in policy for managing security controls. It goes beyond a simple change to network design to fundamentally rethinking security philosophy. With the deployment of sophisticated attack vectors and uncovering long-term exploitations, this step change may give network security a lead over the malicious attackers.

Traditionally, architecture-imposed security controls between an internal system and the outside world will focus on blocking unauthorized access. However, security moved from active management to passive monitoring once authenticated and inside the boundary. Any malicious process inside the system could exploit further weaknesses to escalate privileges and hide its presence.

In a Zero Trust Architecture, there is no assumption of authenticated trust inside the boundary. Trust must be earned; connections are assumed suspect until proven legitimate. This includes all networks, even hardened wired local networks; there are no exceptions. Every interaction must employ robust authentication techniques.

Zero Trust Architecture relies on network designers having a fully defined architecture regarding systems, services, devices, and users. Any omission or incomplete definition can lead to weaknesses. This approach requires a robust definition of all devices and users for efficient controls and minimal vulnerabilities.

Establishing trust comes from building confidence in the communications between a device or user and a service. Monitoring and inspecting transactions allow the system to build up a picture of the trustworthiness of the network link, using the results to determine whether to grant access to the service. Thus, trust is established within the architecture on a case-by-case basis for every interaction instead of trust decisions being concentrated at the boundaries of a firewall or VPN connection.

Continuous monitoring of services and devices is necessary, with action taken on detecting issues, whether operational health problems or detection of suspect activities. Security should be proportionate to the importance of services and data. Access to critical services and essential data requires strict control using the principle of least privilege to minimize access to information and resources. It is crucial to note that implementing a zero-trust model will require selecting services designed to operate in a zero-trust environment.

ZERO TRUST POLICIES

The fundamental principle of using a zero-trust model is that every action a user or device performs is subject to a policy decision that determines if it will be permitted. This invisible operation verifies each access attempt to data or resources. Access is forbidden if the policy criteria are not met, and a security action is initiated. This will severely restrict an attacker's activities that have compromised a user account or network device.

Zero trust policies provide the mechanism for determining which users and devices are granted authorized access to which services and data. Each policy comprises a set of rules applied to access requests from users or devices that meet defined criteria. The compliance assessment with the rules results in an action assignment to that access request.

The criteria identify the nature of the access request to determine which rules will apply. The requirements can be a single individual named user or any access requests originating from a region of the world. The following are examples are typical criteria:

- Access control identifier
- Specific email address
- Email address domain
- Specific IP address
- Source IP address range
- Validated client certificate
- Validated access service token
- Login method and credentials

The rules are the set of conditions upon which access decisions are made. These can be affirmative rules such that access is granted if the rule is met, such as access from a pre-authorized device. Alternatively, the rules may be written, so access is denied if the rule is met, such as an access request from a specific country or region. The rules can be constructed as logical operations using Boolean algebra for complex constructs. Rules can be a single condition but typically include multiple functions to provide robust security. Typical structures can consist of:

- All conditions must be met before granting access.
- More than A conditions from a set of N conditions (N > A) must be met before granting access.
- Only one condition from a set of conditions must be met before granting access.
- More than A conditions from a set of N conditions (N > A) and less than B conditions from a set of N conditions (N > B) must be met before granting access.
- Only one condition from a set of conditions must be met before denying access.
- More than A conditions from a set of N conditions (N > A) must be met before denying access.
- All conditions must be met before denying access.

Policy examples:

- All requests from email addresses from the domain “@trusted.com” with an Australian IP address are granted access to a service.
- Access requests to financial data from the users in the accounting department logged on using multi-factor authentication (MFA) from a company-assigned end-user device (as recognized by its MAC address) from a predefined static IP address will be granted access.
- Access requests to commercially sensitive company data from users who have not logged on from a company-assigned end-user device (as recognized by its MAC address) will be denied access.

LMNTRIX IDENTITY PROTECTION

LMNTRIX IDENTITY

The LMNTRIX Identity managed protection services are comprised of five key elements that provide a comprehensive intelligence-led identity threat identification, management, and response solution.

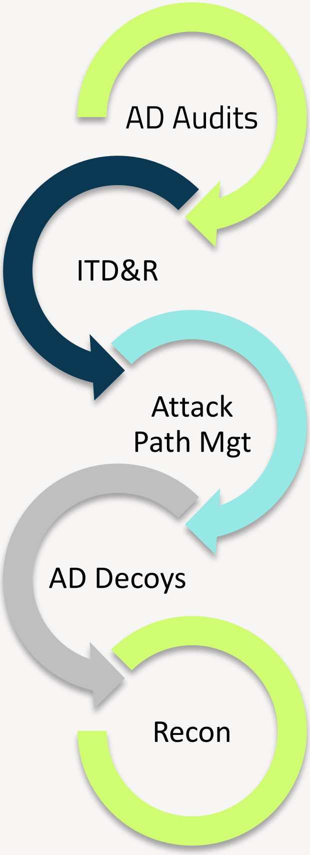

Figure 12 - LMNTRIX Identity Service

1. Active Directory Audits:

- LMNTRIX conducts initial and bi-annual audits of Active Directory (AD), providing a detailed, actionable report highlighting all AD misconfigurations and exposures.

2. Identity Threat Detection & Response (IDDR):

- Our IDDR service collects and consolidates identity information and access logs, integrating seamlessly with the LMNTRIX Extended Detection and Response (XDR) solution. Specialized identity, access, and AD detection rules trigger alerts, monitored 24/7 by expert analysts in our Security Operations Center (SOC), transforming alerts into incidents when necessary.

3. Attack Path Management:

- LMNTRIX Attack Paths identifies, monitors, and manages chains of exploitable privileges and user behaviors. Customized tools, seamlessly integrated into the LMNTRIX XDR solution, provide results that feed into the SOC incident management service.

4. Active Directory Decoys:

- Leveraging the LMNTRIX DECEIVE service, we deploy Active Directory (AD) decoys to generate fictitious credentials and services. Integrated into the LMNTRIX XDR solution, these decoys are monitored for reconnaissance activities, with results seamlessly integrated into the SOC incident management service.

5. Identity Reconnaissance (ID Recon):

- The LMNTRIX RECON service offers comprehensive monitoring for credential breaches, focusing on any identity or credential-based information leaks.

ABOUT LMNTRIX

Often, the difference between preventing a cyber attack or suffering a crippling loss is simply knowing where to look for the signs of a compromise. Even the most advanced attackers leave traces of their presence, so an effective defense must not only be vigilant but also ever-adaptive in response to changes in attacker tactics. A critical element in this age of constantly evolving threats is a detailed view of an organization's entire potential attack surface. Unfortunately, log collection solutions are simply outgunned against today's advanced threat actors as they either lack the data or the ability to analyze their data in a manner that allows rapid attack detection.

LMNTRIX has reimagined cybersecurity, once again turning the tables in favor of the defenders. We have cut out the bloat of SIEM, log analysis, false positives, and associated alert fatigue and created new methods for confounding even the most advanced attackers. We combine deep expertise with cutting-edge technology, leading intelligence, and advanced analytics to detect and investigate threats with great speed, accuracy, and focus. We believe that in a time of continuous compromise, you need continuous response – not incident response.

As a company, we stand in defiance of the unwanted human presence within corporate networks by attacking the root of the problem—the adversary's ability to gain entry and remain undetected. Our real-time hunt operations identify signs of planned and active attacks and take action to neutralize them, forming the basis of our comprehensive Active Defense approach to limiting security exposure.

We are a partner who becomes an extension of your internal team, can augment your MSSP, or be a full-service SOC as a service security solution.

LMNTRIX Active Defense is a three-tier outcome-based solution (Industry refers to it as Managed Detection & Response (MDR) and our platform as Extended Detection & Response (XDR).

(1) LMNTRIX XDR (AWS Data Lake and Platform)

(2) LMNTRIX TECHNOLOGY STACK (Deployed deep within Customer Networks)

(3) LMNTRIX CYBER DEFENSE CENTRE (Security Analyst Driven).

LMNTRIX XDR natively unifies Machine and Underground Intelligence, NGAV, EDR, NDR, UEBA, and Deception Everywhere with completely automated attack validation, investigation, containment, and remediation on a single, intuitive platform. Backed by a 24/7 Managed Detection and Response service – at no extra cost – LMNTRIX provides comprehensive protection of the environment for even the smallest security teams. It is a single investigative platform for insights into threats on enterprise, cloud, hybrid, and industrial control systems (ICS) networks. The LMNTRIX XDR delivers unique advantages over current network security solutions. It is a holistic and multi-vector platform with an unlimited retention window of full-fidelity network traffic, innovative security visualizations, and the ease and cost-savings of an on-demand deployment model.

LMNTRIX XDR is based on multiple detective, responsive, and predictive capabilities that integrate and share information to build a security protection system that is more adaptive and intelligent than any one element. The constant exchange of intelligence between the Active Defense components and the wider cybersecurity community enables LMNTRIX to keep abreast of the TTP of the most persistent, well-resourced, and skilled attack groups.

Figure 13 - LMNTRIX XDR

LMNTRIX TECH STACK is a powerful, proprietary threat detection stack embedded within the client environment behind existing controls. TECHNOLOGY STACK comprises multiple detective systems, combining threat intel application and correlation, static-file analysis, user and entity behavior analytics (UEBA), and anomaly detection techniques to find threats in real time. In addition, it eliminates alert fatigue, determining which alerts to escalate through multi-platform consensus.

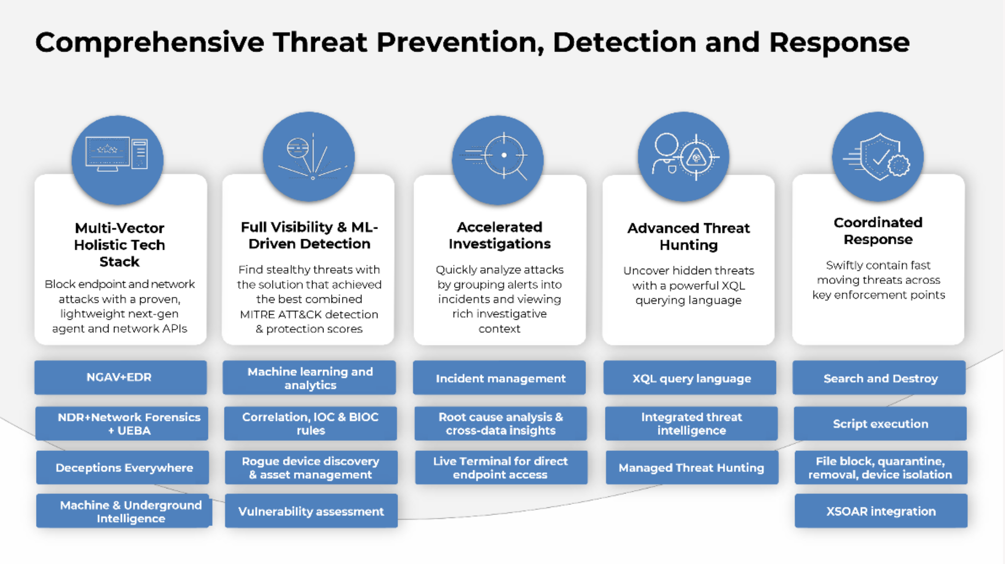

Comprehensive Threat Prevention, Detection and Response

Multi-Vector Holistic Tech Stack
Block endpoint and network attacks with a proven, lightweight next-gen agent and network APIs
- NGAV+EDR
- NDR+Network Forensics + UEBA
- Deceptions Everywhere
- Machine & Underground Intelligence

Full Visibility & ML-Driven Detection
Find stealthy threats with the solution that achieved the best combined MITRE ATT&CK detection & protection scores
- Machine learning and analytics
- Correlation, IOC & BIOC rules
- Rogue device discovery & asset management
- Vulnerability assessment

Accelerated Investigations
Quickly analyze attacks by grouping alerts into incidents and viewing rich investigative context
- Incident management
- Root cause analysis & cross-data insights
- Live Terminal for direct endpoint access

Advanced Threat Hunting
Uncover hidden threats with a powerful XQL querying language
- XQL query language
- Integrated threat intelligence
- Managed Threat Hunting

Coordinated Response
Swiftly contain fast moving threats across key enforcement points
- Search and Destroy
- Script execution
- File block, quarantine, removal, device isolation
- XSOAR integration

Figure 14 - LMNTRIX XDR – A Comprehensive Threat Prevention, Detection & Response Platform

LMNTRIX CYBER DEFENSE CENTER (CDC)

A global network of Cyber Defense Centers comprising trained and certified hunters and intrusion analysts provides constant vigilance and on-demand analysis of your digital assets and networks. Our intrusion analysts actively probe and monitor your networks and endpoints 24x7, using the latest intelligence and proprietary methodologies to look for signs of compromise. When a suspected breach is detected, the team performs an in-depth analysis of potentially affected systems to confirm the breach. Additionally, when data theft or lateral movement is imminent, our endpoint containment feature makes immediate action possible by quarantining affected hosts, whether they are on or off your corporate network. This significantly reduces or eliminates the consequences of a breach.

Figure 15 - LMNTRIX Cyber Defense Centre

APPENDIX - IDENTITY THREAT DETECTION & RESPONSE RULES

The following infrastructure-specific rules for implementing ITDR offer comprehensive identity threat coverage and represent the minimum recommended ruleset for use when selecting an identity protection service. The rule names and descriptions are provided by ELASTICSEARCH BV, provider of Elastic, the leading platform for search-powered solutions.

Recommended Microsoft 365 Rules	
Rule Name	Description
Attempts to Brute Force a Microsoft 365 User Account	This rule identifies attempts to brute force a Microsoft 365 user account. An adversary may attempt a brute force attack to obtain unauthorized access to user accounts.
Microsoft 365 Exchange Anti-Phish Policy Deletion	This rule identifies the deletion of an anti-phishing policy in Microsoft 365. By default, Microsoft 365 includes built-in features that help protect users from phishing attacks. Anti-phishing policies increase this protection by refining settings to detect and prevent attacks better.
Microsoft 365 Exchange Anti-Phish Rule Modification	This rule identifies the modification of an anti-phishing rule in Microsoft 365. By default, Microsoft 365 includes built-in features that help protect users from phishing attacks. Anti-phishing rules increase this protection by refining settings to detect and prevent attacks better.
Microsoft 365 Exchange DKIM Signing Configuration Disabled	v when a DomainKeys Identified Mail (DKIM) signing configuration is disabled in Microsoft 365. With DKIM in Microsoft 365, messages sent from Exchange Online will be cryptographically signed. This will allow the receiving email system to validate that the messages were generated by a server that the organization authorized and were not spoofed
Microsoft 365 Exchange DLP Policy Removed	This rule identifies when a Data Loss Prevention (DLP) policy is removed in Microsoft 365. An adversary may remove a DLP policy to evade existing DLP monitoring.
Microsoft 365 Exchange Malware Filter Policy Deletion	This rule identifies when a malware filter policy has been deleted in Microsoft 365. A malware filter policy will alert administrators that an internal user sent a message containing malware. This may indicate an account or machine compromise that must be investigated. Deletion of a malware filter policy may be done to evade detection.

Microsoft 365 Exchange Malware Filter Rule Modification	This rule identifies when a malware filter rule has been deleted or disabled in Microsoft 365. An adversary or insider threat may want to modify a malware filter rule to evade detection.
Microsoft 365 Exchange Management Group Role Assignment	This rule identifies when a new role is assigned to a management group in Microsoft 365. An adversary may attempt to add a role to maintain persistence in an environment.
Microsoft 365 Exchange Safe Attachment Rule Disabled	This rule identifies when a safe attachment rule is disabled in Microsoft 365. Safe attachment rules can extend malware protections to include routing all messages and attachments without a known malware signature to a special hypervisor environment. An adversary or insider threat may disable a safe attachment rule to exfiltrate data or evade defenses.
Microsoft 365 Exchange Safe Link Policy Disabled	This rule identifies when a Safe Link policy is disabled in Microsoft 365. Safe Link policies for Office applications extend phishing protection to documents that contain hyperlinks, even after they have been delivered to a user.
Microsoft 365 Exchange Transport Rule Creation	This rule identifies a transport rule creation in Microsoft 365. As a best practice, Exchange Online mail transport rules should not be set to forward emails to domains outside your organization. An adversary may create transport rules to exfiltrate data.
Microsoft 365 Exchange Transport Rule Modification	This rule identifies when a transport rule has been disabled or deleted in Microsoft 365. Mail flow rules (also known as transport rules) are used to identify and take action on messages that flow through your organization. An adversary or insider threat may modify a transport rule to exfiltrate data or evade defenses.
Microsoft 365 Global Administrator Role Assigned	In Azure Active Directory (Azure AD), permissions to manage resources are assigned using roles. The Global Administrator role enables users to access all administrative features in Azure AD and services that use Azure AD identities, like the Microsoft 365 Defender portal, the Microsoft 365 compliance center, Exchange, SharePoint Online, and Skype for Business Online. Attackers can add users as Global Administrators to maintain access and manage all subscriptions and their settings and resources.
Microsoft 365 Inbox Forwarding Rule Created	This rule identifies when a new Inbox forwarding rule is created in Microsoft 365. Inbox rules process messages in the Inbox based on conditions and take actions. The rules will forward the emails to a defined address in this case. Attackers can abuse Inbox Rules to intercept and exfiltrate email data without making organization-wide configuration changes or having the corresponding privileges.

Microsoft 365 Potential ransomware activity	This rule identifies when Microsoft Cloud App Security reports that a user has uploaded files to the cloud that might be infected with ransomware.
Microsoft 365 Teams External Access Enabled	This rule identifies when external access is enabled in Microsoft Teams. External access lets Teams and Skype for Business users communicate with other users outside their organization. An adversary may enable external access or add an allowed domain to exfiltrate data or maintain persistence in an environment.
Microsoft 365 Teams Guest Access Enabled	This rule identifies when guest access is enabled in Microsoft Teams. Guest access in Teams allows people outside the organization to access teams and channels. An adversary may enable guest access to maintain persistence in an environment.
Microsoft 365 Unusual Volume of File Deletion	This rule identifies that a user has deleted an unusually large volume of files, as reported by Microsoft Cloud App Security.
Microsoft 365 User Restricted from Sending Email	This rule identifies when a user has been restricted from sending email due to exceeding the sending limits of the service policies per the Security Compliance Center.
New or Modified Federation Domain	This rule identifies a new or modified federation domain, which can be used to create trust between O365 and an external identity provider.
O365 Email Reported by User as Malware or Phish	This rule detects the occurrence of emails reported as Phishing or Malware by Users. Security Awareness training is essential to stay ahead of scammers and threat actors, as security products can be bypassed, and the user can still receive a malicious message. Educating users to report suspicious messages can help identify gaps in security controls and prevent malware infections and Business Email Compromise attacks.
O365 Excessive Single Sign-On Logon Errors	This rule identifies accounts with a high number of single sign-on (SSO) login errors. Excessive login errors may indicate an attempt to brute force a password or SSO token.
O365 Exchange Suspicious Mailbox Right Delegation	This rule identifies the assignment of rights to access content from another mailbox. An adversary may use the compromised account to send messages to other accounts in the target organization's network while creating inbox rules so that messages can evade spam/phishing detection mechanisms.

O365 Mailbox Audit Logging Bypass	This rule detects the occurrence of mailbox audit bypass associations. The mailbox audit is responsible for logging specified mailbox events (like accessing a folder or a message or permanently deleting a message). However, actions taken by some authorized accounts, such as accounts used by third-party tools or accounts used for lawful monitoring, can create a large number of mailbox audit log entries and may not be of interest to your organization. Because of this, administrators can create bypass associations, allowing certain accounts to perform their tasks without being logged. Attackers can abuse this allowlist mechanism to conceal actions taken, as the mailbox audit will log no activity done by the account.
Potential Password Spraying of Microsoft 365 User Accounts	This rule identifies a high number (25) of failed Microsoft 365 user authentication attempts from a single IP address within 30 minutes, which could indicate a password-spraying attack. An adversary may attempt a password-spraying attack to obtain unauthorized access to user accounts.
Recommended Amazon Web Services Rules	
Rule Name	Description
AWS IAM Assume Role Policy Update	This rule identifies attempts to modify an AWS IAM Assume Role Policy. An adversary may attempt to modify the AssumeRolePolicy of a misconfigured role to gain the privileges of that role.
AWS IAM Brute Force of Assume Role Policy	This rule identifies a high number of failed attempts to assume an AWS Identity and Access Management (IAM) role. IAM roles are used to delegate access to users or services. An adversary may attempt to enumerate IAM roles to determine if a role exists before attempting to assume or hijack the discovered role.
AWS IAM Deactivation of MFA Device	This rule identifies the deactivation of a specified multi-factor authentication (MFA) device and removes it from association with the username for which it was originally enabled. In AWS Identity and Access Management (IAM), a device must be deactivated before it can be deleted.
AWS IAM Group Creation	This rule identifies the creation of a group in AWS Identity and Access Management (IAM). Groups specify permissions for multiple users. Any user in a group automatically has permission assigned to the group.
AWS IAM Group Deletion	This rule identifies the deletion of a specified AWS Identity and Access Management (IAM) resource group. Deleting a resource group does not delete resources that are group members; it only deletes the group structure.
AWS IAM Password Recovery Requested	This rule identifies AWS IAM password recovery requests. An adversary may attempt to gain unauthorized AWS access by abusing password recovery mechanisms.

AWS IAM User Addition to Group	This rule identifies the addition of a user to a specified group in AWS Identity and Access Management (IAM).
AWS Management Console Brute Force of Root User Identity	This rule identifies a high number of failed authentication attempts to the AWS management console for the Root user identity. An adversary may attempt to brute force the password for the Root user identity, as it has complete access to all services and resources for the AWS account.
AWS Management Console Root Login	This rule identifies a successful login to the AWS Management Console by the Root user.
AWS Root Login Without MFA	This rule identifies attempts to log in to AWS as the root user without using multi-factor authentication (MFA). Amazon AWS's best practices indicate that MFA should protect the root user.
AWS SAML Activity	This rule identifies when SAML activity has occurred in AWS. An adversary could manipulate SAML to maintain access to the target.
AWS STS GetSessionToken Abuse	This rule identifies the suspicious use of GetSessionToken. Tokens could be created and used by attackers to move laterally and escalate privileges.
AWS Security Group Configuration Change Detection	This rule identifies a change to an AWS Security Group Configuration. A security group is like a virtual firewall; modifying configurations may allow unauthorized access. Threat actors may abuse this to establish persistence, exfiltrate data, or pivot in an AWS environment.
AWS Security Token Service (STS) AssumeRole Usage	This rule identifies the use of AssumeRole. AssumeRole returns a set of temporary security credentials that can be used to access AWS resources. An adversary could use those credentials to move laterally and escalate privileges.
Recommended Microsoft Azure Rules	
Rule Name	Description
Azure Active Directory High-Risk Sign-in	This rule identifies high-risk Azure Active Directory (AD) sign-ins by leveraging Microsoft's Identity Protection machine learning and heuristics. Identity Protection categorizes risk into three tiers: low, medium, and high. While Microsoft does not provide specific details about how risk is calculated, each level brings higher confidence that the user or sign-in is compromised.

Azure AD Global Administrator Role Assigned	In Azure Active Directory (Azure AD), permissions to manage resources are assigned using roles. The Global Administrator role enables users to access all administrative features in Azure AD and services that use Azure AD identities, like the Microsoft 365 Defender portal, the Microsoft 365 compliance center, Exchange, SharePoint Online, and Skype for Business Online. Attackers can add users as Global Administrators to maintain access and manage all subscriptions and their settings and resources.
Azure Active Directory High-Risk Sign-in	This rule identifies high-risk Azure Active Directory (AD) sign-ins by leveraging Microsoft's Identity Protection machine learning and heuristics. Identity Protection categorizes risk into three tiers: low, medium, and high. While Microsoft does not provide specific details about how risk is calculated, each level brings higher confidence that the user or sign-in is compromised.
Azure Active Directory PowerShell Sign-in	This rule identifies a sign-in using the Azure Active Directory PowerShell module. PowerShell for Azure Active Directory allows for managing settings from the command line, intended for users who are members of an admin role.
Azure Application Credential Modification	This rule identifies when a new credential is added to an application in Azure. An application may use a certificate or secret string to prove its identity when requesting a token. Multiple certificates and secrets can be added to an application, and an adversary may abuse this by creating an additional authentication method to evade defenses or persist in an environment.
Azure Automation Account Created	This rule identifies when an Azure Automation account is created. Azure Automation accounts can be used to automate management tasks and orchestrate actions across systems. An adversary may create an Automation account to maintain persistence in their target's environment.
Azure Conditional Access Policy Modified	This rule identifies when an Azure Conditional Access policy is modified. Azure Conditional Access policies control access to resources via if-then statements. For example, if a user wants to access a resource, they must complete an action, such as using multi-factor authentication. An adversary may modify a Conditional Access policy to weaken their target's security controls.
Azure External Guest User Invitation	This rule identifies an invitation to an external user in Azure Active Directory (AD). Azure AD is extended to include collaboration, allowing you to invite people outside your organization to be guest users in your cloud account. Unless a business needs to provide guest access, it is best practice to avoid creating guest users. Guest users could potentially be overlooked indefinitely, leading to a potential vulnerability.

Azure Global Administrator Role Addition to PIM User	This rule identifies an Azure Active Directory (AD) Global Administrator role in addition to a Privileged Identity Management (PIM) user account. PIM is a service that enables you to manage, control, and monitor access to important resources in an organization. Users assigned to the Global administrator role can read and modify any administrative setting in your Azure AD organization.
Azure Blob Permissions Modification	This rule identifies when the Azure role-based access control (Azure RBAC) permissions are modified for an Azure Blob. An adversary may modify the permissions on a blob to weaken their target's security controls, or an administrator may inadvertently modify the permissions, which could lead to data exposure or loss.
Azure Privilege Identity Management Role Modified	Azure Active Directory (AD) Privileged Identity Management (PIM) is a service that enables you to manage, control, and monitor access to important resources in an organization. PIM can be used to manage the built-in Azure resource roles such as Global Administrator and Application Administrator. An adversary may add a user to a PIM role to maintain persistence in their target's environment or modify a PIM role to weaken their target's security controls.
Azure Service Principal Addition	This rule identifies when a new service principal is added in Azure. An application, hosted service, or automated tool that accesses or modifies resources needs an identity created. This identity is known as a service principal. For security reasons, using service principals with automated tools is always recommended rather than allowing them to log in with a user identity.
Azure Storage Account Key Regenerated	This rule identifies a rotation to storage account access keys in Azure. Regenerating access keys can affect any applications or Azure services dependent on the storage account key. Adversaries may regenerate a key as a means of acquiring credentials to access systems and resources.
Azure Service Principal Credentials Added	This rule identifies when new Service Principal credentials have been added in Azure. In most organizations, credentials will be added to service principals infrequently. Hijacking an application (by adding a rogue secret or certificate) with granted permissions will allow the attacker to access data normally protected by MFA requirements.
Azure Kubernetes Rolebindings Created	This rule identifies the creation of role binding or cluster role bindings. You can assign these roles to Kubernetes subjects (users, groups, or service accounts) with role bindings and cluster role bindings. An adversary who has permission to create bindings and cluster-bindings in the cluster can create a binding to the cluster-admin ClusterRole or other high-privileged roles.

Multi-Factor Authentication Disabled for an Azure User	This rule identifies when multi-factor authentication (MFA) is disabled for an Azure user account. An adversary may disable MFA for a user account to weaken the account's authentication requirements.
Possible Consent Grant Attack via Azure-Registered Application	This rule detects when a user grants permissions to an Azure-registered application or when an administrator grants tenant-wide permissions to an application. An adversary may create an Azure-registered application that requests access to data such as contact information, email, or documents.
Azure Active Directory High-Risk User Sign-in Heuristic	This rule identifies high-risk Azure Active Directory (AD) sign-ins by leveraging Microsoft Identity Protection machine learning and heuristics.
Recommended Google Cloud Platform Rules	
Rule Name	Description
GCP IAM Custom Role Creation	This rule identifies an Identity and Access Management (IAM) custom role creation in the Google Cloud Platform (GCP). Custom roles are user-defined and allow for the bundling of one or more supported permissions to meet specific needs. Custom roles will not be updated automatically and could lead to privilege creep if not carefully scrutinized.
GCP IAM Role Deletion	This rule identifies an Identity and Access Management (IAM) role deletion in the Google Cloud Platform (GCP). A role contains a set of permissions that allows you to perform specific actions on Google Cloud resources. An adversary may delete an IAM role to inhibit access to accounts utilized by legitimate users.
GCP IAM Service Account Key Deletion	This rule identifies the deletion of an Identity and Access Management (IAM) service account key in the Google Cloud Platform (GCP). Each service account is associated with two sets of public/private RSA key pairs used to authenticate. If a key is deleted, the application will no longer be able to access Google Cloud resources using that key. A security best practice is to rotate your service account keys regularly.
GCP Service Account Creation	This rule identifies when a new service account is created in the Google Cloud Platform (GCP). A service account is a special type of account used by an application or a virtual machine (VM) instance, not a person. Applications use service accounts to make authorized API calls, authorized as either the service account itself or as G Suite or Cloud Identity users through domain-wide delegation. Service accounts can present a security risk if they are not tracked and managed properly. An adversary may create a new service account to use during their operations to avoid using a standard user account and attempt to evade detection.

GCP Service Account Deletion	This rule identifies when a service account is deleted in the Google Cloud Platform (GCP). A service account is a special type of account used by an application or a virtual machine (VM) instance, not a person. Applications use service accounts to make authorized API calls, authorized as either the service account itself or as G Suite or Cloud Identity users through domain-wide delegation. An adversary may delete a service account to disrupt their target's business operations.
GCP Service Account Disabled	This rule identifies when a service account is disabled in the Google Cloud Platform (GCP). A service account is a special type of account used by an application or a virtual machine (VM) instance, not a person. Applications use service accounts to make authorized API calls, authorized as either the service account itself or as G Suite or Cloud Identity users through domain-wide delegation. An adversary may disable a service account to disrupt their target's business operations.
GCP Service Account Key Creation	This rule identifies when a new key is created for a service account in the Google Cloud Platform (GCP). A service account is a special type of account used by an application or a virtual machine (VM) instance, not a person. Applications use service accounts to make authorized API calls, authorized as either the service account itself or as G Suite or Cloud Identity users through domain-wide delegation. If private keys are not tracked and managed properly, they can present a security risk. An adversary may create a new key for a service account to attempt to abuse the permissions assigned to that account and evade detection.
GCP Storage Bucket Configuration Modification	This rule identifies when the configuration is modified for a Google Cloud Platform (GCP) storage bucket. An adversary may modify the configuration of a storage bucket to weaken the security controls of their target's environment.
GCP Storage Bucket Permissions Modification	This rule identifies when the Identity and Access Management (IAM) permissions are modified for a Google Cloud Platform (GCP) storage bucket. An adversary may modify the permissions on a storage bucket to weaken their target's security controls, or an administrator may inadvertently modify the permissions, which could lead to data exposure or loss.
Recommended Google Workspace Rules	
Rule Name	Description
Google Workspace API Access Granted via Domain-Wide Delegation of Authority	This rule detects when a domain-wide delegation of authority is granted to a service account. Domain-wide delegation can be configured to grant third-party and internal applications access to the data of Google Workspace users. An adversary may configure domain-wide delegation to maintain access to their target's data.

Google Workspace Admin Role Assigned to a User	Assigning the administrative role to a user will grant them access to the Google Admin console and grant them administrator privileges allowing them to access and manage various resources and applications. An adversary may create a new administrator account for persistence or apply the admin role to an existing user to carry out further intrusion efforts. Users with super-admin privileges can bypass a single sign on if enabled in Google Workspace.
Google Workspace Admin Role Deletion	This rule detects when a custom admin role is deleted. An adversary may delete a custom admin role to impact system administrators' permissions or capabilities.
Google Workspace Custom Admin Role Created	This rule detects when a custom admin role is created in Google Workspace. An adversary may create a custom admin role to elevate the permissions of other user accounts and persist in their target's environment.
Google Workspace Password Policy Modified	This rule detects when a Google Workspace password policy is modified. An adversary may attempt to modify a password policy to weaken an organization's security controls.
Google Workspace Role Modified	This rule detects when a custom admin role or its permissions are modified. An adversary may modify a custom admin role to elevate the permissions of other user accounts and persist in their target's environment.
Google Workspace User Group Access Modified to Allow External Access	User groups in Google Workspace are created to help manage users' permissions and access to various resources and applications. The security label is only applied to a group when users within that group are expected to access sensitive data and/or resources, so administrators add this label to manage security groups better easily. Adversaries with administrator access may modify a security group to allow external access from members outside the organization. This detection does not capture all modifications to security groups but only those that could increase the associated risk.
MFA Disabled for Google Workspace Organization	This rule detects when multi-factor authentication (MFA) is disabled for a Google Workspace organization. An adversary may attempt to modify a password policy to weaken an organization's security controls.