Defending the Enterprise with Conditional Access Anywhere

An introduction to adaptive and threat-aware conditional access that helps organisations reduce risk by understanding identity everywhere.

TABLE OF CONTENTS

INTRODUCTION

More than ever, both identity and security teams need practical approaches to protect their users and devices from cyberattacks, breaches, and insider threats, without disrupting their business or overloading analysts. Today, identity and access products lack insight into threats while security tools tend to generate many inconclusive alerts that require manual investigation.

Similarly, enforcement methods remain limited to simple Allow or Deny responses that lack an understanding of behaviour and the constantly changing contexts of risk. Additionally, as organisations shift to the cloud, many security teams have lost consistent visibility into their users and assets. Teams often lack visibility and control over behaviours in the cloud, or in the best case, rely on separate and siloed solutions that lack context.

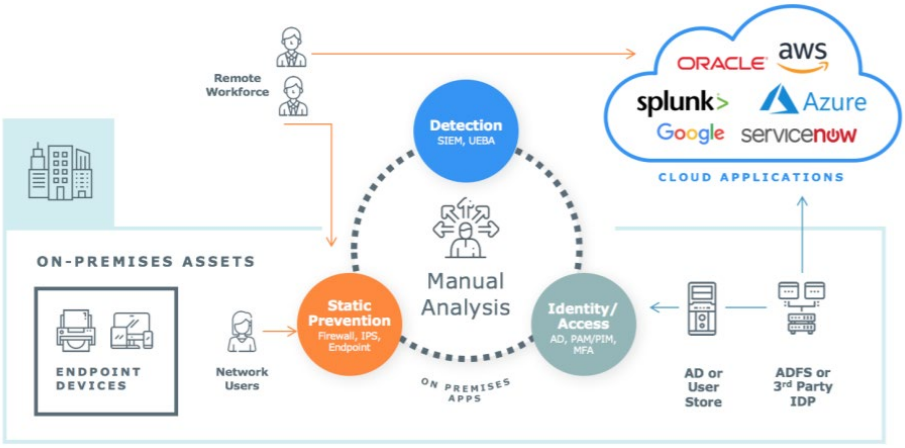

The approach discussed in this document, bridges these perspectives and enables organisations to deliver real-time conditional access and security controls that prevent threats based on identity, behaviour, and risk. Conditional access pairs user behaviour to detect threats with a contextual automated response that redirects risky user behaviour and proactively stop threats without disrupting the business.

The solution automatically learns an entity's role, tracks behaviours over time, and applies flexible policies that can automatically adapt as circumstances change. Flexible response options can make gradual responses to changing risk, automatically close incidents related to benign anomalies, or block once a threat is verified.

This solution white paper will dive deeper into the following areas:

- Identity and Risk Insights
- Threat Detection and Analytics
- Conditional Access Anywhere
- Flexible Deployment and Journey

IDENTIFY AND RISK INSIGHTS

Data breaches often begin by an attacker finding an initial weakness or vulnerability in the target environment such as an employee using a weak or compromised password or a stealthy admin using an unmanaged endpoint to name a few. In order to defend their environment, security teams need full visibility into their attack surface, including all users and accounts.

See and Verify All Accounts

Security staff need visibility into all of the accounts operating on the network. In addition to the traditional human users of a network, security teams must be aware of the many programmatic service accounts in use. These accounts often have high privileges and can be a valuable target for attackers.

Conditional access sensors directly analyse traffic to detect the true nature of an entity. This can reveal when an entity might be pretending to be something they are not, such as a human attacker trying to use a programmatic account or a workstation masquerading as a server. This same type of analysis is applied to other classifications such as identifying managed and unmanaged devices, shared devices, and more.

Likewise, the solution automatically detects the many types of users within an environment; and identifies privileged users and administrators based on observed data. This intelligence includes the ability to reveal “stealthy administrators” who may have important privileges but are not part of the official Administrators Group within Active Directory.

Conditional access can also focus on high-value users such as executives who are often targets of attackers. Visibility can likewise be aligned to users based on their functional role or organisational unit in the enterprise.

Risk Scoring

This approach assimilates all of the available perspectives and context into an actionable risk score for every entity (user, account, device) in the network. This can include weaknesses that make the entity more vulnerable to attack, observed risky behaviours, or signs of malicious behaviour.

This score is expressed as a number between 0-10 and represents the likelihood that the activities or posture of an entity can lead to a successful breach by a malicious attacker, or that an insider may be going rogue.

Risk scores are constantly evaluated and updated based on changes in the environment. Some elements of risk score decay slowly over time, while others can be resolved quickly. For example, when a user changes from a weak to a strong password, the user’s risk score decreases in response immediately. Additionally, important users or assets can get a boost in risk scores such as accounts with administrative rights, power users as executives, or even servers with specific critical roles.

Risk scores are used throughout this approach to automate responses and enable staff to investigate incidents or identify problem users in the network quickly. Risk scores can even be used outside the approach by using the API to share risk context with other systems and orchestration platforms.

Insights

Insights Finding and preemptively resolving weaknesses is one of the most important aspects of strong cybersecurity. However, weaknesses can come in a variety of forms – user-specific traits such as password issues, device configuration problems, support for outdated and weak protocols, Active Directory settings, and vulnerability to a wide range of attack techniques.

The Insights page is dedicated to tracking the security posture of the network while making it easy to find and monitor the users, devices, and accounts in the network that pose the greatest risk. Insights make it easy to find a vary array of problems.

Just as an example, this includes finding devices that share the same local admin credentials, users with weak passwords, devices that allow RDP or RPC, accounts being shared by multiple users, and many more. This visibility allows organisations to reduce their attack surface where possible, and closely monitor any assets that could be inviting to an attacker.

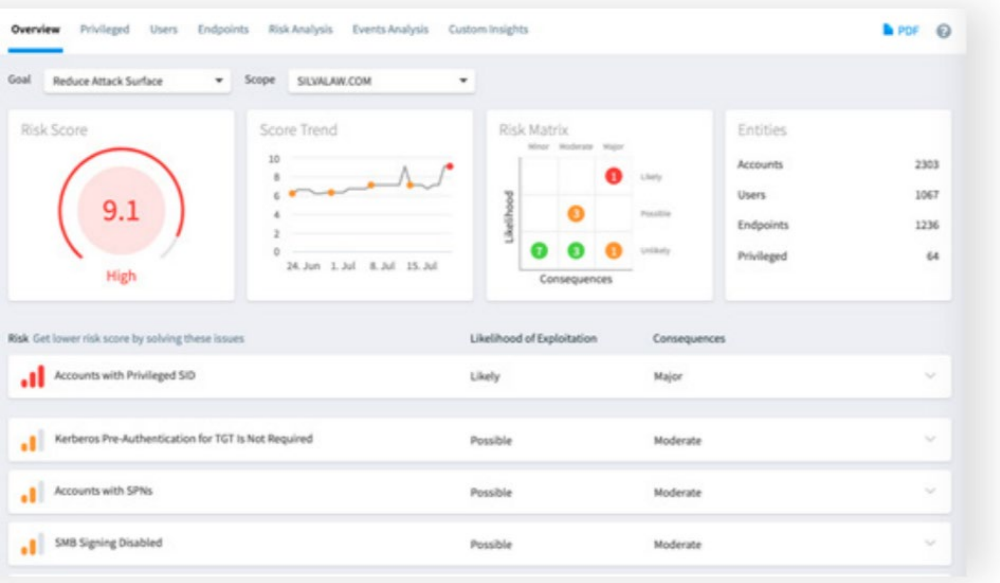

The Insights page contains a variety of pre-built views for seeing risk. Staff can immediately jump to views dedicated to Privileged Users, End Users, Endpoints, as well as a Risk Assessment of all entities. Alternatively, staff can access views tied to specific use cases such as checking overall hygiene of the Active Directory, or attack surface analysis of the network or a particular segment. Each view is interactive, making it easy to drill down into additional information.

The list below includes just some of the weaknesses that this approach can automatically identify:

- A user or admin using an unmanaged endpoint
- Stale privileged account
- Device compatible with versions of NTLM
- SMB signing is disabled
- Devices vulnerable to skeleton key attack
- Multiple devices with the same local administrator
- Accounts only using a DES key
- Password that never expires
- Stealthy administrators
- Using a weak password
- Has a high-risk score
- Using an exposed password

The Insights page also lets staff easily track changes over time. The 'Privileged Users' view could easily reveal when new administrative accounts have been created that otherwise may have been missed. Similarly, in the 'Endpoints' tab, staff can easily see if more unmanaged endpoints have entered the network over a particular period of time.

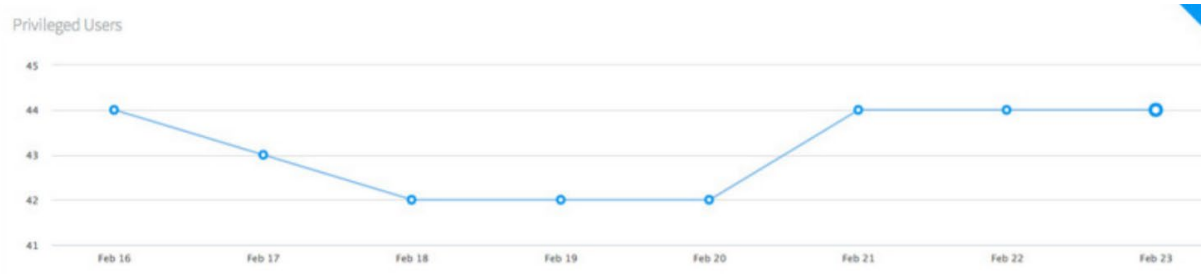

The Insights page then provides a detailed list of each relevant user or device for further investigation. This list is called the Entity Table and provides deep insight into each entity. The table shows if the entity is human, programmatic, or a device along with an organisational unit and risk score.

The table can then offer a wide array of additional detailed traits based on the particular entity, for example, the system could distinguish a particular user as being an executive, a watched user, or a user with multiple devices. A device could be further identified as being an App Server, or a DNS server, or a device with a vulnerable OS. Staff can also create their customised Insights based on virtually any attribute tracked by the solution.

Type	Primary	Secondary	Department	Org. Unit	Attributes	Score
	Amanda Hutchinson	SILVALAW.COM\AHutchinson	Civil	silvalaw.com/NYC_HQ/Corp_...		8.5
	Brian White	SILVALAW.COM\BWhite	Finance	silvalaw.com/NYC_HQ/Corp_...		7.8
	Rita Parker	SILVALAW.COM\RParker	HR	silvalaw.com/NYC_HQ/Corp_...		7.2
	Carol Carson	SILVALAW.COM\CCarson	Security	silvalaw.com/NYC_HQ/Corp_...		7.1
	Paul Taylor	SILVALAW.COM\PTaylor	IT	silvalaw.com/NYC_HQ/Corp_...		7.1
	Sharon Jefferson	SILVALAW.COM\SJefferson	M&A	silvalaw.com/NYC_HQ/Corp_...		7.1
	Laura Henson	SILVALAW.COM\LHenson	IT	silvalaw.com/NYC_HQ/Corp_...		7.1

Risk Analysis

The Insights page also provides staff with a dedicated risk analysis view. This visualisation shows risk scores in relation to the impact on the network, for example, a user may have a high-risk score, but if that user has relatively limited privileges on the network, then the impact could be low. The graph is broken into quadrants with the top-right section representing entities with both high risk as well as high impact. This again provides a very helpful way for staff to hone in on the users that need the most immediate attention.

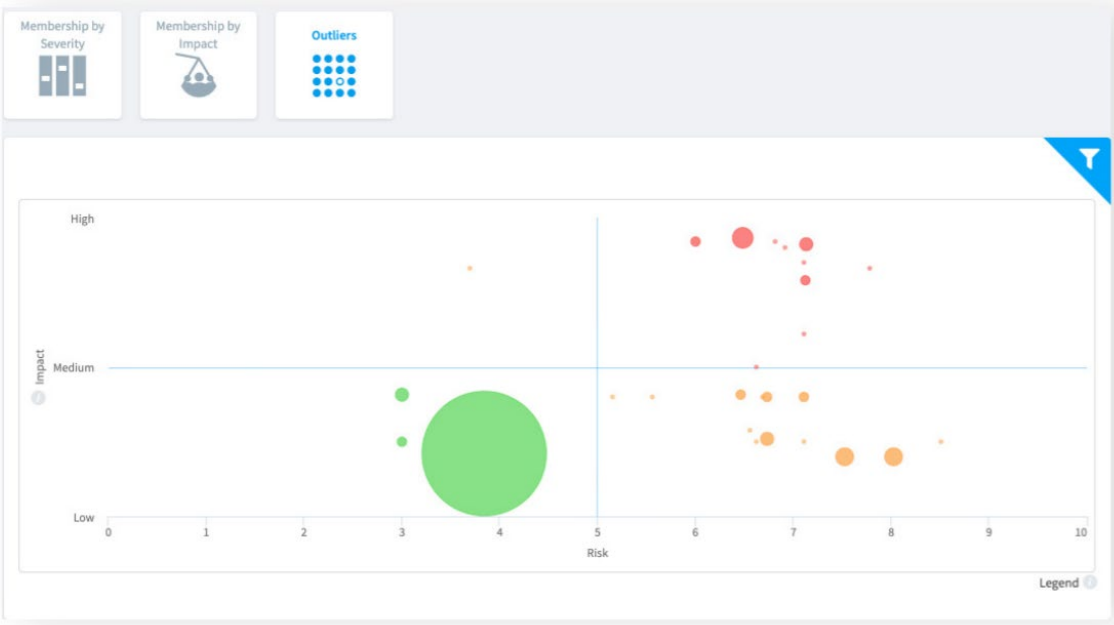

The Risk Analysis page also breaks out risk by Group and Organisational Unit. This allows staff to easily identify the groups and individuals that are contributing the most to the risk of the enterprise. The Outliers view provides additional detail by showing the risk of individual users or entities in the context of impact. For example, a user may have a high-risk score based on a series of unusual or potentially malicious behaviours and have a high impact score based on his access to a high-value database or application.

Reporting

This method provides built-in and customisable reporting for virtually all of the many traits tracked by the platform. This can vastly simplify the sharing of information within the organisation while improving the speed of regular tasks such as compliance reporting. Reports can be configured based on time (weekly, monthly, etc.), and can be emailed or downloaded in PDF format.

THREAT DETECTION AND ANALYTICS

Sensors allow for the direct analysis of traffic travelling to and from the authentication infrastructure. This allows the platform to directly detect a wide range of threats in the environment and to track the behaviour of all entities over time continuously; alongside automatically learning normal behaviour patterns for each entity and identify risky or anomalous behaviour.

Enhance Entity Classification

In addition to information from Active Directory, this approach additionally performs a direct analysis of traffic. This means that every entity is classified with certainty and also recognise when an entity is being impersonated, such as a human user impersonating a service account. This level of analysis also reveals a wealth of context such as being able to associate specific endpoints and device traits with user accounts, such as being able to find an administrator who is using an unmanaged device.

Anomalous Behaviour

The solution continually learns and tracks the behaviour of every user, device, and account in the environment to recognise any abnormal behaviour.

Anomalous behaviour is not necessarily malicious behaviour, but it can often be the first indicator that something is wrong in the environment. For instance, an end-user accessing unusual resources, from an unusual location, or at an unusual time can be a sign that the user is potentially compromised by an attacker or malware.

By learning normal behaviour across a wide variety of traits, anomalies can be flagged, the user's level of risk can be scored, it can be compared to peer's behaviour and then it can challenge the user to confirm identity to ensure that the behaviour is valid.

Malicious Behaviour

Malicious actions and techniques used by attackers to move laterally within a network and extend the intrusion can be directly revealed. This includes techniques such as Pass-the-Hash, Golden Ticket Attacks, Active Directory Harvesting, and attempts to elevate privilege via forged Privileged Account Certificate (PAC), to name a few.

This analysis also reveals the use of common attack tools such as Mimikatz or the use of insecure protocols such as NTLM. These detections can be particularly valuable as they can reveal a more deterministic sign of attack.

Malicious Behavior	Anomalous Behavior	Entity Classification	Security Posture
++ Brute Force	++ Assets accessed	++ Human vs programmatic	++ Weak password
++ Account Scanning	++ Applications or services used	++ Workstation vs Server	++ Exposed password
++ Pass-the-Hash /Ticket	++ Time	++ Managed vs unmanaged	++ Shared account
++ AD Harvesting	++ Location	++ Etc...	++ Stale privileged accounts
++ Forged PAC File	++ Device		++ NTLM use
++ Etc...	++ Etc...		++ Etc...

Investigating Incidents

The Incidents page provides administrators with quick access to the information they need to get work done. Administrators can customise the time range and then further filter based on severity, status (new, open, resolved), and type of user (human or programmatic) or device (managed vs unmanaged). Staff can also search for incidents of specific type or identify incidents that involve a specific user, account, or device.

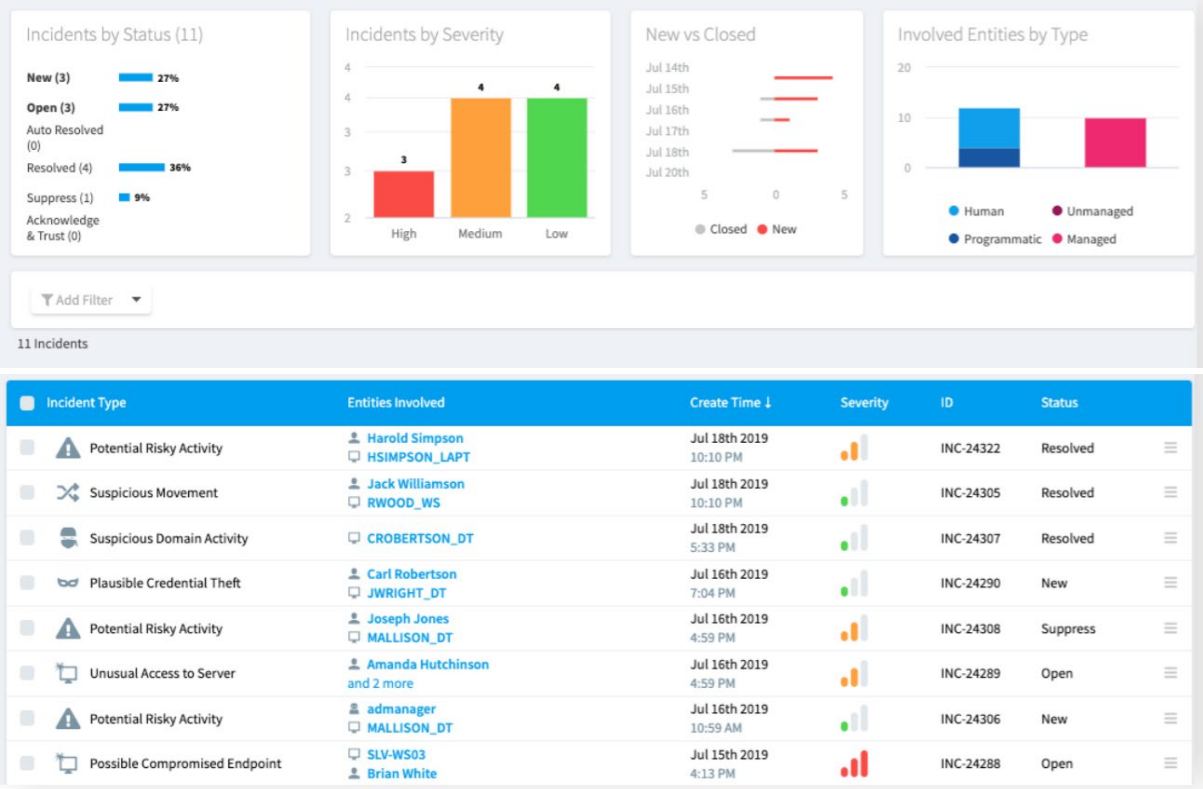

Clicking on a particular incident provides a detailed narrative of the incident and its progression over time. For example, in the Suspicious Movement incident on the following page, the detail shows when the incident was created and that a particular user was observed accessing and an unusual server, and then the next day using an unusual device.

Staff can click to learn more about a particular event within the Incident details. The right-hand side of the screen also shows details about the users or devices involved in the incident along with their overall risk score and last time seen. This view also provides recommendations on next steps and provides a place for administrators to leave comments on the incident.

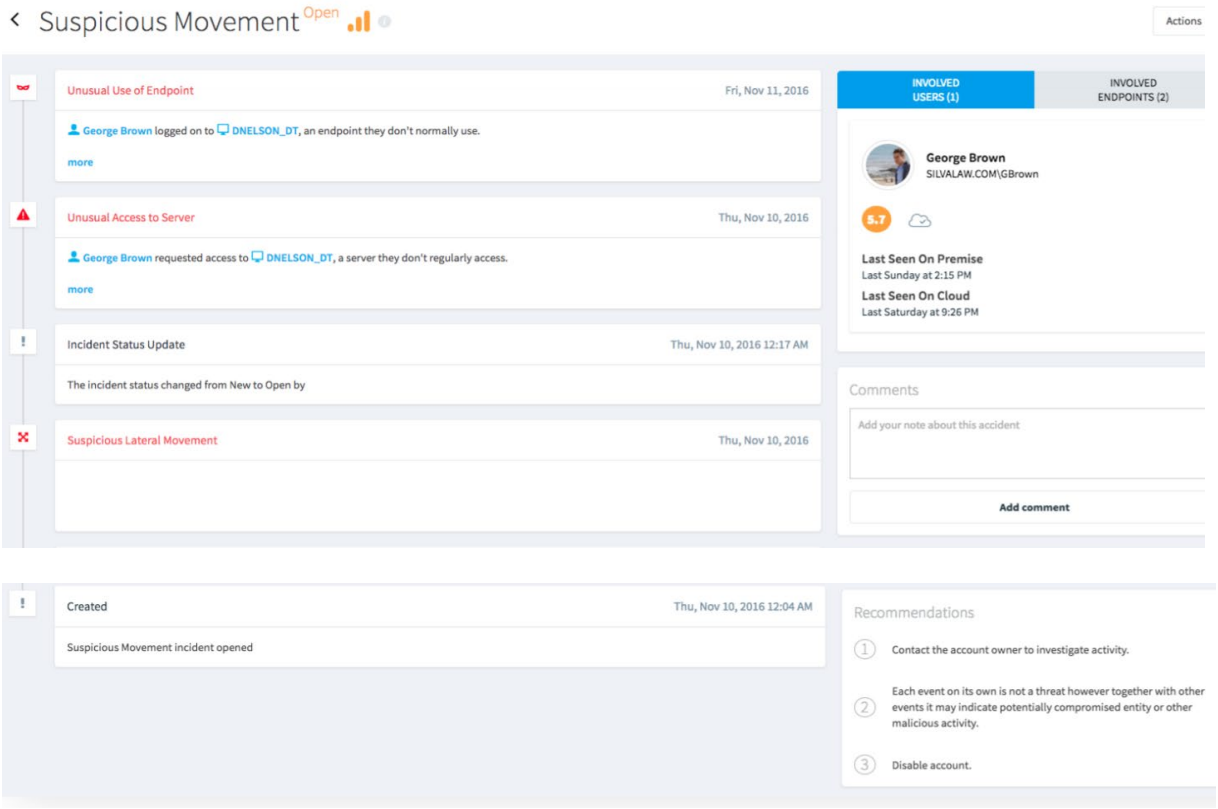

Once an event is investigated, staff can further manage the incident by marking it as resolved, dismissed, or as a false positive. Dismissing an incident will suppress the incident, although conditional access will continue to track event in the background.

If the same event occurs in the future, the incident will be generated again. If an incident is marked a false positive, will learn that the behaviour is allowed and will not generate new incidents in the future.

Threat Hunter

The intuitive interface lets analysts query and correlate across any combination of attributes and network traffic events tracked by the solution. Analysts are free to follow their intuition and ask open-ended questions that cut across user and device attributes, access and authentication methods, account changes, time, location, and more When analysts see something interesting, Threat Hunter can provide any related events and a chronological view to put the details of the hunt into full context.

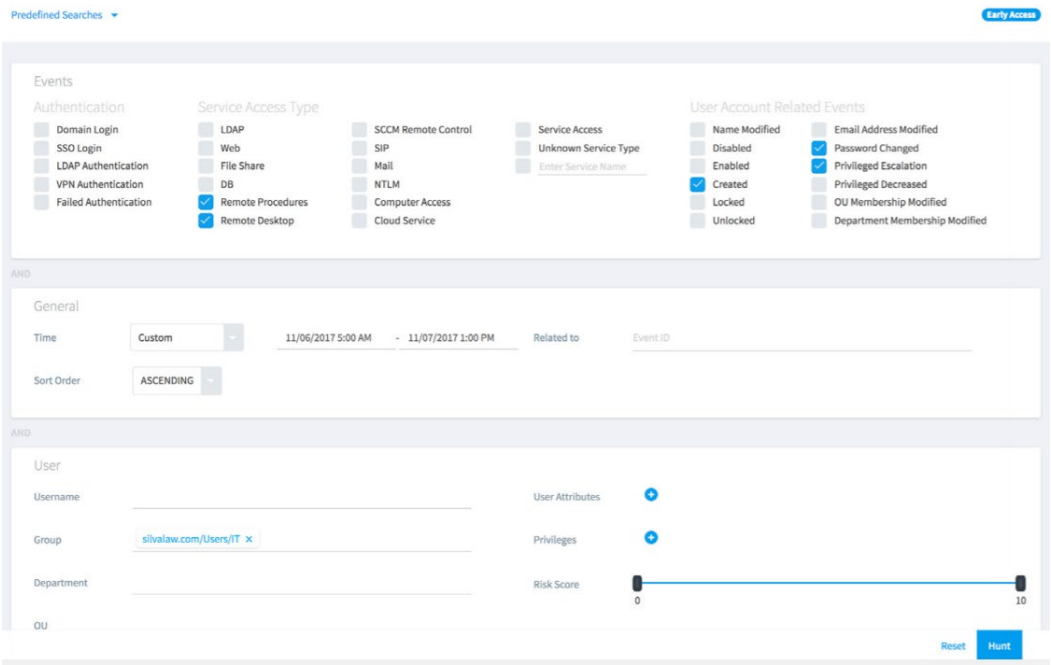

CONDITIONAL ACCESS ANYWHERE

Detecting threats is a critical component of the approach, but ultimately the solution is about turning this intelligence into business-appropriate action. The overarching goal is to deliver responses that are automated without requiring analyst time or impacting valid end users.

To deliver on this goal, a highly adaptable policy engine that can block, challenge users, verify threats, and deliver a variety of real-time graded responses is leveraged. This enables organisations to build powerful conditional access policies to align a wide range of security contexts to real-time and business-appropriate responses.

The Policy Engine

The Policy Engine is the key to conditional access. It brings observed behaviour, user role, risk scores, the target being accessed, access method, and many other factors into a single action-oriented context. Just as importantly, policies have the ability to gain new information and adapt over time.

When abnormal or risky user behaviour is detected, the Policy Engine can automatically challenge users to confirm their identity. Then based on the response, the policy engine could take further action such as isolating the user via NAC, blocking access, or notifying staff. This real-time and adaptive approach to conditional access ensures that actions remain appropriate to the situation without requiring constant attention from analysts.

The Policy Engine takes in a variety of inputs such as any detection rules, user-defined rules or changes to an entity's attributes. Responses can include the ability to block a user, force a password change, or challenge a user with multi-factor authentication.

The results of a challenge can likewise change the user's risk score and also drive further responses. This allows the policy engine and the organisation's response to interact with the user and adapt to the situation logically.

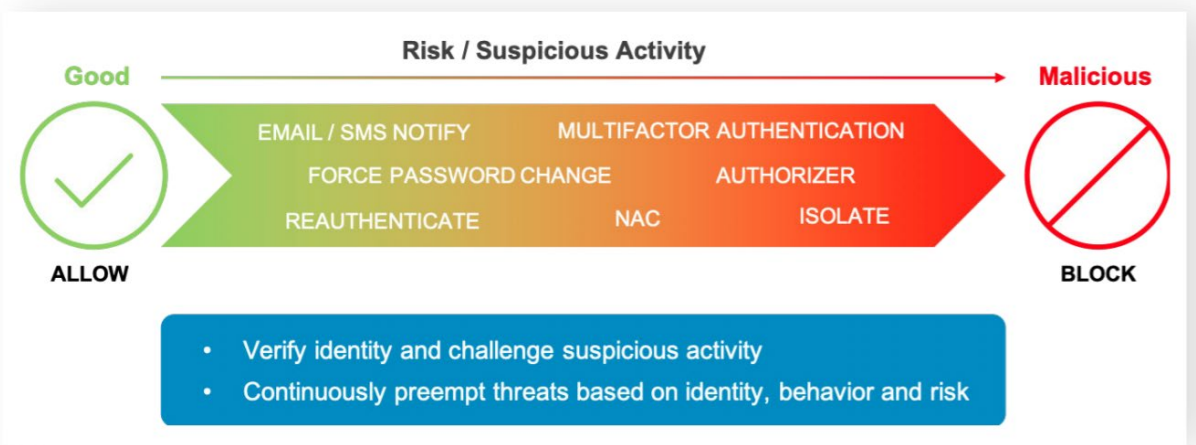

How the Policy Engine Works

Policies are built on a combination of triggers, conditions, and actions. Triggers are the core activity of the policy rules such as an unusual use of an endpoint. Conditions allow the policy to be targeted to a specific situation or use case. For example, a policy for unusual use of endpoint could specifically look for unusual behaviour of devices belonging to executives. Actions specify the automated response or security control to be used when a rule is matched.

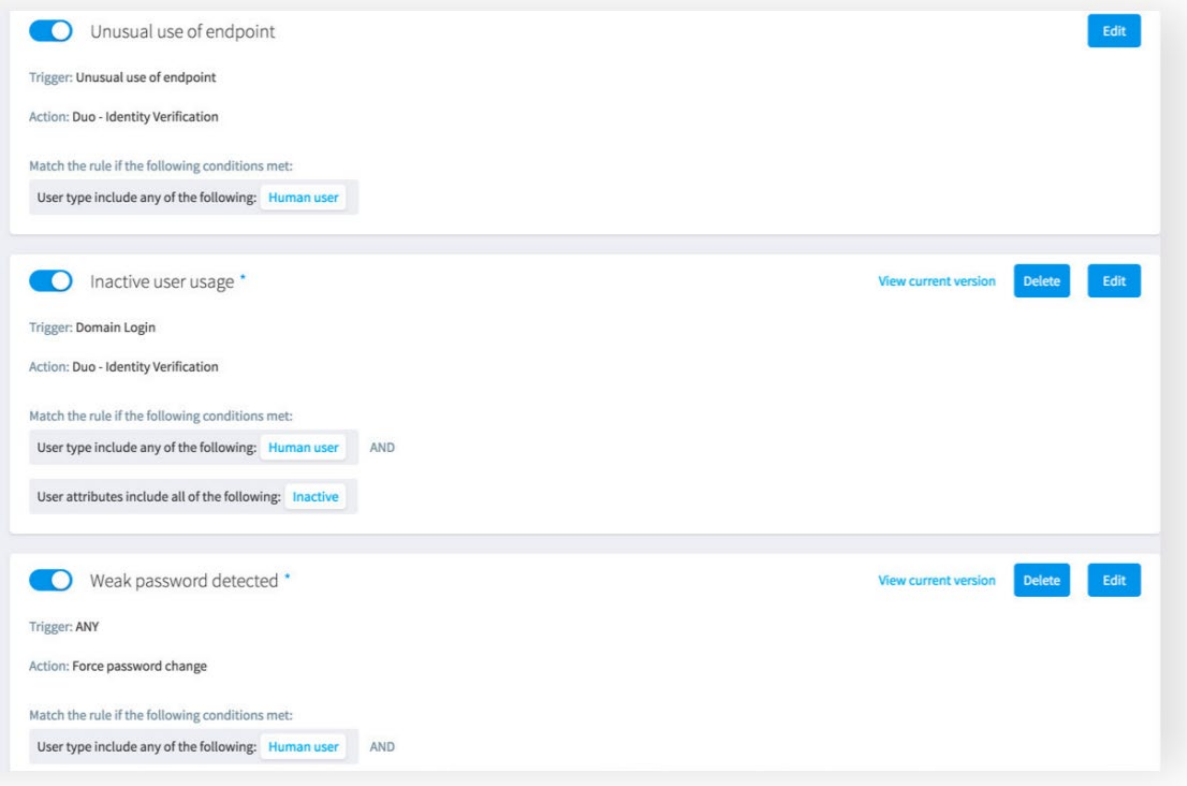

Incorporating identity, role, target and behaviour into the Policy Engine ensures business processes can continue while containing security threats. The table below provides just a few examples of conditional access policies that can easily be created using this approach.

Action	Condition	Trigger	User
Approver required	Any or Specific	Login from new Location	Third Party Vendor, Consultant
Isolate using 3rd party NAC	Any	Risk score > 9	Any
Add user to SSO risk group to limit access to cloud applications	Any	Pass-the-hash detected in on-premise network	Any
Block	Workstation	Remote Desktop	Admin
MFA - Verify identity	Critical Server group	Login	Any user not from jump host
Change password on next login	Any	Weak password detected	Employee

Extending MFA to Any Resource

Organisations are able to extend multi-factor authentication to virtually any resource in the enterprise. Using the solution's network-based approach, teams can add conditional access and MFA controls to custom, legacy, and home-grown applications without making any changes to the applications being protected.

MFA controls can also be added to assets such as databases or workstations without the need to install agents on the devices being protected. Additionally, by integrating with Active Directory Federation Services (ADFS), any federated services accessed via a web browser such as Office 365 or even applications configured for single sign-on (SSO) can be protected.

FLEXIBLE DEPLOYMENT AND JOURNEY

Flexible Deployment and Journey

This approach provides customers with enormous flexibility in terms of how the solution is deployed. At the highest level, the solution can be deployed in three different ways:

- **Sensor-less Deployment** - In this configuration, information is gathered by querying Active Directory servers for important security-related traits and integrating logs from other enterprise sources via APIs. This option provides insights and analysis into user accounts, identifies privileged users, stealthy admins, and a wide variety of password-related issues.
- **Passive Sensor Deployment** - This option leverages a passive sensor, which can be deployed in the network or on the Active Directory servers themselves. This allows traffic travelling to and from the Active Directory infrastructure to be directly analysed. This approach includes all the insights of the sensor-less approach and adds the ability to analyse user behaviours and detect active threats in real-time such as lateral movement, attack tools, and dangerous protocol use.
- **Active Sensor Deployment** - This is the most robust deployment option and includes all the previously described functionality as well as the ability to enforce conditional access policies such as adaptive MFA, blocking of threats and more. An active sensor can be deployed as an in-line network sensor or directly on the Active Directory server.

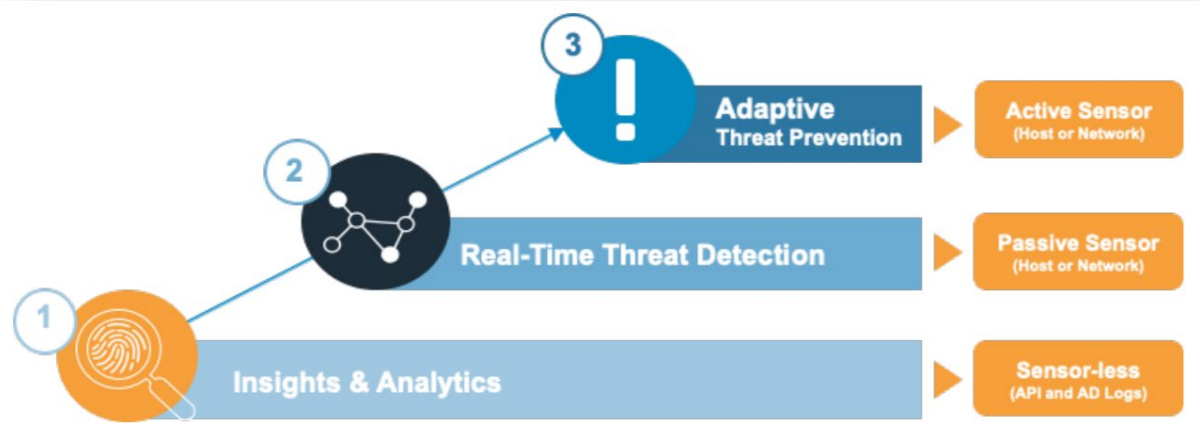

Only one of these options is required at a time. This flexible architecture means that the needs of any environment can be easily aligned while retaining the option to grow as needed. Many customers will deploy initially using active sensors to leverage all the features of the approach, while others can begin with insights and analytics, which requires little more than active directory credentials.

CONCLUSION

This paper is certainly not an exhaustive list of the solution's capabilities. We encourage you to continue to learn more by either seeing a demo or testing the solution in your environment, where we can show in detail how conditional access can help the unique needs of your network.

WHO ARE SECRUTINY?

We are Incident Response specialists who spend 95% of our time making sure our clients don't need to respond to incidents. Secrutiny was founded by three people – all veterans who came to realise there is much confusion in the industry… tools were often overly complex, misconfigured and isolated, and adding further 'layers to the security onion' at significant cost was not the answer.

Through responding to 300+ incidents, we learnt the way to help organisations NOT to be breached is to support them in achieving better security and risk reduction with what they already have; adding capability, where necessary, based on evidence and risk appetite.

We're here to answer any questions you may have and always happy to share our experiences. Reach out to us and we'll respond as soon as we can.

www.secrutiny.com

enquiries@secrutiny.com

0203 8232 999