Security Evolved:

An Expert Guide on Identity and Access Threat Prevention

TABLE OF CONTENTS

GLOSSARY

DNS	Domain Name System
MFA	Multi-Factor Authentication
NTLM	Windows NT LAN Manager
VPN	Virtual Private Network
SIEM	Security Information and Event Management
SMB	Server Message Block
SSO	Single Sign-On

EXECUTIVE SUMMARY

Enterprises and the threats that target them have fundamentally evolved over the past decade. In response, the security industry has generated an enormous amount of point solutions and technologies to try and keep pace. However, for all this innovation and change, the underlying enforcement architecture has remained largely unchanged.

Today the job of real-time policy enforcement and threat prevention still primarily falls to network firewalls and endpoint security products much as it has for the past 20 years. While essential for security, these technologies make very binary allow/block decisions based on a specific and often static event or metadata within a network session or suspicious file.

On the other hand, many of the innovations in security such as machine-learning, AI and behavioural detection algorithms are often inconclusive, reactive, and limited to detection only. As a result, they feed into an increasingly overwhelmed incident response process that attempts to catch up to threats that have already happened.

The challenge is that no matter how fast this process runs, security operations is always inherently a step (or more) behind the attacker. This creates a devil's bargain for security teams where they must either block or allow based on incomplete information, and then try to use additional context after the fact to clean up what was missed.

Instead, we propose a new approach that augments the existing architecture instead of replaces it. This approach brings full enterprise and business context to real-time enforcement decisions. Identity, behaviour, devices, anomalies, and risk all play a real-time role. Just as importantly, enforcement and access options can be graded based on the risk to the business, and policies can actively seek out and adapt to new information.

These changes allow security to evolve in the following ways:

- From reactive response to real-time enforcement.
- From black-and-white rules to policies based on enterprise risk and context.
- From rigid, static rules to policies that continue to adapt to real-time changes.
- From perimeter context to a full enterprise context.
- From end-user enforcement to end-user engagement.

This paper lays out the key concepts of this architecture and how it fits with and extends all of an organisation's security investments.

FINDING THE STRATEGIC LOCATION

As attacks have evolved and matured, more of the battle between attacker and defender occurs inside the network. Attackers often look for any potential initial victim and then begin an ongoing process of using that victim's identity to move deeper into the network, compromise more hosts, and repeat until they gain access to the data or systems they want. This means that the majority of an attack occurs inside the enterprise.

To adapt, we need to bring security enforcement to the inside of the network where the action is. One option would be to extend the perimeter firewall concept to the inside of the network. However, this becomes very expensive and complex and still relies on binary allow/deny segmentation rules.

Instead, we propose to bring cyber security visibility and enforcement to the authentication infrastructure. An organisation's Active Directory infrastructure is the natural nerve centre of an enterprise, governing how users and accounts access applications and assets. From this central point, we can observe identity, application, network, and behavioural traits all in one context and create new logical segmentation strategies based on identity and risk.

Active Directory also governs whether access should be granted or not. By adding cyber security context to this level, we empower security to make real-time enforcement decisions before data is compromised. This opens the door to new types of segmentation based not simply on network boundaries, but on policies that understand the context of identity, behaviour, risk, and virtually any enterprise trait.

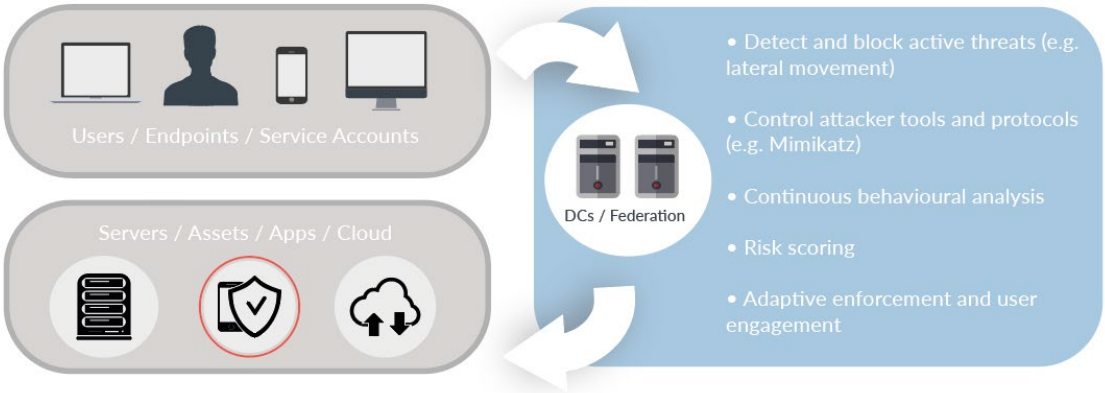

Figure 1: Bringing Threat Prevention to the Identity and Access Infrastructure

Better yet, enforcement at the authentication infrastructure opens up more flexible enforcement options. Instead of blocking a connection outright, we could instead choose to reduce a user's privileges. This could keep the user productive while limiting higher risk behaviours. Additionally, we can interactively challenge suspicious or risky behaviour in real time. For example, a user behaving suspiciously could be required to pass a multi-factor authentication challenge before access is granted to a critical server, see *Figure 2* below.

The authentication infrastructure controls who can connect to what and how they can do it. By bringing advanced threat prevention to this strategic area, we can truly align security to the needs of the organisation.

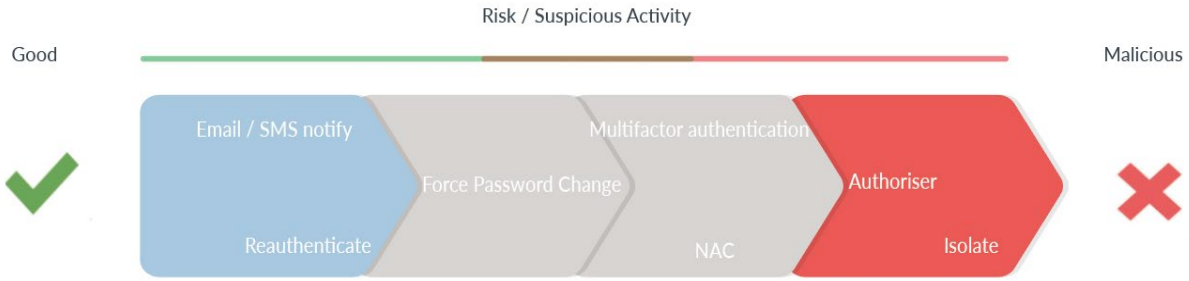

Figure 2: Authentication Infrastructure

Lastly, this approach vastly simplifies many of the deployment challenges of monitoring internal user behaviour. In the past, organisations had two rather cumbersome options. They could deploy agents on every host and/or attempt to tap all the traffic in their internal network. Both options were highly complex, time-consuming and prone to error. However, the directory infrastructure is the logical centre of the enterprise, where all users and accounts must go to gain access to resources. By bringing a security layer to the authentication infrastructure, organisations can get visibility into all behaviour by just updating their internal Domain Name System (DNS) settings. While this is not the only deployment option (passive taps and host-based agents on the domain controllers are other options), it illustrates the advantage of bringing security to the root of enterprise identity.

EXTENDING ENTERPRISE CONTEXT

Once we have established a strong strategic position, we need to gather the right information to make smart decisions. This includes information about the user's identity, privileges and devices, what applications and assets are normally used, and a wide variety of behavioural traits. We should bring this same context to bear whether the enterprise applications are local, in the cloud or a hybrid of the two. Lastly, we should bring in context from the rest of an organisation's security investments, and work with others to automate additional responses. Let's take a look at each of these in order.

Identity and Privilege

Virtually all modern attacks rely on compromising a victim's identity to spread within the network and do damage. Privileged users such as network administrators are the ultimate prize in this regard as their credentials can give an attacker nearly omnipotent control over the network. As a result, understanding a user's role and privileges in the organisation is critical not only for setting smart security policy but also for detecting the signs that the user has been compromised and stopping active attacks.

Key criteria for understanding Identity and Privilege

- Human vs service accounts
- Permissions and privileges
- Password weakness
- Shared password or devices
- Managed or unmanaged devices
- Group membership
- Business privileges

This means we need to know as much as we can about a given account. Is the account a human user or a service account? What organisational unit or groups does the account belong to? For example, it is not uncommon for users outside of the traditional administrator groups to have permissions that allow them to control an administrator account. Even though they are not an "official" administrator, their actual permissions could allow an attacker to gain administrator privileges in the network.

Next, we need to understand the security posture of the entity. Is the account using a weak password or a password that was compromised in another breach? Is the account's password shared with other accounts? Is the user connecting from an unmanaged device?

Collecting this information requires flexibility. Some of the necessary data can be found within active directory itself, and some will only be seen in network traffic. As a result, it is critical to analyse both sources to ensure we accurately identify the “who” in every event.

Key aspects of behavioural context

- *Applications, assets, and protocols typically used*
- *Device history*
- *Geolocation*
- *Normal working times*
- *Authentication or MFA Failures*
- *Stale account usage*

Adding Behavioural Context to Real-Time Response

Next, it’s imperative to understand what a user or account actually does. This could be a single action such as connecting to a sensitive asset, or it could be a pattern of behaviour learned over weeks and months. What locations does a user typically connect from? What are the user’s normal working hours? What devices, applications and assets are normally required? How does a user’s behaviour compare to other users in the same role? Deviations from normal behaviour can potentially indicate the user is doing something risky or possible even compromised by an attacker.

While this information provides critical context, an organisation may not be willing to block a connection simply because the user acted strangely. This is a critical area where the security architecture must adapt to the needs of the business. Responses must be able to automatically challenge anomalous behaviour to help distinguish the unusual from the malicious, and likewise, have the flexibility to gradually respond to the detected level of risk.

Bringing enforcement controls to the authentication infrastructure provides new options and flexibility to meet these needs. For example, anomalous behaviour could trigger a multi-factor authentication (MFA) challenge to verify the user’s identity. If successful, the event can be automatically closed without requiring work from an analyst, and in parallel, the system retrained based on the learned behaviour. If it fails, the user could be demoted, blocked, or require a human authoriser to approve the access. The most important part is that the enforcement options be able to gradually scale in relation to the observed risk. Only then will the security architecture be able to support the needs of both the business and security operations.

STOPPING MALICIOUS BEHAVIOUR

The ability to compromise valid credentials and identities has become a fundamental aspect of virtually every phase of a modern attack. A recent analysis showed that over 80% of breaches involved the use of compromised credentials. And while a compromised account may present in the form of anomalous user behaviour, there are many hacker techniques that are unambiguously malicious. For example, using tools like Mimikatz have become standard practice for attackers to steal credentials and move laterally within a network by using techniques such as Pass-the-Hash, Pass-the-Ticket, or various relay attacks. Attackers will also perform a wide variety of reconnaissance such as account scanning and credential spraying to find new targets or credentials, while techniques such as Golden Ticket attacks can allow an attacker to achieve near-permanent persistence within a network.

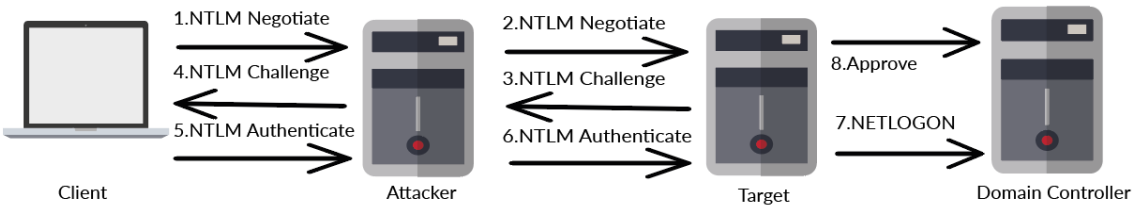

Figure 3: Relay attacks abuse credentials by intercepting and relaying valid challenges and responses in Windows NT LAN Manager (NTLM), Server Message Block (SMB) and other protocols.

These techniques are the difference between a threat that is limited to a single host and a threat that exposes the entire enterprise and all of its assets. Simply put, unless organisations can find and stop these techniques in real time, they can't prevent damage. By bringing actual cyber threat detection and prevention to the seat of enterprise identity, organisations can reclaim the advantage instead of continually playing catch up.

MAKING SECURITY COORDINATED, CONSISTENT AND COMPLETE

As applications and computing continue to evolve, it is up to security to keep pace. This can be easier said than done because the introduction of new technology doesn't necessarily mean that an old one goes away. For example, security may need to protect access to cloud-based applications as well to local applications or even legacy, custom-coded applications.

Adding a layer of protection in front of the authentication infrastructure can once again provide the necessary flexibility in this regard. For example, suppose we wanted to add MFA protections to a custom application or any network resource such as a local server, file share, or critical workstation. Instead of writing custom code on each app or asset, a security layer in front of the Active Directory could provide a centralised point of policy enforcement.

When a connection is attempted, the authentication proxy could trigger the MFA challenge based on the app, asset, user, or virtually any other attribute. This ensures that security policy and enforcement remains consistent for all assets without having to do custom work on each one.

This strategy extends to both private and public cloud applications and data centres as well. An authentication proxy should naturally align with an organisation's Active Directory, whether hosted locally or deployed in the cloud. Additionally, by integrating with the enterprise single sign-on (SSO) solution and federation service, security teams could extend the reach of the architecture even further to track access to any number of public cloud resources.

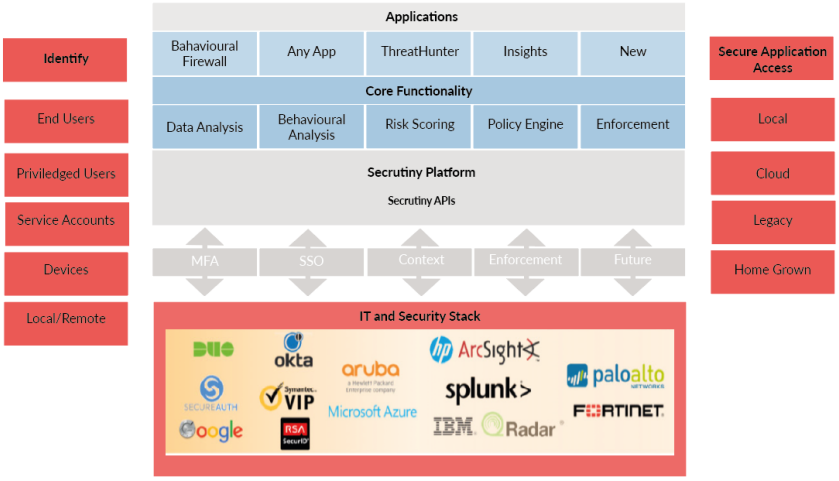

Figure 4: Bringing Consistency and Coordination Across Identity, Access and Security

This ability to integrate, share, and extend context across point solutions is key to a business-centric security architecture. We have previously described how SSO and MFA context could be used, but the same would apply to the enterprise virtual private network (VPN), Security Information and Event Management (SIEM), firewalls, and much more. This bi-directional flow ensures that the teams can have the most complete context for making decisions while retaining the control to enforce in real time. Instead of making decisions based on individual sessions or incidents, we can leverage the combined intelligence of all of an organisation's security investments, making it a true conditional access control based on Identity and threat.

CONCLUSION AND NEXT STEPS

This document lays out the foundation and some of the requirements for building a modern approach to enterprise security that can adapt to the changing needs of the business. While this is certainly not the only approach, we believe that bringing a real-time security layer based on identity, behaviour, and risk to the level of the authentication infrastructure provides a way to achieve both incredible flexibility as well as the ability to stop threats before damage is done.

This is our mission at Secrutiny. We have developed the industry's first security platform that brings a real-time layer of threat prevention to an organisation's authentication infrastructure. This allows organisations to stop cyber attacks, insider threats, and enforce policies across the extended enterprise based on identity, behaviour, and risk.

We do this by working in real-time with your Active Directory and domain controllers. From this strategic vantage point, Secrutiny can observe and continually learn the behaviours that are most essential to your enterprise, while proactively detecting the malicious actions of attackers.

Our built-in intelligence automatically surfaces weaknesses such as password issues, risky user behaviours and all their associated devices. The platform automatically identifies all privileged users, including those with privileges outside of the traditional administrator groups in Active Directory.

We continually track the actions and behaviours of all of these entities for signs of malicious or abnormal behaviour. When we see something that is suspicious, we can contextually challenge the behaviour before access is granted to verify identity and, if needed, progressively scale back the user's access based on risk and the enterprise policy. This allows organisations to automatically resolve any false positives without preventing valid access and logically aligning real-time security enforcement with the needs of the business.

If you would like to learn more about our approach, we recommend scheduling a brief demonstration where we can show the platform in action.

ABOUT SECRUTINY

Secrutiny was founded by three people – all veterans who came to realise that there is too much "snake oil" and "propaganda" in the industry. Through responding to 300+ incidents, they 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.

As the cyber industry continues to be propagated by fear, uncertainty and doubt, Secrutiny has a different approach. We focus on the things that matter, helping organisations find a lasting solution to their security fatigue, identifying the threat and asking why those threats are a risk to the business. We don't believe in adding technical stress for the sake of ticking boxes.

80% of risk reduction can be achieved with what you have.

For more information, visit https://secrutiny.com or contact us for an introductory.

☎ 0203 8238 999 ✉ enquires@secrutiny.com in LinkedIn