



THE IDENTITY UNDERGROUND REPORT



A COMPREHENSIVE ANALYSIS OF THE IDENTITY
ATTACK SURFACE'S INHERENT WEAKNESSES

YOUR DEFENSES ARE SKY HIGH...

We can see what's going on above the ground – and what we see, we can protect.

When it comes to identity protection, the user accounts and configurations we're aware of lie in full view above the ground. We can, therefore, defend them effectively against identity threats.

Unfortunately, this aboveground knowledge is painfully limited.

Beneath the known identity attack surface exists an underground world of misconfigurations, forgotten user accounts, legacy settings, malpractices, and insecure built-in features. In this report we refer to these as **Identity Threat Exposures (ITEs)**.

Attackers use these ITEs as co-conspirators to perform credential theft, privilege escalation and lateral movement. What's more, due to the common practice of syncing AD user accounts to the cloud IdP, this underground exposure could also provide attackers with direct access to your SaaS environment.

BUT UNDERGROUND YOU'RE EXPOSED.

A dramatic, low-key photograph of a flashlight beam shining through a dark, rocky tunnel. The beam is bright and focused, creating a strong contrast with the surrounding darkness. The tunnel walls are rough and textured, and the overall atmosphere is mysterious and suspenseful.

ARE YOU READY?

Let's expose this underground world into the light of day.

REPORT HIGHLIGHTS

67%

of organizations exposed their SaaS apps to compromise with insecure on-prem password sync.

37% of admins authenticate in NTLM, enabling attackers to access cleartext passwords.

Legacy NTLMv1 is used by 7% of admin users, exposing their passwords to compromise.

13% of user accounts are stale and do not perform any activity, allowing attackers to compromise them and evade detection.

109 new shadow admins are, on average, introduced by a single AD misconfiguration, enabling attackers to reset a true admin's password.

31% of all users are service accounts with high access privileges and low visibility.

12% of admin accounts are configured to have unconstrained delegation, exposing their environments to privilege escalation attacks.

7% of users regularly perform admin-level access, even though they are not included in any admin group.

THE ATTACK PATH FROM UNDERGROUND TO THE CLOUD

The vast majority of organizations today employ a hybrid identity infrastructure, with Active Directory (AD) for on-prem resources and a cloud IdP for SaaS.

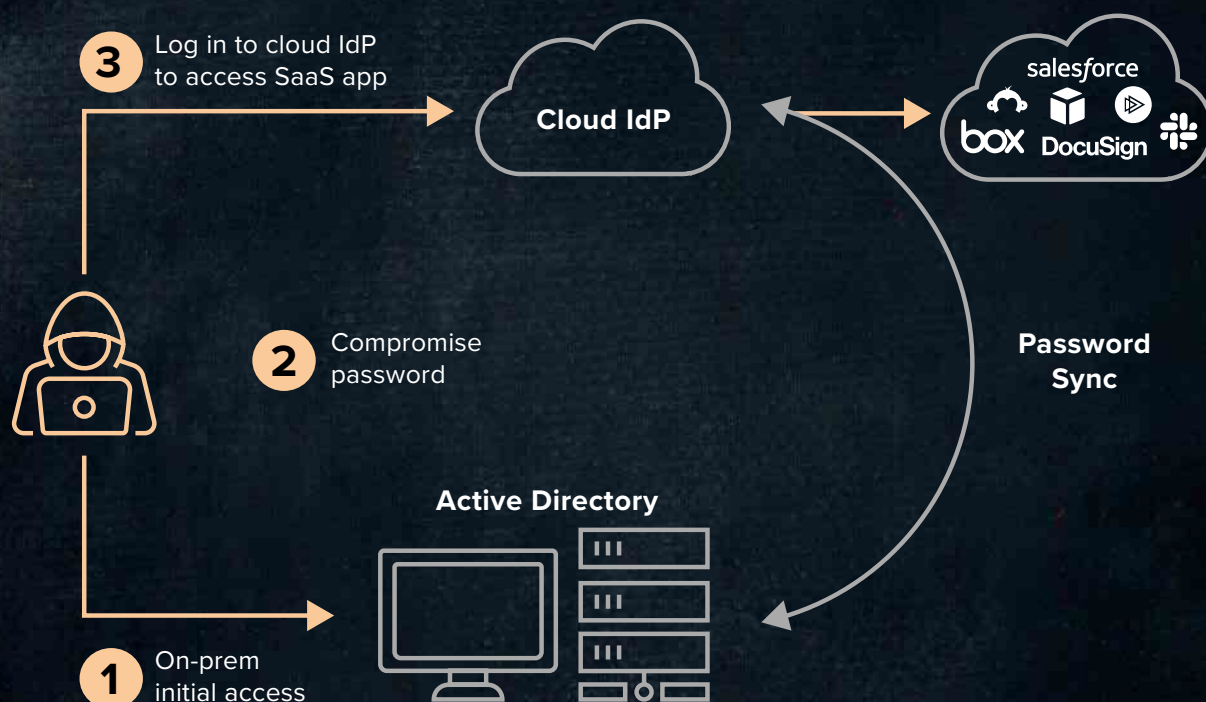
The common practice is for AD to sync users' hashes to the cloud IdP, so users can access SaaS apps with the same credentials as on-prem resources.

This significantly increases the SaaS environment's potential attack surface, as any attack that results in the adversary gaining cleartext passwords paves the way to cloud assets.

So any ITEs that enable attackers to get users' cleartext passwords provide adversaries with direct access to the SaaS environment. ITEs that expose weakly decrypted password hashes (NTLM, NTLMv1, admins with SPN) or enable attackers to reset user passwords (shadow admins) are already extensively exploited by adversaries.

67% of organizations sync the majority of their users from AD to their cloud IdP, making every password exposure in AD an access vector to the SaaS environment.

Diagram #1






TABLE OF CONTENTS

6	Executive Summary
7	Introducing Identity Threat Exposures (ITEs)
10	Password Exposers
18	Privilege Escalators
24	Lateral Movers
30	Protection Dodgers
33	Recommendations

EXECUTIVE SUMMARY

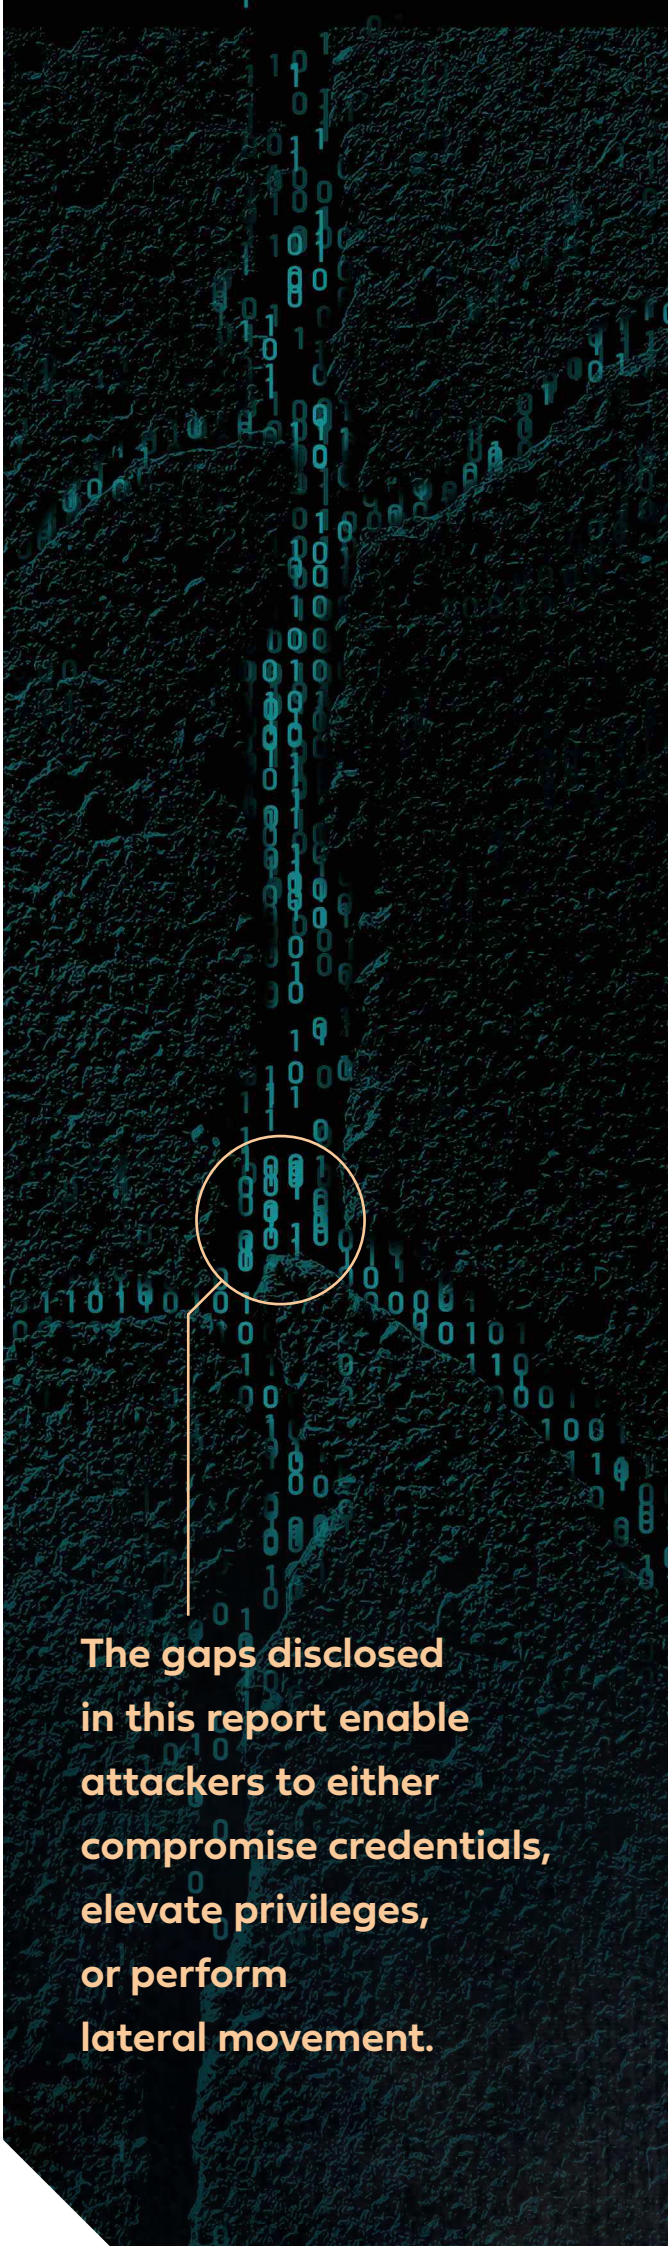
Critical parts of your identity attack surface are unknown, unprotected and underground.

This report is the first attempt to map out the most critical identity security weaknesses in the hybrid enterprise environment. These Identity Threat Exposures (ITEs), gathered from hundreds of live production environments, are the key weaknesses that allow attackers to access credentials, escalate privileges and move laterally, both on-prem and in the cloud.

They are also extremely hard to eliminate, as they stem from either misconfiguration, legacy infrastructure or built-in features.

Misconfigurations are inevitable in a high-scale production environment. Legacy infrastructure is often required for apps and systems that cannot be updated or migrated to the cloud. And built-in features are a reality that cannot be altered.

The ITEs we've gathered here are behind the steep increase in lateral movement, which is now a feature of almost every attack. For many reasons, a comprehensive insight into the resilience of the identity attack surface isn't yet part of the security team's playbook. We hope that by taking the insights from this report, these teams can now throw a spotlight on critical weaknesses and take action against them.



The gaps disclosed in this report enable attackers to either compromise credentials, elevate privileges, or perform lateral movement.

INTRODUCING IDENTITY THREAT EXPOSURES (ITEs)

What are Identity Threat Exposures (ITEs)?

ITEs are security weaknesses that expose an environment to identity threats: credential theft, privilege escalation, or lateral movement. An ITE can result from a misconfiguration, malpractice, legacy identity infrastructure, or even built-in features.

ITEs in this report: prevalent, impactful, and available for attackers to exploit.

While there are multiple ITEs of different types, we've only included those that introduce a risk every organization is likely to experience. For that, we've applied the following inclusion criteria:

- **Prevalence:** The ITE appeared in a large portion of the tested environments, making it a common phenomenon.
- **Availability:** The ITE is easily discoverable by adversaries that have gained initial access to the targeted environment.
- **Impact:** The ITE makes it significantly easier for adversaries to gain access to systems and resources, escalate their privileges or move laterally.

We describe each ITE with the following attributes:

- **MITRE mapping:** the MITRE technique used against the ITE.
- **Type:** the ITE's root cause – either misconfiguration, legacy infrastructure, or built-in feature.
- **Compromise impact:** the potential ramifications of an attacker actively abusing this ITE.
- **Availability:** how an attacker can discover the existence of this ITE and abuse it.
- **Visibility and protection:** how existing security controls impact this ITE.

WHAT ITE TYPES ARE THERE?

We classify ITEs into four groups, based on what attackers can achieve by using them:

Password Exposers:

ITEs that allow adversaries to access a user account's cleartext password.

Lateral Movers:

ITEs that enable adversaries to use compromised accounts to perform undetected lateral movement.

Privilege Escalators:

ITEs that enable adversaries to escalate any access privileges they already possess.

Protection Dodgers:

ITEs that make security controls less effective at monitoring and protecting user accounts.

Table #1 summarizes the ITEs in this report:

Category	Related MITRE ATT&CK	Examples
Password Exposers	Credential access	NTLM authentication NTLMv1 authentication Admins with SPN
Privilege Escalators	Privilege escalation	Shadow admins Unconstrained delegation
Lateral Movers	Lateral movement	Service accounts Prolific users
Protection Dodgers	There isn't an exact MITRE ATT&CK technique that maps to this category. It allows attackers to go undetected for long periods of time.	New user accounts Shared accounts Stale users



PASSWORD EXPOSERS

Password Exposer ITEs enable attackers to discover a user account's cleartext password. Members of this group expose the password hash to common compromise techniques, so attackers can crack it offline and use it for future attacks on both on-prem and SaaS resources.

PASSWORD EXPOSERS IN THIS REPORT:



**NTLM
authentication**



**NTLMv1
authentication**



**Admins
with SPN**

NTLM AUTHENTICATION

Related MITRE technique:

ID TA0006 Credential Access

Type: Legacy Infrastructure

NTLM authentication exists in any Windows domain. While Kerberos is often used as the default authentication protocol, NTLM is still widespread.

Compromise Impact: Exposure of a user account's cleartext password

Once attackers decrypt a user account's hash, they obtain its cleartext password, giving them a direct route into every resource it has access to, both on-prem and in the cloud.

Availability: High

There are many tools to obtain an NTLM hash from either a machine's memory (Mimikatz, ProcDump, etc.) or network traffic (Responder and other MITM TTPs).

Visibility and Protection: Low

As the fallback when Kerberos authentication fails, NTLM is a core part of authentication infrastructure and cannot be easily eliminated from an environment.

CLOUD ALERT



The exposed passwords from NTLM hashes can also be used to access the SaaS environment.



**The use of NTLM
exposes users' cleartext
passwords, yet it is
actively used by 64%
of user accounts.**


NTLM AUTHENTICATION

The following stats show the average usage of NTLM authentication in AD environments today:




NTLM/Kerberos breakdown by account usage

64% of user accounts regularly authenticate with NTLM



NTLM/Kerberos breakdown by authentication traffic

37% of admins and **46%** of non-admin users regularly authenticate with NTLM



NTLM AUTHENTICATION

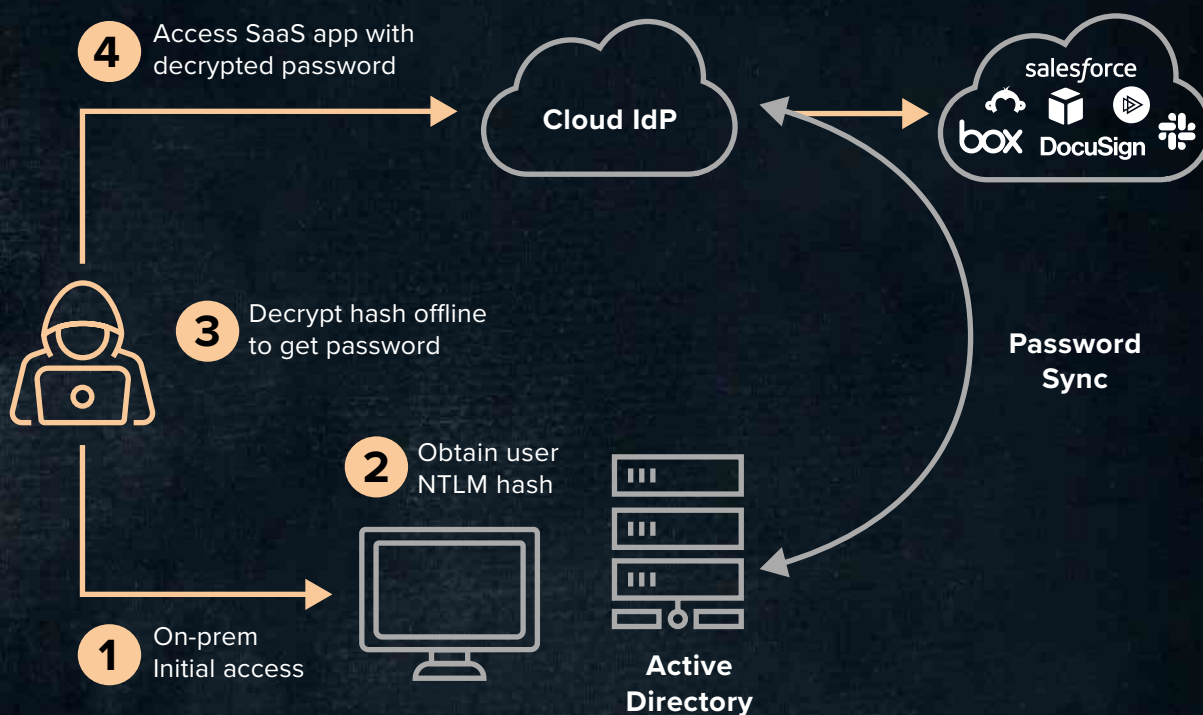
How does NTLM expose the SaaS environment to malicious access?

Password exposure through the use of NTLM introduces a potential risk to the SaaS environment due to the high rates of synced users.

Diagram #2 shows the flow of such an attack:

1. Attacker gains initial access to a machine within the targeted environment.
2. Once inside, the attacker captures a user's NTLMv2 hash, either from the machine itself or from the network traffic.
3. Attacker decrypts the hash offline to obtain the cleartext password.
4. With the newly obtained password, the attacker opens a browser tab and connects directly to the compromised user's SaaS environment.

Diagram #2



NTLMv1 AUTHENTICATION

Related MITRE technique:

ID TA0006 Credential Access

Type: Legacy infrastructure

NTLMv1 is the first version of the NTLMv2 protocol that is still in common use, primarily (though not exclusively) in environments that rely on legacy applications.

Compromise Impact: Exposure of a user's cleartext password

Residence of NTLMv1 in an environment exposes users' cleartext passwords to adversaries, because it implements a relatively low complexity encryption algorithm on the user hash. As a result, adversaries that intercept the encrypted hash can easily decrypt it offline and obtain the user's cleartext password.

Availability: High

NTLMv1 can be easily detected via traffic interception or a direct registry query. Once detected, any form of forced authentication reveals the user's hash for offline cracking.

Visibility and Protection: Low

Most organizations don't have visibility into the presence of NTLMv1 in their environments. While querying AD can reveal how many machines support NTLMv1, there is no easy way to gain insight into its actual usage.

CLOUD ALERT



The exposed passwords from NTLMv1 hashes can also be used to access the SaaS environment.


Organizations are unaware that NTLMv1 is still commonly used by admin accounts in their environments.



NTLMv1 AUTHENTICATION

Average percentage of NTLMv1 authentications per organization size

	Small	Medium	Large
Client machines:	3.4%	5.2%	6.7%
Server machines:	1.8%	2.3%	1.5%
Admin accounts:	4.4%	6.5%	6.9%



NTLMv1 AUTHENTICATION

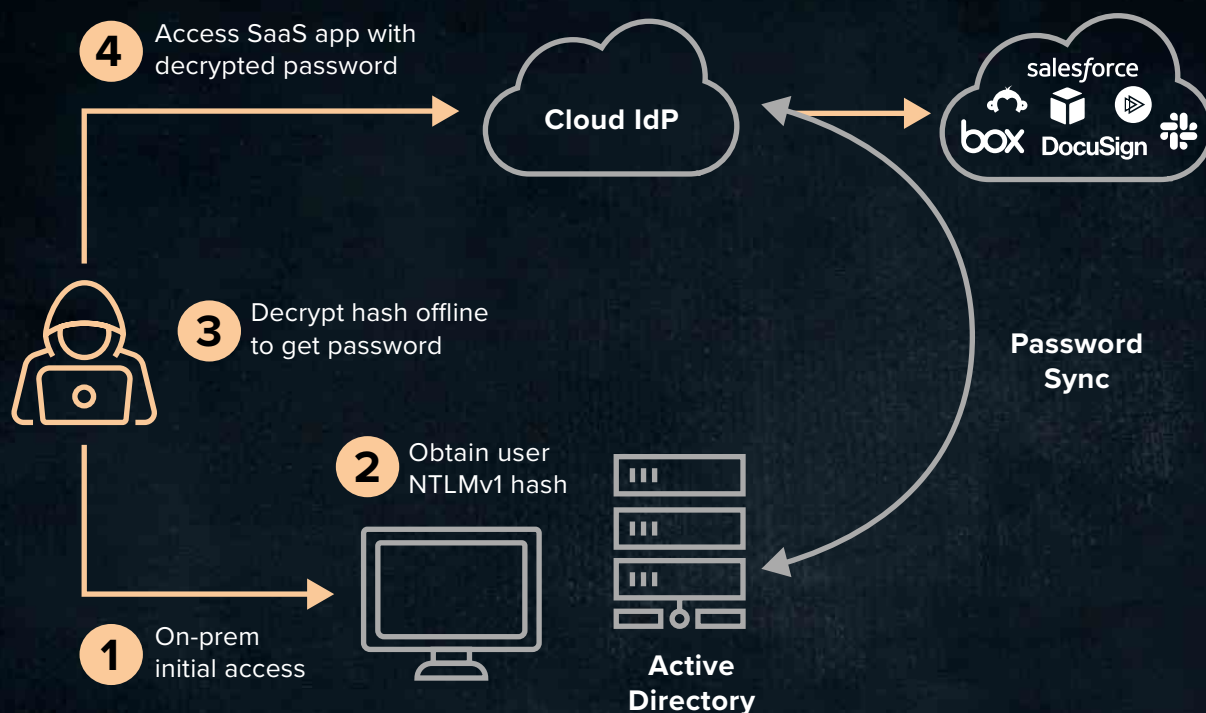
How does NTLMv1 expose the SaaS environment to malicious access?

Password exposure through the use of NTLMv1 introduces a potential risk to the SaaS environment due to the high rates of synced users.

Diagram #3 shows the flow of such an attack:

1. Attacker gains initial access to a machine within the targeted environment.
2. Once inside, the attacker captures a user's NTLMv1 hash, either from the machine itself or from the network traffic.
3. Attacker decrypts the hash offline, obtaining the cleartext password.
4. With the newly obtained password, the attacker opens a browser tab and connects directly to the compromised user's SaaS environment.

Diagram #3



ADMINS WITH SPN

Related MITRE technique

ID TA0006 Credential Access
ID: T1558.003 Kerberoasting

Type: Built-In feature

Service Principal Name (SPN) is the unique identifier of a service instance. Attackers can identify these accounts and request a service ticket, which is encrypted with the service account's hash. This can be taken offline and cracked.

Compromise Impact: Exposure of the service account's cleartext password

Once attackers decrypt the service account's hash, they obtain its password and gain a route into every resource this service account has access to.

Availability: High

There are various open-source tools attackers can use to identify accounts with an SPN and for obtaining the ticket and cracking it offline.

Visibility and Protection: Low

As a built-in identity infrastructure feature, SPN cannot be removed from accounts that require it. Moreover, there is no easy way to discern between a legitimate request for a service ticket and a malicious one.

CLOUD ALERT




The exposed passwords from SPN hashes can also be used to access the SaaS environment.

Admin accounts with an SPN provide adversaries with the means to crack their passwords offline.

ADMINS WITH SPN

Average percentage of admins with SPN per organization size



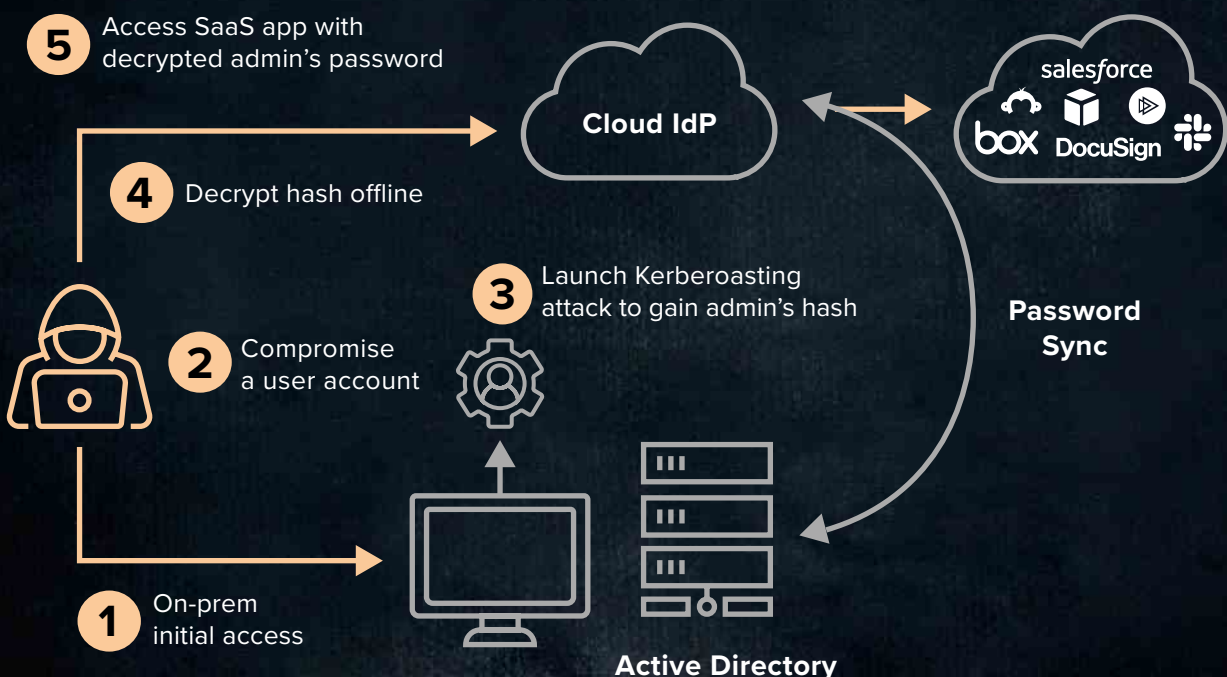
Leveraging this ITE for malicious SaaS access

Password exposure through admins with SPN introduces a potential risk to the SaaS environment due to the high rates of synced users.

Diagram #4 shows the flow of such an attack:

1. Attacker gains initial access to a machine within the target environment.
2. Once inside, the attacker compromises a domain user account.
3. Attacker identifies accounts within the domain with associated SPN and launches a Kerberoasting attack.
4. Attacker gets the admin's ticket, dispatches from the environment, extracts the admin's NTLM hash from the ticket, and decrypts it to get the admin's password.
5. The attacker opens a browser tab and connects directly to the compromised admin's SaaS environment.

Diagram #4





PRIVILEGE ESCALATORS

Privilege Escalator ITEs enable attackers to gain additional access privileges on top of what they have already acquired. This is caused by misconfiguration in user creation or by using insecure legacy settings.

PRIVILEGE ESCALATORS IN THIS REPORT:



**Shadow
admins**



**Unconstrained
delegation**

SHADOW ADMINS

Related MITRE Technique:

TA00004 Privilege Escalation

Type: Misconfiguration

Shadow admins are user accounts that have been inadvertently assigned full or partial admin privileges, or configuration/reset privileges over admin accounts.

Compromise Impact: Escalating an attacker's access privileges

Compromising a shadow admin enables an attacker to control an account that has high access and configuration privileges, paving the way to further access and compromise of additional resources.

Availability: High

Detecting an AD hierarchy misconfiguration that could lead to an account becoming a shadow admin is easy and can be done with a simple script.

Visibility and Protection: Low

Identity and security teams have low visibility into the number and privileges of shadow admins in their environments. As a result, these accounts are not subject to the monitoring and protection measures typically enforced on admin accounts. This makes undiscovered shadow admins a critical risk.

CLOUD ALERT




Once adversaries use a shadow admin to reset the password of a true admin, they can use the new password to access the SaaS environment.

3% of non-admin users are shadow admins.

SHADOW ADMINS


Average percentage of shadow admins out of non-admin users per organization size



How are shadow admins created?

There are various AD misconfigurations that can assign administrative privileges to a non-admin.

This diagram shows the leading ones.




No one is safe against a shadow burst



30% of organizations experienced a burst of at least 70 new shadow admins within 30 days

What is a shadow burst?

A shadow admin burst occurs when a single misconfiguration triggers a sudden leap in the number of shadow admins. The average number of new shadow admins created in a burst by a misconfiguration is 70, while the largest number we've observed was 1000.



New shadow admins per single misconfiguration.

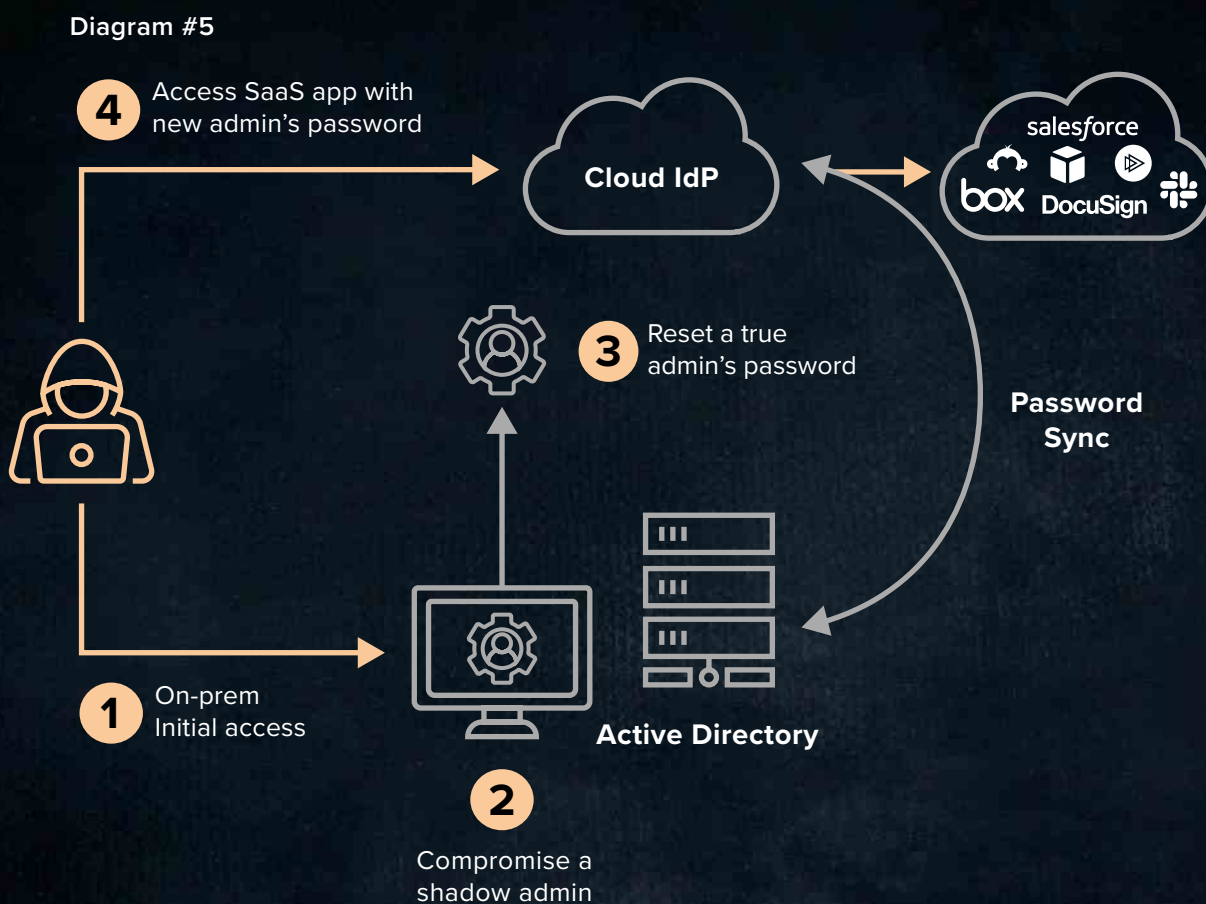
SHADOW ADMINS

How do shadow admins expose the SaaS environment to malicious access?

Aside from escalating privileges in the AD managed environment, the use of shadow admins introduces a potential risk to the SaaS environment due to the high rates of synced users.

Diagram #5 shows the flow of such an attack:

1. Attacker gains initial access to a machine within the targeted environment.
2. Once inside, the attacker compromises a user account and discovers it's a shadow admin.
3. Attacker uses the shadow admin to reset the password of a true admin.
4. Attacker uses the new password to access the admin's SaaS environment.





UNCONSTRAINED DELEGATION

Related MITRE Technique

TA00004 Privilege Escalation

Type: Legacy Configuration

Unconstrained delegation is the insecure legacy version of delegation which was later followed by constrained and resource-constrained delegation. It allows a compromised account to access all the same resources as the delegating account. This capability is mostly required for machine accounts that access other machines on behalf of a user; for example, when an app server accesses a database to fetch data for an app user.

Compromise Impact: Attacker gains the privileges of the delegating account

When an admin account logs in to a machine that has unconstrained delegation, its TGT remains stored in the machine's memory. This allows the attacker to establish a new session with the privileges of the user account's TGT.

Availability: High

Attackers can access the machine names that support unconstrained delegation with a straightforward PowerShell/cmd command.

Visibility and Protection: Medium


Identity teams can easily discover which accounts support unconstrained delegation. Disabling it eliminates the risk altogether; however, this is not always possible due to operational concerns. In this case, discerning between a legitimate delegation and a malicious one can prove incredibly difficult.

Attackers can easily gain admin access privileges by abusing unconstrained delegation.



UNCONSTRAINED DELEGATION

Average percentage of admin accounts with unconstrained delegation per organization size





LATERAL MOVERS

Lateral Mover ITEs enable attackers to perform lateral movement without being detected by existing security measures. User accounts within this category are often excluded from the visibility and protection of the security solutions in place.

LATERAL MOVERS IN THIS REPORT:



Service accounts



Prolific users

SERVICE ACCOUNTS

Related MITRE technique

TA0008 Lateral Movement

Type: Built-In Feature

Service accounts are user accounts created for machine-to-machine communication. They are created either within the installation of on-prem software or manually by admins to automate repetitive tasks. Service accounts typically have high privilege access to perform their machine-to-machine tasks, effectively making them admin accounts.

Compromise Impact: Access privileges to multiple resources

The compromise of a service account could give attackers access to multiple resources, making it an ideal tool for lateral movement and mass propagation within the target environment.

Availability: High

Service accounts can be easily detected by their naming conventions and SPN association. Default naming conventions is a specific issue with the service accounts created during software installation.

Visibility and Protection: Low

A recent [whitepaper](#) by Osterman Research revealed that only 4% of organizations have full visibility into their service accounts. Service accounts cannot be protected with MFA, and the lack of visibility into their activities eliminates the possibility of protecting them in a PAM vault with password rotation.

Service accounts are the default targets for lateral movement due to their high access privileges, low visibility, and protection challenges.




SERVICE ACCOUNTS

Average percentage of service accounts per organization size



Service accounts NTLM/ Kerberos usage breakdown

46% of service accounts regularly authenticate with NTLM



46% of service accounts regularly authenticate with NTLM

56%

of organizations unknowingly sync more than half of their service accounts to their SaaS directory.

SERVICE ACCOUNTS

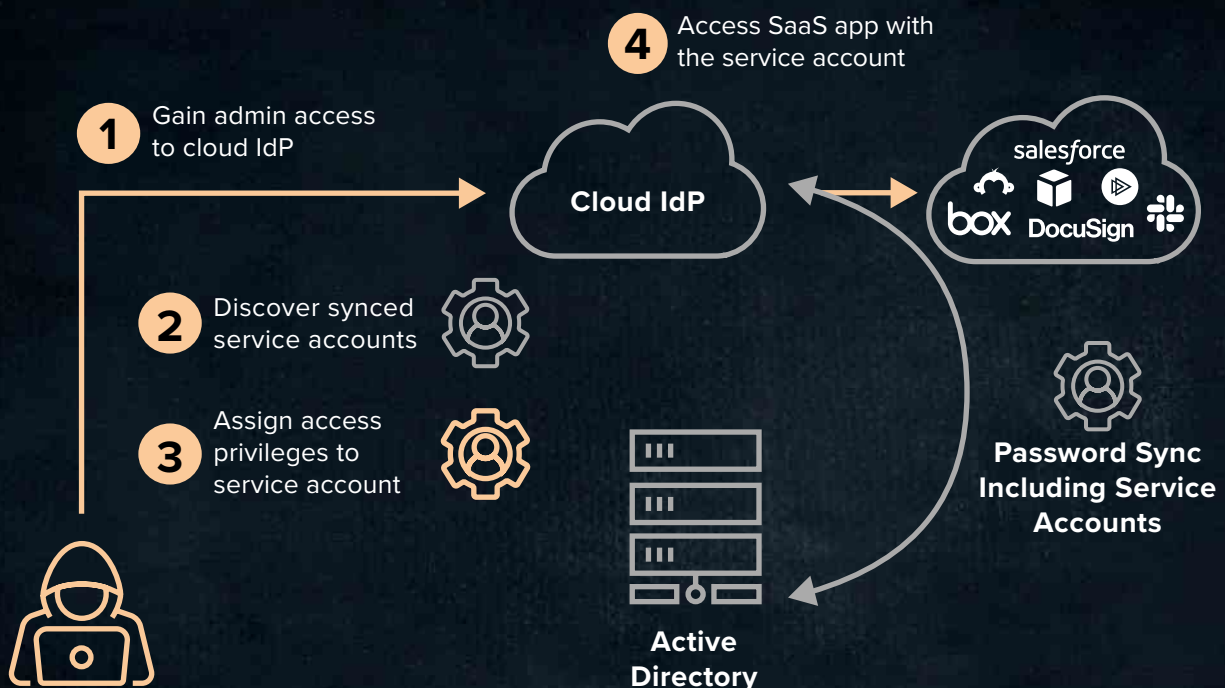
How do service accounts expose the SaaS environment to malicious access?

Service accounts are typically synced from AD to the cloud IdP. This increases the SaaS attack surface, since it creates a multitude of dormant accounts in the cloud IdP. These accounts can't be used to access SaaS resources by default, since they are unknown to the SaaS management team. However, an attacker that has gained admin access privileges to the cloud IdP can activate them and assign them access privileges.

Diagram #6 shows the flow of such an attack:

1. Attacker gains admin access privileges to the cloud IdP management console.
2. Once inside, the attacker searches for synced service accounts (naming conventions are a useful guide) until finding one.
3. Attacker configures an access policy for the chosen service account and assigns it access privileges to SaaS apps.
4. Attacker uses the service account to access and act within the SaaS environment.

Diagram #6



PROLIFIC USERS

Related MITRE Technique

TA0008 Lateral Movement

Type: Misconfiguration

Prolific users are standard user accounts, as defined by all AD parameters, that have access privileges to an exceedingly high number of machines.

Compromise Impact: Access privileges to multiple machines

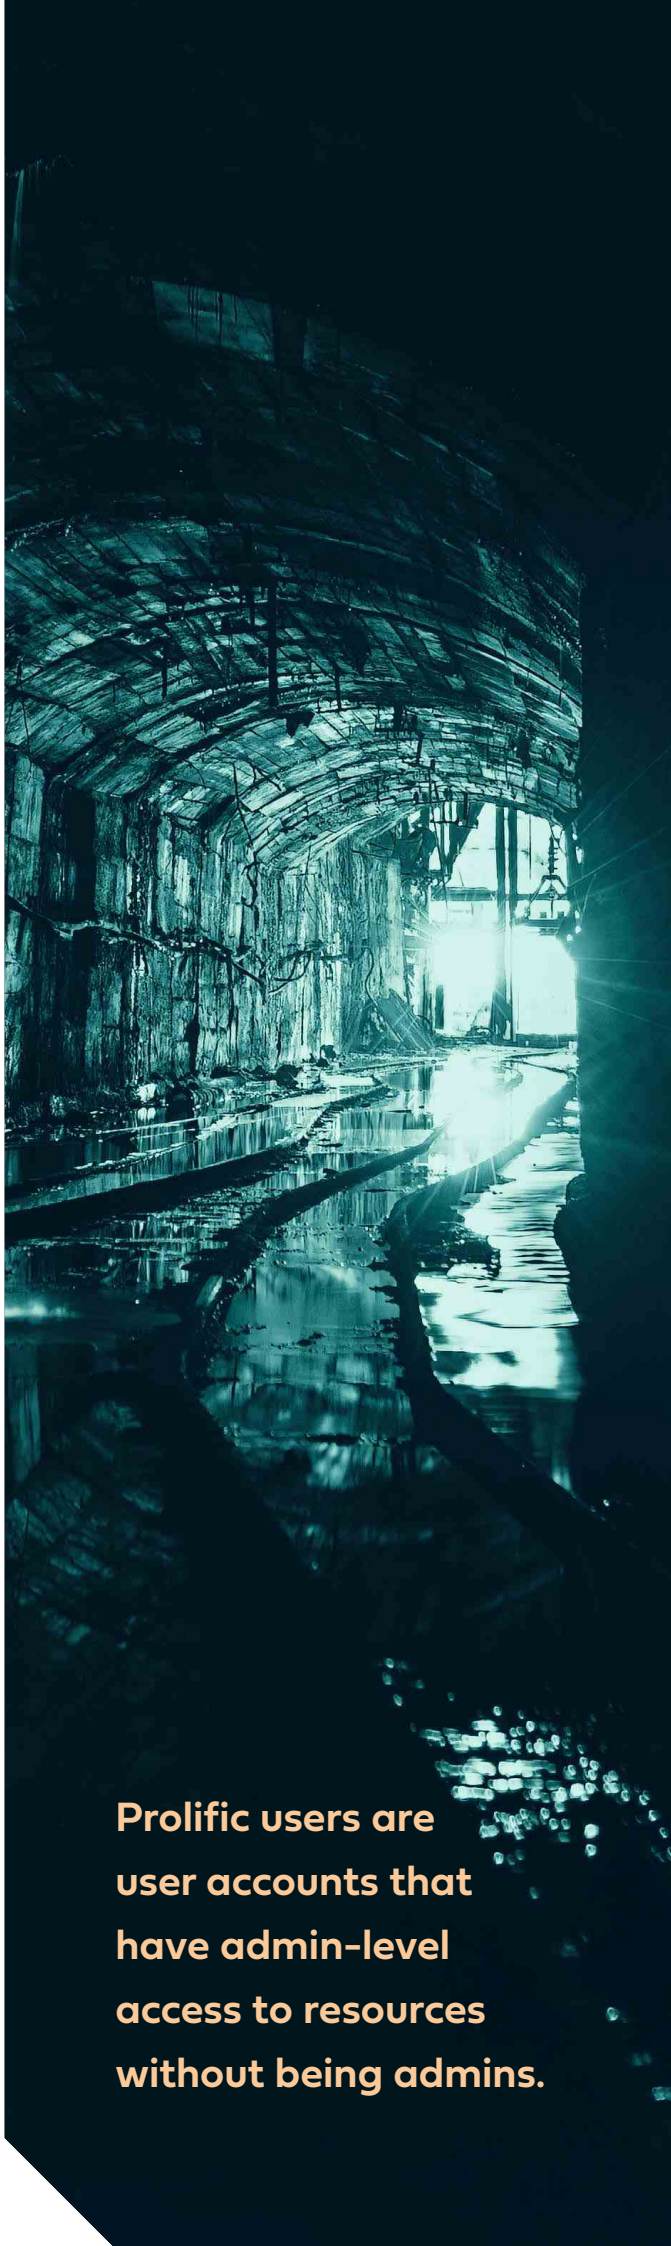
Once compromised, attackers gain a direct route into the same resources as these prolific user accounts, facilitating a rapid and efficient lateral movement process.

Availability: Medium

There is no straightforward way to know in advance if a user account is prolific or not. However, given their relatively large number, attackers stand a good chance of finding one simply by trying to use a standard compromised account to move laterally.

Visibility and Protection: Low


Prolific users are not subject to the same monitoring and protection measures placed over admin users. Technically, they are not even admins, since they are not included in any administrative user group. This makes them a highly lucrative target for compromise, as they yield a similar result as the compromise of an admin account and are less likely to be protected.



Prolific users are user accounts that have admin-level access to resources without being admins.

PROLIFIC USERS

Average percentage of prolific user accounts per organization size



Clearly, the larger the environment, the higher the risk of prolific users. This makes sense, since they are tightly related to misconfigurations or poor judgment in associating users to the groups they should be in.

How should we interpret these results?

Users that fall into this category are either:

- **Regular domain users** that are actively accessing an unrestricted number of resources.
- **Actual admins** that are accessing resources and performing tasks without using an account from one of the admin groups.

That way or the other, the risk created by a prolific user is clear. Compromise of this type of account offers attackers access to multiple resources. Moreover, since accessing many resources is standard behavior for a prolific user, it won't be flagged as an anomaly by the SIEM, UEBA, or any other behavioral analysis tools.



PROTECTION DODGERS

Protection Dodger ITEs potentially downgrade the protection of a user account. This group differs from the others because they are an inherent part of the user management routine and not, strictly speaking, security gaps. We chose to include them since they introduce key protection challenges.

PROTECTION DODGERS IN THIS REPORT:



New users



Shared accounts




Stale users

NEW USER ACCOUNTS

Newly created accounts are not immediately scoped in an environment's security policies and practices, nor subject to the organization's security measures.


Average percentage of new accounts created per week per organization size



SHARED ACCOUNTS

These are user accounts associated with more than one human user, which prevents them from being protected by MFA. There is an industry consensus that the practice of sharing a user account between multiple users is highly insecure, yet it is still widespread.


Average percentage of shared accounts per organization size



STALE USERS

These are user accounts that are no longer active, yet are still valid and can be used to access resources. These users might be unknown to the security team and are not subject to MFA, monitoring, or any other security measures.

Average percentage of stale users per organization size

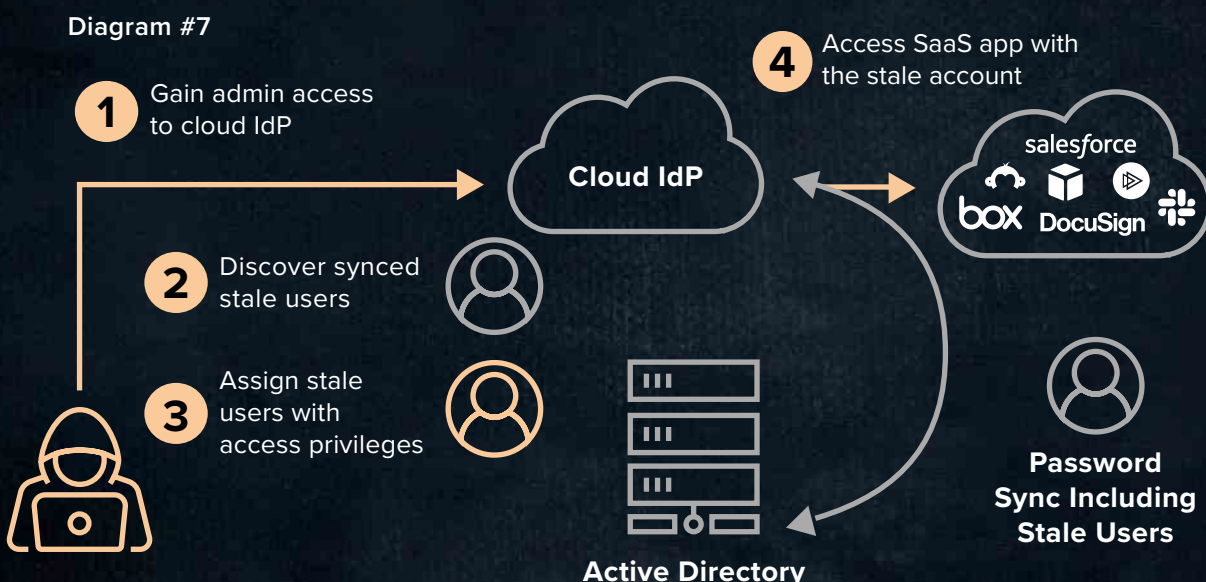


How do stale users expose the SaaS environment to malicious access?

Stale accounts are typically synced from AD to the cloud IdP. These accounts can't be used to access SaaS resources by default, since they are unknown to the SaaS management team. However, an attacker that has gained admin access privileges to the cloud IdP can activate them and assign them access privileges.

Diagram #7 shows the flow of such an attack:

1. Attacker gains admin access privileges to the cloud IdP management console.
2. Once inside, the attacker searches for accounts that are not included in any access policy. The large amount of stale accounts ensures high success rates for this search.
3. Attacker configures an access policy for the chosen stale account, assigning it access privileges to SaaS apps.
4. Attacker uses the account to access and act within the SaaS environment.



RECOMMENDATIONS

1 Know where you're exposed

Make sure you have visibility into the ITEs in your environment. If you're syncing AD users to your cloud IdP, ensure it follows Microsoft's best practices and does not create a mass of idle users.

2 Eliminate risk where you can

Work closely with the identity team to weed out the ITEs that result from malpractices or misconfigurations and establish a process to address them as soon as – or before – they appear.

3 Contain and monitor existing risks

For ITEs that cannot be eliminated, such as service accounts or the use of NTLM, ensure the SecOps team has a process in place to monitor these accounts closely for any sign of compromise.

4 Take preventative measures

Apply identity segmentation rules or MFA policies to prevent user accounts from falling victim to featured ITEs where possible. Enforce access policies on your service accounts that would block them from accessing any destination beyond their pre-designated resources.

5 Connect the identity and security teams

The responsibility for identity protection is distributed between the identity and the security teams, where the latter's knowledge enables them to prioritize which ITEs to resolve, while the former can put these fixes into effect.



ABOUT SILVERFORT



Silverfort is the first Unified Identity Protection platform that provides all the required capabilities to prevent and detect identity threats in a single solution. With its native integration with all leading IAM solutions, Silverfort delivers comprehensive MFA, service account protection, identity threat detection and response (ITDR), identity segmentation, and identity security posture management (ISPM) across the hybrid environment. Silverfort extends modern identity protection to any user accessing any resource, including those that could have never been protected before, such as legacy apps, command-line access, service accounts, and many more. With these capabilities, Silverfort enables organizations to gain real-time protection against the use of compromised credentials, on-prem and in the cloud.

For more information, visit silverfort.com