



Securing Active Directory and Azure AD

The issues and paths to transformation

White Paper

Introduction

Since its launch more than 20 years ago, the Active Directory service has become a market standard that is present in almost all the information systems of organizations. In recent years, two significant trends have brought it back into the spotlight.

The first is the result of the high exposure of this component to cyber threats. Since the Active Directory is the cornerstone of the information system in terms of rights and privileged accounts, it is a top-priority target for hackers, who are looking to break into the information system and gain broad access. In this way, they can use it to deploy malware or to access, and then divulge, information. In recent years, many organizations have launched major remediation projects to face up to this threat.

The second trend is due to the rise in the use of collaborative services, which has increased sharply with the spike in home working. To enable all the new usages in the modern workplace, user management has extended its scope of action to include the cloud, thanks to Azure AD. Most cases do not consist of a switch from full on-premises to full cloud, but more of an extension of the existing configuration in the form of hybrid architectures. This shift must take into account the security considerations to avoid exposing the organization.

Microsoft and Wavestone teamed up to analyze the trends observed in the field, to list the necessary considerations and to provide a few keys and best practices to be adopted when making structural changes.



Contents

1

What level of maturity was observed in the field?	4
Which architectures and what level of cyber security maturity?	5
Feedback on the attacks observed by the CERT-W	11

2

How can the situation be improved?	19
Enterprise Access Model	20
Securing tier 0 and implementing the control plane	25
Securing Azure AD subscriptions	33
Migrating from Active Directory to Azure Active Directory	43
Improving the security situation	52

3

How to combat a cyber attack	54
Understanding the difficulties of rebuilding Active Directory in order to better anticipate a crisis	55
Simply rebuilding AD is not enough	59

Conclusion	60
------------	----

What level of maturity was observed in the field?

This chapter aims to present the observed current level of security for Active Directory and Azure AD.

To begin with, we will look at the main types of architectures we observed and the security issues, and then we will share the lessons learned from the attacks encountered by the Wavestone CERT (CERT-W).

Which architectures and what level of cyber security maturity?

There are three main types of architecture:

On-premises AD architecture

Traditional architecture is on the decline amongst customers (<10%). For these customers, sovereignty is usually an important factor.

Hybrid AD/Azure AD architecture

A majority of customers (80% to 95%) have adopted this architecture, which has been driven by the rise of Office 365. But there are variations in the implementations, especially in the synchronization or not of password hashing.

Active Directory is organized around domains, grouped into one or more forests. In large organizations, it is quite common to have more than 100 domains or forests.

This complex architecture can be explained by the weight of history and the multitude of changes that the enterprise has gone through, such as mergers, acquisitions, restructuring, etc.

Since Active Directory is not perceived as an application that “adds value to the business”, new scopes are integrated with minimal changes (the cheapest and fastest way) and without giving any global thought to the optimization of the architecture. Trust relationships between domains or forests are added, so that users can be recognized everywhere in the IS.

A low level of security

In most organizations, keeping AD in a secure state often consists simply of installing security patches and dealing with the obsolescence of the OS.

“Only 25% of authentications still take place on-premises”

100% Azure AD architecture

This architecture is only found in new entities that are created through a 100% digital prism (<5%). It requires an information system that meets a number of prerequisites.

On- premises Active Directory, or the weight of history

“In 80% of the audits, the IS was breached in less than 24 hours”

Wavestone, 2020 AD audits

Functional upgrades are rarely implemented, usually due to the ignorance or the fear of the possible impacts of extending the system. Consequently, organizations do not benefit from the new functionality that limits security risks (for example, the implementation of protected users groups, authentication silos and Kerberos armoring). Similarly, too few organizations deactivate obsolete protocols.

It is quite common for organizations to have hundreds of domain or enterprise administrator accounts, while best practices and the application of the principle of least privilege require these accounts to be limited to fewer than five. This situation can be explained by the difficulty of managing change (the withdrawal of rights can be perceived as a form of demotion or dispossession). But also by the demand for service accounts with excessive rights that make it easier to integrate new applications.

Changes to the architecture, and in particular the establishment of trust relationships, are defined exclusively from a functional perspective, without assessing the risks of the propagation of an attack between the domains and the forests. In the course of its audits, Wavestone regularly comes across abandoned domains with a two-way trust relationship. And this is what determines the overall level of security!

Likewise, the implementation of Azure AD is intended to meet functional needs, without any consideration of security. Very few organizations have defined a management process for privileged accounts that is well adapted to the specifics of Azure AD. In another example, only 30%(*) of global administrator accounts have activated multi-factor authentication, while this functionality is native and free.

(*) Microsoft, August 2021

The recent rise in the awareness of security issues

Against the backdrop of the current explosion of attacks that exploit the AD configuration errors, it has become common for executive committees to question the CIO or the Information System Security Officer about the level of security of AD, and to approve investments of hundreds of thousands, or even millions, of euros in rework and security projects.

“53% of major corporations are running projects to make AD secure”

CESIN, 2021 Barometer

A large majority have launched projects to make AD secure, but the audits conducted by Wavestone show us that:

“Less than 10% of customers correctly implemented good security practices”

Wavestone, 2020 AD audit

Even if administration practices may have evolved, there are often configuration faults that allow access to tier 0 (this concept is described in detail in Chapter 2). For example, means of compromising security and increasing privileges can be concealed in dangerous access rights (ACL) configured on tier 0 objects, or result from the misuse of accounts with high privileges. The AD security model is described in Chapter 2.

A 100% Azure AD architecture: the target for all?

Today, no large organizations are in a position to completely replace their on-premises AD with a service that is supported 100% by Azure AD. Few organizations are capable of having an IS that meets all the technological prerequisites of this transformation (workstations managed by Mobile Device Management (MDM) and Azure AD, absence of applications using

NT LAN Manager (NTLM), Kerberos or Lightweight Directory Access Protocol (LDAP) authentication, creation of users in the cloud, etc.). Using a managed AD service could be an option to guarantee backward compatibility, without having to manage tier 0.

Using the cloud could be perceived as delegating risk management to Microsoft, but this is only partial. Customers remain responsible for the configuration of the platform, identities and data. Unfortunately, awareness of this fact remains low. Any new subscription to the Teams solution includes the creation of an Azure AD subscription.

32/100

The average Secure Score

34.6 is the highest score in the technology sector

24 is the score when creating an Office 365 E3 subscription

To control their security, it is important for organizations to take possession of the tools on offer that did not exist on-premises. However, these tools may require advanced licenses. The Secure Score (detailed in a Focus later on in this document), which is the target that organizations aim for, should be between at least 60 and 70.

Azure AD: a *cloud* directory

Azure Active Directory (Azure AD), released in November 2011, is an Identity as a Service (IDaaS) solution.

Azure AD provides organizations with the functionality required to manage the authentication of modern applications (SAML and WE-Federation, OAuth2, OpenID Connect and FIDO2).

It is not simply Active Directory in the cloud, but an entirely new solution.

Azure AD was designed on a cloud architecture, based on micro-services and distributed across several geographic zones.

An Azure AD tenant is automatically created when an organization subscribes to a Microsoft cloud service, such as Azure or Office 365.

Graph API

Azure AD proposes an Application Programming Interface (API), called Graph API, to interrogate and update the objects in the directory.

Graph API is a gateway that unifies several other REST APIs, including those of Exchange Online, OneDrive, Endpoint Manager or Security Graph.

A miracle cure?

Even if most common attacks today target AD, migration to Azure AD is not a miracle cure. Azure AD simplifies the implementation of passwordless strong authentication or risk-based conditional access control, but privileged accounts must still be tightly controlled.

It is imperative to implement a security project to control this new building block and to benefit from the transformation in order to adopt state of the art administration practices.

In particular, it is advised to:

- / Audit the Azure AD configuration, regularly validate the members of the privileged roles and the applications authorized to interact with Azure AD;
- / Develop specific detection scenarios to limit the time for which intruders remain invisible in the IS.

345 million

Azure AD manages more than 345 million active users every month, with an average of 30 billion authentication requests per day.

Is NTLM still the Achilles heel of security?

Applications use NT LAN Manager (NTLM) to authenticate users and, possibly, to make sessions secure, when requested by the application.

The NTLM protocols are obsolete authentication protocols that use a challenge-response method to enable clients to mathematically prove that they possess the NT password hash. Current and past versions of Windows support several versions of this protocol, including NTLMv2, NTLM and the LM protocol.

NTLM

NTLMv2

NTLMv1

LM

Kerberos has been the default authentication protocol since Windows 2000. However, if, for any reason, the Kerberos protocol is not negotiated, then the applications linked to Active Directory will attempt to use one of the NTLM protocols, if available.

All the versions of NTLM are vulnerable to widely documented attacks. This is the reason why the NTLM protocol is not supported in Azure Active Directory, and can be deactivated in Azure AD DS and in Active Directory.

Beyond weak cryptographic support, the absence of any server authentication can allow hackers to steal the identity of a server. Consequently,

applications that use NTLM may be vulnerable to reflection attacks, in which hackers steal the authentication exchange between a user and a legitimate server and use it to authenticate themselves on another computer, or even on the user's computer.

The NTLM protocol is not supported in Azure Active Directory.

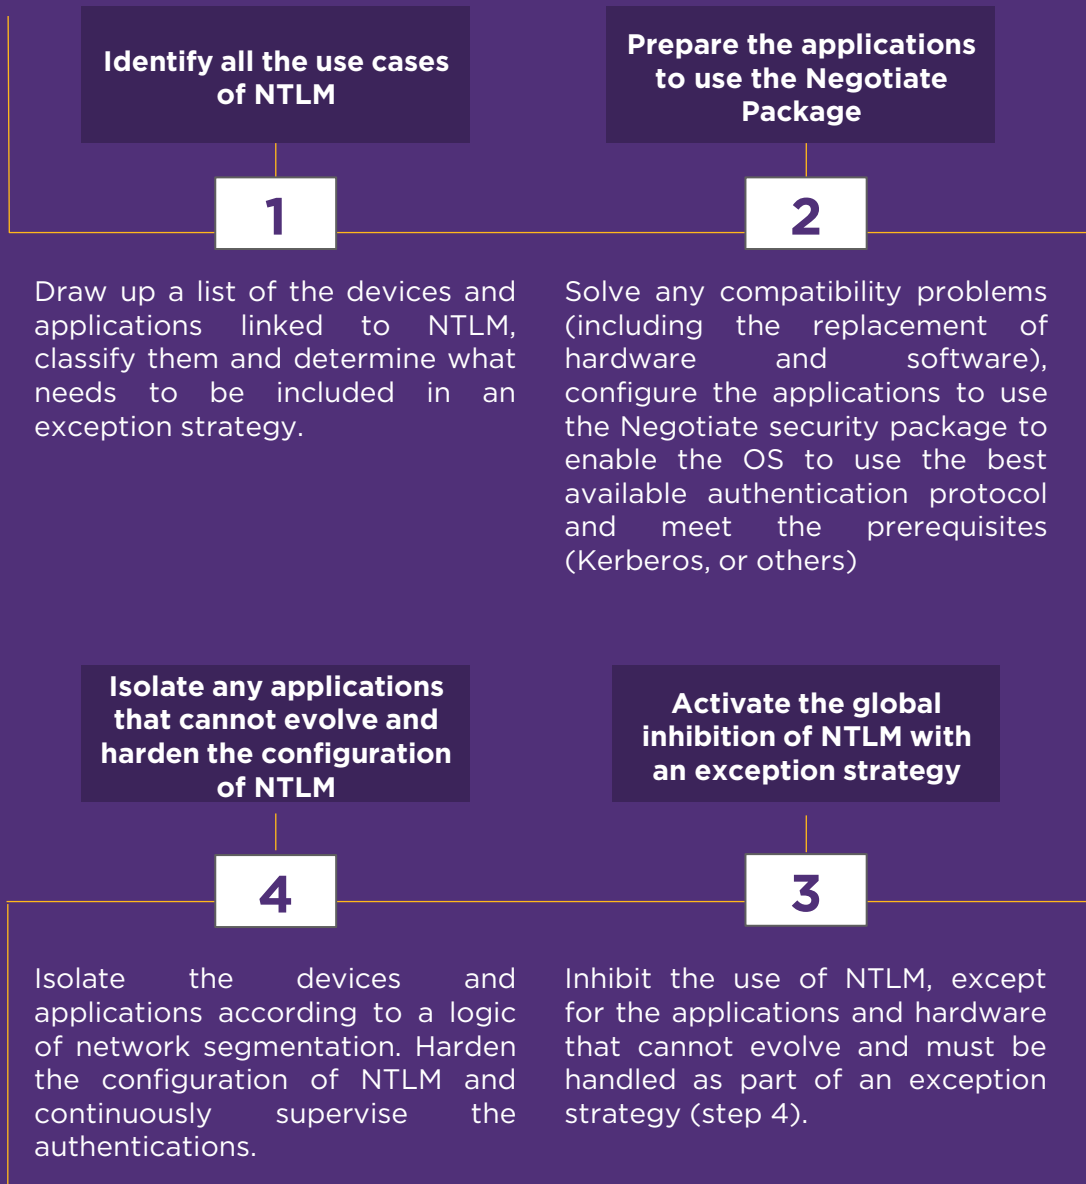
The complete removal of NTLM from an environment undeniably improves security by avoiding Pass-The-Hash (PTH) type attacks. However, it does not prevent other classes of attack, such as clear password theft or the theft of Kerberos Ticket-Granting Tickets (TGT).

Organizations are encouraged to implement Kerberos or to use modern authentication protocols (OpenID Connect, SAML, etc.) for their existing applications, because Microsoft does not expect the NTLM protocol to be improved.

Why is NTLM still used?

NTLM is still widely used due to legacy applications that have not evolved, but also due to poor application configurations, which support Kerberos nevertheless.

Doing away with NTLM in four steps



It is inadvisable to universally inhibit the use of NTLM, without first mapping out the legacy applications and conducting an impact analysis.

Feedback on the attacks observed by the CERT-W

AD at the heart of the threat from ransomware and the hackers' methods

Since it is at the heart of information system security, AD has become the preferred target of large-scale cyber attacks.

The [2020 Wavestone CERT benchmark](#) is unequivocal.

“AD was compromised in 95% of the cyber crises handled by the CERT-W”

These findings are shared by the French National Agency of Information Systems Security

(ANSSI): *An analysis of the methods used in recent attacks highlighted a rise in the targeting of Active Directory directories.*


This observation can be explained by the central role that AD plays in the enterprise IS.

Acquiring high AD privileges often enables hackers to take control of the entire Windows ecosystem, or even to reach other environments through the workstations of developers or administrators. Further to this breach, sensitive data can be extracted or business activities and services can be durably disrupted, in particular by ransomware attacks.

2019 and 2020 saw an unprecedented increase in ransomware attacks.



A large majority of the crisis interventions by the CERT-W in 2020 in all sectors were in response to ransomware attacks.



An observation shared by the Microsoft Detection and Response Team, or DART.

More recently, a high proportion of demands for ransom payments have been made for both the decryption of encrypted data and the non-disclosure of stolen data.

This combination of locking down the IS and leaking data has gradually spread as a result of the actions of the Maze group in North America.

Attackers make use of the possibility to execute distributed code offered by AD, and the Group Policy Objects (GPO) or the local administration rights on machines linked to the AD, to deploy loads that encrypt the system. GPOs and direct access to systems through legitimate administration tools, which are sometimes used simultaneously, allow the entire Windows environment to be encrypted in a matter of hours.

Are backups safe from hackers?

While it is still unusual for hackers to specifically target backup infrastructures, they may be part of the collateral damage caused by attacks. In the absence of a dedicated backup infrastructure that is not integrated in the live AD forest, the deployment of a load that encrypts the entire environment will also encrypt the backup systems.

Groups of hackers only specifically targeted backups in a minority of the attacks investigated by the CERT-W, on either the backup server system platforms or through administration consoles that implement single authentication with AD.

Since this method maximizes the chances of receiving payment of the ransom, this trend will probably become stronger in the years to come.

Study of a ransomware attack based on an intervention by the CERT-W in 2020

This attack draws inspiration from an intervention by the CERT-W, following the deployment of ransomware in an environment comprising several tens of thousands of machines. Faced with the vulnerabilities and configuration faults that are frequently exploited, this analysis highlights the Tactics, Techniques and Procedures (TTP) of a ransomware operator. All the information could be used to identify the parties has been anonymized.

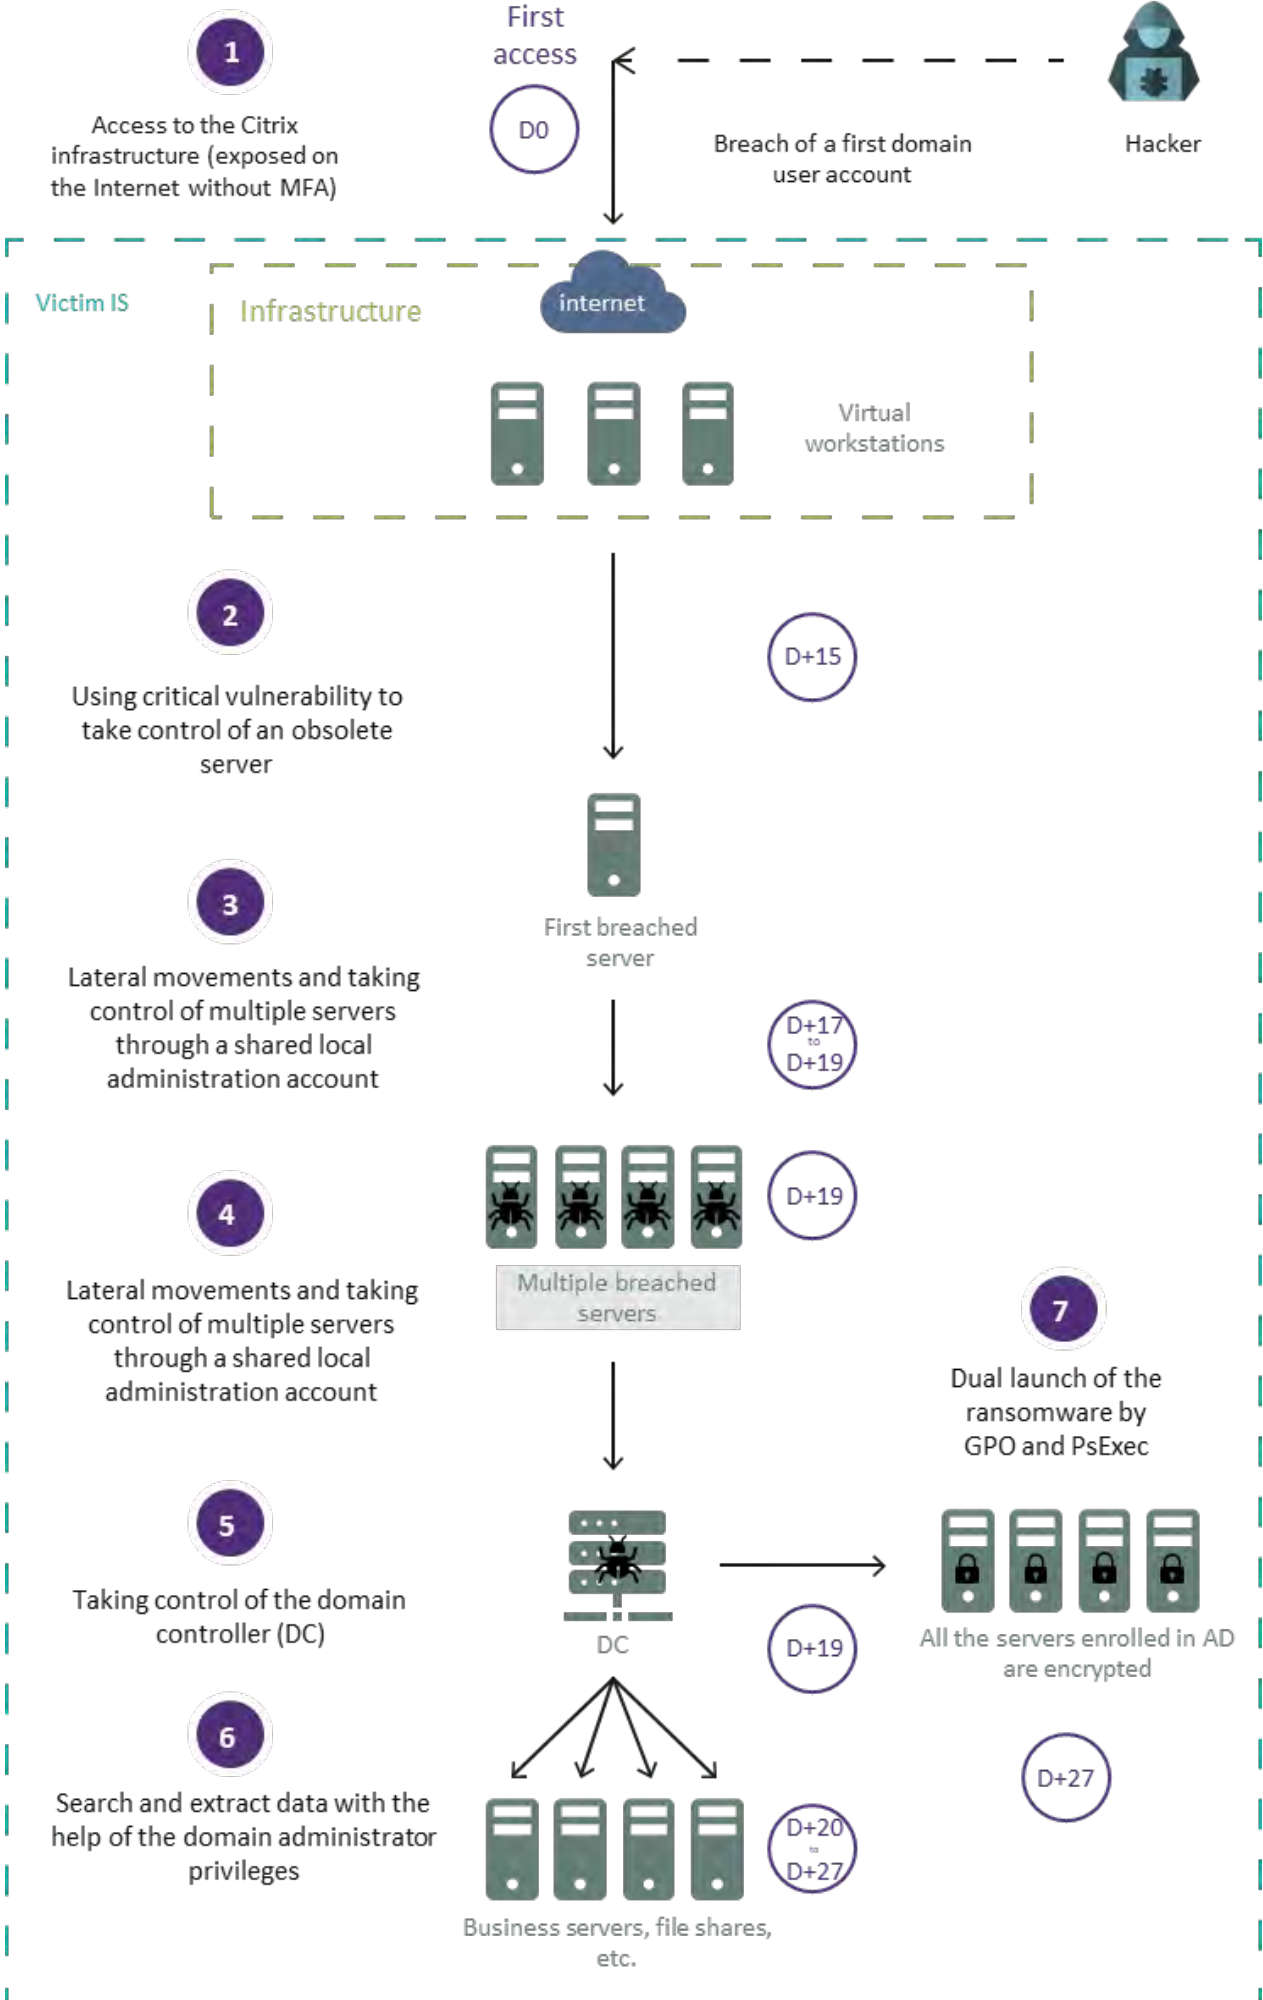
Once the security of a first resource has been breached, increasing the privileges in an Active Directory environment forms an integral part of the methods used by groups of hackers. In the majority of ransomware attacks, and in particular in attacks that use the Ransomware-as-a-Service model, no advanced attack techniques are observed. In this model, the cyber criminals are affiliated with a ransomware supplier in order to benefit from the encryption agent and the associated services (payment and contact infrastructure, publications site, etc.), in exchange for a share of the ransom. In this way, the attackers operate mainly opportunistically in order to gain a short-term return on investment.

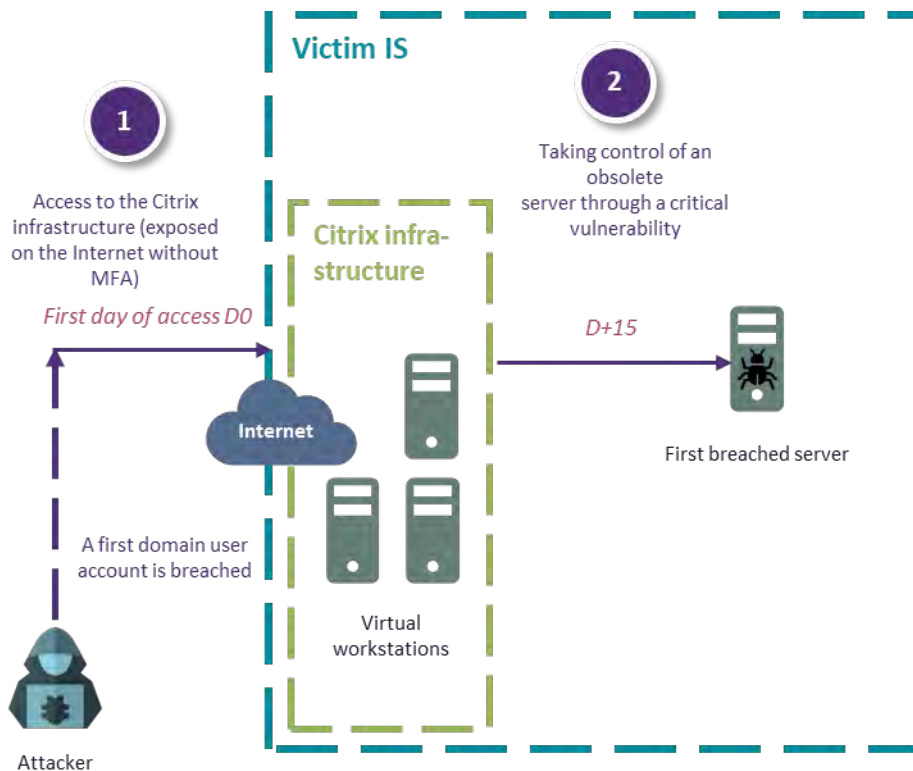
As a general rule, at least one of the phases of the chain of attack could have been avoided by adopting good basic cyber hygiene and computer security practices.



Looking beyond non-targeted and opportunistic attacks, groups of ransomware assailants use more advanced resources and skills in their operations. This trend, known as Big Game Hunting, enables groups of cyber criminals to target organizations with higher levels of IT security.

Diagram of the studied attack





Initial access

The intruder breaks into the IS by first breaching the security of a domain account through an infrastructure of virtual desktops exposed on the Internet. In the absence of multi-factor authentication, the breached identifiers can be reused by default.

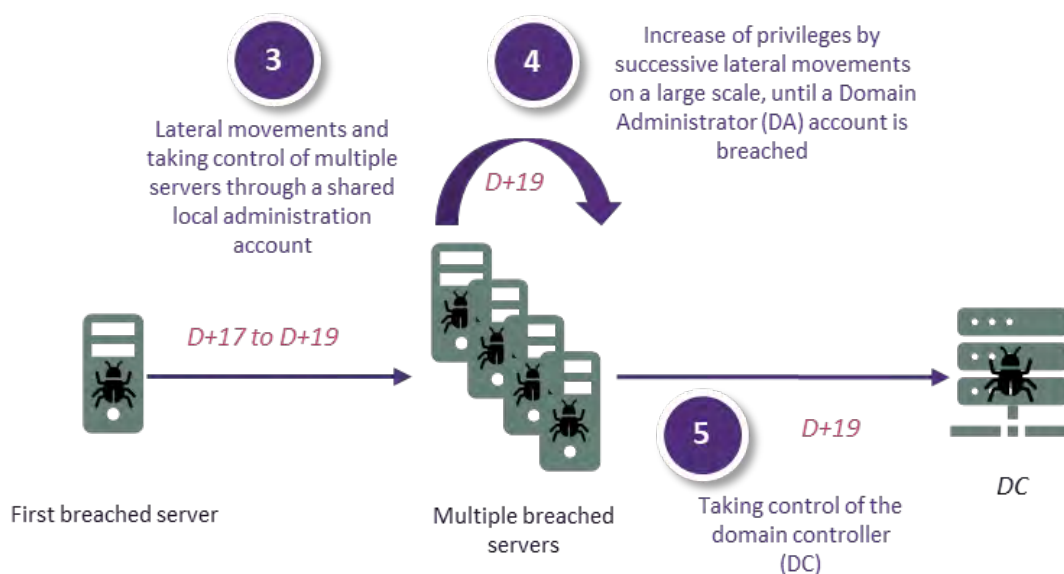
The exposure of the access service on the Internet without MFA exposes the IS to illegitimate access through the reuse of identifiers

While phishing and critical vulnerabilities remain the main channels of infection, reusing identifiers on exposed remote access services, without any multi-factor authentication,

accounts for almost one in five cases of initial breach investigated by the CERT-W. Examples include cases of attacks by brute force on the Remote Desktop Protocol (RDP) services exposed on the Internet. The assailant can then escape from the restricted desktop using a command interpreter, executed through authorized software, gain access to the system and start actively reconnoitering. A delay is often observed between the first malevolent access and the start of the active reconnoitering phase. This delay can be explained by the time it takes the group of specialized hackers to sell the access to a ransomware operator.

Hackers can use critical vulnerability on a Windows 2003 server (for which support has expired and no more security patches are provided) to take remote control of the server and set up a base camp.

The absence of the LAPS¹ solution makes lateral movements easier



Lateralization and increase of privileges

The attacker can then raise their privileges using rebounds from the local administrator's account on the breached server, whose password is shared with accounts on a multitude of other servers.

Consequently, a first privileged domain account is breached, then attempts are made to make lateral movements on a large scale on several thousand machines from this account, until a domain administrator's account is breached.

Obsolete servers must be isolated and must not lower the level of security of the entire IS

1. Microsoft's Local Administrator Password Solution (LAPS) offers the possibility to automate the management of the passwords of local accounts on the machines integrated in AD and guarantees that the password of one local account per machine is unique.

Despite the signature-based anti-viral solution installed on the targeted systems, utilities that are natively present on the Windows systems allow the memory of the LSASS memory, which contains the authentication secrets of connected users, to be extracted.

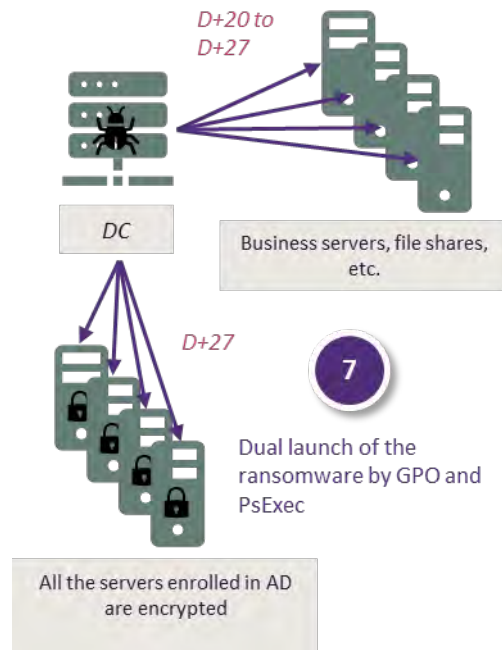
Using accounts with domain administrator privileges to perform routine tasks makes it possible to raise privileges using successive lateral movements.

If structured administration per level (the tiering model) is not implemented, the AD domain is exposed to increases in privileges

Once the domain has been breached, the attacker enters a post-operation phase in order to identify and extract the enterprise's sensitive data. The domain administrator privileges grant very wide access to the enrolled resources (network file shares, mailboxes, etc.).

6

Search and extract data with the help of the domain administrator privileges



Deployment of the ransomware

In the final step, the malevolent load is deployed by group strategy and an automatic process, executed from a domain controller, to allow for rapid propagation in the victim's IS.

As well as encrypting the breached machines, the encryption agent also deletes the Windows log files stored on the machines and any copies that are stored locally.

Supply chain attacks: an emerging trend

Supply chain attacks are a growing channel of infection, in addition to the more common vectors of breach. Despite the fact that they have already been documented for several years, attacks through suppliers and service providers have increased in the last 2 years, in terms of both numbers and scale. This growth can be explained by improvements in the IT security of corporate infrastructures, which force assailants to breach third parties in order to reach their final targets.

The privileges on the AD directory granted to third-party solutions, which are often set too high by publishers for simplicity's sake, can have disastrous consequences in the event of a supply chain attack.

Also, supply chain attacks offer a particularly attractive return on investment for groups of attackers, because compromising an initial entity can potentially open up points of access to many of its customers. These attacks, which are only made by hackers with the most advanced technical capacities, have been used by ransomware operators for financial gain and by States for spying.

The SolarWinds attack in 2020 made a strong impression. This attack demonstrated the potential scale of supply chain attacks, because 18,000 SolarWinds customers were impacted. This attack, which used highly sophisticated intrusion methods (injection of a load on compilation, deployment of dedicated infrastructures for each end target, etc.), also relied on conventional lateralization techniques and, in part, exploited ordinary vulnerabilities.

The hackers very quickly created secret doors, discreetly opened channels of communication and camouflaged and concealed their traces, while looking for means of obtaining higher privileges. But they also used known techniques, such as reusing compromised passwords or lateral movements with identical administrator accounts on all the machines.



How can the situation be improved?

This chapter explains the new AD Enterprise Access Model security model, then makes recommendations on how to make on-premises AD and Azure AD secure and describes a road map for modernization with Azure AD.

Enterprise Access Model

In 2012, Microsoft introduced the tier-based administration model, which aims to partition the authentication secrets in an Active Directory environment.

The principle of implementation consists of creating a partition between the administrators, according to the resources they manage. This helps to protect the authentication secrets and prevents breaches at a lower trust level from being propagated to a higher trust level.

This concept is based on the Bell LaPadula model, which was introduced in the 1970s.

This model separates the administration of resources into three levels, according to their criticality. This means that the administrators who manage the users' workstations are separated from those who manage the servers and those who manage the enterprise identity repository, which is Active Directory in this case.

The documents [Mitigating Pass-the-Hash and Other Credential Theft, versions 1 and 2](#) describe this administration model in detail, associated with an Enhanced Security Admin Environment (ESAE), widely referred to as the hardened forest.

In December 2020, Microsoft upgraded this administration model to take into account cloud and hybrid environments. The new [Enterprise Access Model](#) is an evolution of the preceding model, in which the concept of the tiered model remains, but has been restructured, with new terminology.

The ESAE approach has been withdrawn from the general recommendations, because it is complex and costly to implement. But the implementation of an administration forest may still be appropriate in [certain cases](#), and in particular for disconnected environments.



Understanding the preceding tiered model

It is essential to understand the principles of the tiering model in order to fully grasp the best security practices in a Microsoft environment.


Technical assets are used to isolate the tiers.

Tier 0 is the most privileged level and includes the accounts, groups, domain controllers and resources that have direct or indirect control over Active Directory. Therefore, tier 0 includes the servers connected to Active Directory (domain controllers), plus any other components with close interaction, such as federation servers, WSUS update servers, application deployment servers, internal PKIs and Azure AD Connect.

Tier 0 administrators can control and supervise the resources on every level (in the AD sense of the term), but they must only interact with tier 0 resources. They must do this on an administration workstation with hardened security that is on this tier.

Tier 1 refers to the servers and applications that are members of the AD domain, and the resources that gravitate around

them. The accounts that control these resources potentially have access to sensitive data. Tier 1 administrators can access the tier 1 resources and can only manage tier 1 resources in Active Directory.



Tier 2 includes the user devices (workstations, printers, etc.). For example, the support hot line and assistance service belong to this tier. Tier 2 administrators can only log onto tier 2 resources and can only manage tier 2 assets in Active Directory.

A new model for hybrid enterprises

The new [Enterprise Access Model](#) was created for hybrid organizations that have on-premises and multi-cloud applications that apply the Zero Trust security principles.

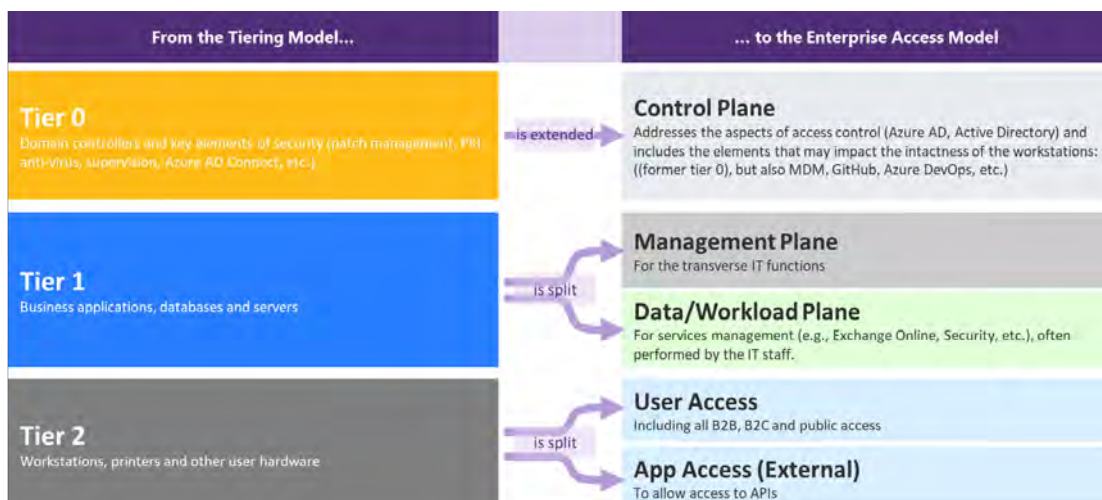
In this new model, security is not exclusively controlled using Active Directory, but also by Azure Active Directory.

The terms have changed, but the principles of separation into levels of privileges (tiering) remain.

Tier 0 has been replaced by the Control Plane. Management of the control plane must be tightly controlled and limited to very high-trust devices. The control plane has been extended beyond the resources found in the former tier 0, to include Azure AD and cloud solutions, such as MDM-type device management tools or development solutions, like GitHub/Azure DevOps.

The notion of management tiers has been introduced in the form of the management plane and the data/Workload plane, which are used to manage applications and data. This was previously known as tier 1. Finally, there are the users and other applications that consume services (applications and data), including internal users, partners, customers, etc.

The Enterprise Access Model does not explicitly mention the user workstations, which were previously in tier 2. Instead, the control plane of Azure AD is used to determine which types of devices can connect to this or that service or application. This is called conditional access control and it forms the cornerstone of a Zero Trust architecture.



A guide to Securing Privileged Access

In its guide to [Securing Privileged Access](#), Microsoft shares a baseline implementation that illustrates the above-mentioned Enterprise Access Model.

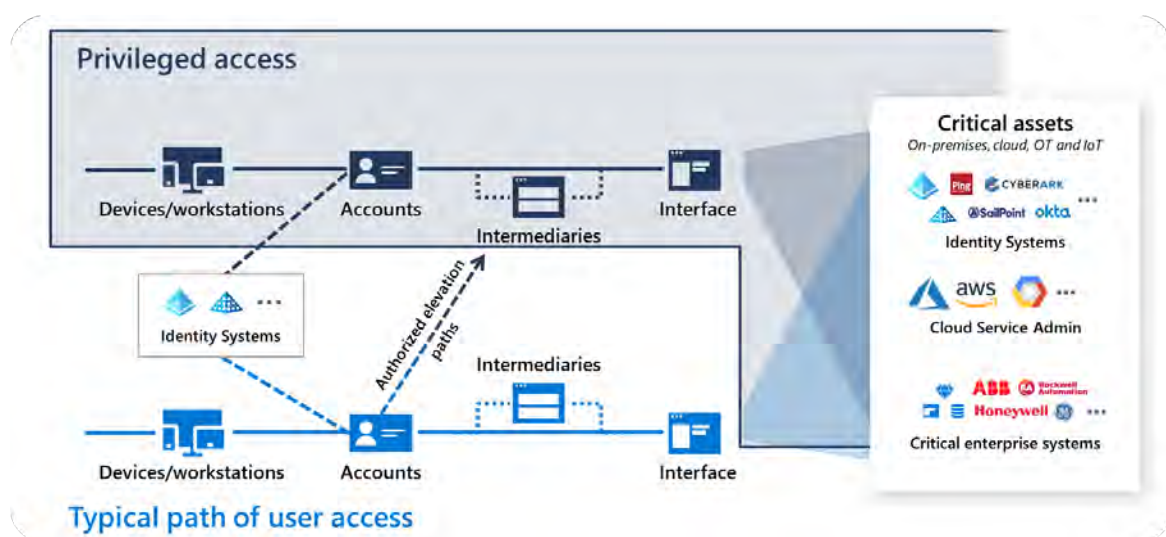
The goal consists of strictly limiting the capability to perform privileged actions to authorized paths, while disrupting the return on investment of the hackers.

In addition to prevention, the access paths are closely monitored to detect any anomalies or devious behaviors.

The strategy consists of encouraging organizations to first use the rich functionality that is natively available in the cloud.

In this Microsoft guide, implementation is based on the security solutions available in Microsoft 365 Enterprise E5. It is a starting point for the implementation of a Zero Trust architecture in a Microsoft environment.

Conditional access control and strong authentication are universally applied to all users. The Microsoft Cloud App Security solution, combined with Identity Protection, is used to supervise sessions.



In this case, the use of privileged accounts (which are flagged as sensitive) is explicitly restricted to specific devices, with the conditional access control in Azure AD Premium.

An Endpoint Detection and Response (EDR) is deployed on the workstations. Microsoft Defender for Endpoint is capable of interacting with Azure AD, through the MDM, in order to escalate the state of health of the devices.

The workstations are managed, and the security profiles are deployed using the Microsoft Endpoint Manager MDM.

In the specialized and administration workstation profiles, the device's user is no longer the administrator of their workstation.

In addition, the list of applications that are authorized to be executed is restricted.

Also, the principles of least privilege and just-in-time access (temporary increase) of the administration rights are implemented with Azure AD Premium.

In this case, this baseline implementation can be used to administer resources in the cloud, but also critical assets, such as Active Directory, through intermediaries (VPN, rebound server, etc.).



Securing tier 0 and implementing the control plane

Irrespective of the reason for launching a project to reinforce the security of AD (action plan following a major security breach, results of intrusion tests or red team exercises, etc.), a significant preparatory phase is necessary before launching such a vast program. The priority consists of focusing on tier 0 and taking account of the other tiers in view of the transformation of the IS (use of the cloud and Azure AD).

Step 1: prepare - 1 to 6 months

Even if the general outlines are known when the project to secure AD kicks off, it is essential to define the boundaries of the project.


Obviously, the duration of the scoping phase depends firstly on the complexity of the environment. The challenges of every project are unique, from the simple case of a single AD forest with a single domain that is operated by a central team, to the more complex case of a large organization that is present on every continental plate and has made numerous acquisitions,

resulting in the creation of trust relationships and a multitude of forests.

It is necessary to secure tier 0. The efforts made will not be wasted, even in the event of a cloud transformation

The three main goals of the scoping phase are to:

1. Map out the legacy environment (forests and the corresponding trust relationships) and identify any vulnerabilities.
2. Determine the security target to be reached and the deadlines (e.g., priority scopes, level of hardening of the target, dedicated character of the infrastructures, forests that can be decommissioned, removal of memberships between the AD backup system and AD itself, etc.).
3. Define the project structure that will allow the targets set in the schedule to be reached.



In these phases, any existing action plans are also taken into consideration (reports by the General Inspectorate, audits already completed, findings of the red team, etc.), so that they can be incorporated in the global project.

The approach to drawing up this map is based on:

- / known market tools (e.g., PingCastle, BloodHound, OAADS, etc.) or frameworks (e.g., the ANSSI checkpoints);
- / interviews with contacts in the businesses or geographies in order to identify autonomous Active Directory infrastructures that may not have been detected by the tools.

Once all this input has been collected, the goals of the transformation project are defined and the paths leading to those goals are identified. This type of project is usually broken down into several sub-projects.

Rationalize the infrastructures (forests and domains to be

decommissioned or migrated) and the trust relationships, wherever possible.

Harden and correct the identified vulnerabilities, update assets whose OS is no longer supported, harden the system and AD, regularly install the security patches, etc.

Implement the tiered model (tier 0 is the priority) and deploy the architectural changes to be made (partitioning, deployment of assets dedicated to tier 0, implementation of dedicated administration workstations, use of administrative silos, etc.).

Define the administration RACI and model: the activities of the teams, types of privileged accounts to be made available to them and administration silos.

Prepare the reconstruction, with the transfer of the backups onto systems that are not members of Active Directory, reconstruction tests, etc.

Recommendations on the use of public tools

Numerous Open Source scripts and tools are available on the Internet to assess the level of security of an Active Directory / Azure AD environment. However, a number of cautionary rules should be obeyed before executing them:

1. Wherever possible, review the source code of the tool to check its behavior.
2. Execute the tool with the lowest required privileges, according to the least privilege principle, and in a restricted environment (e.g., a dedicated virtual machine).
3. Limit outgoing flows according to needs (no flows onto the Internet for an Active Directory configuration review, only to Microsoft resources for an Azure AD review)

Is it necessary to create an administration forest?

As explained previously, Microsoft no longer recommends this model, except in the special cases described [here](#), because it is complex and costly to implement.

What about the administration of the virtualization infrastructure hosting tier 0 assets?

The critical nature of the infrastructure hosting tier 0 assets requires it to be dedicated and its administration to be integrated in tier 0. The use of a shared platform incurs the risk of a lack of control of access to these virtual machines.

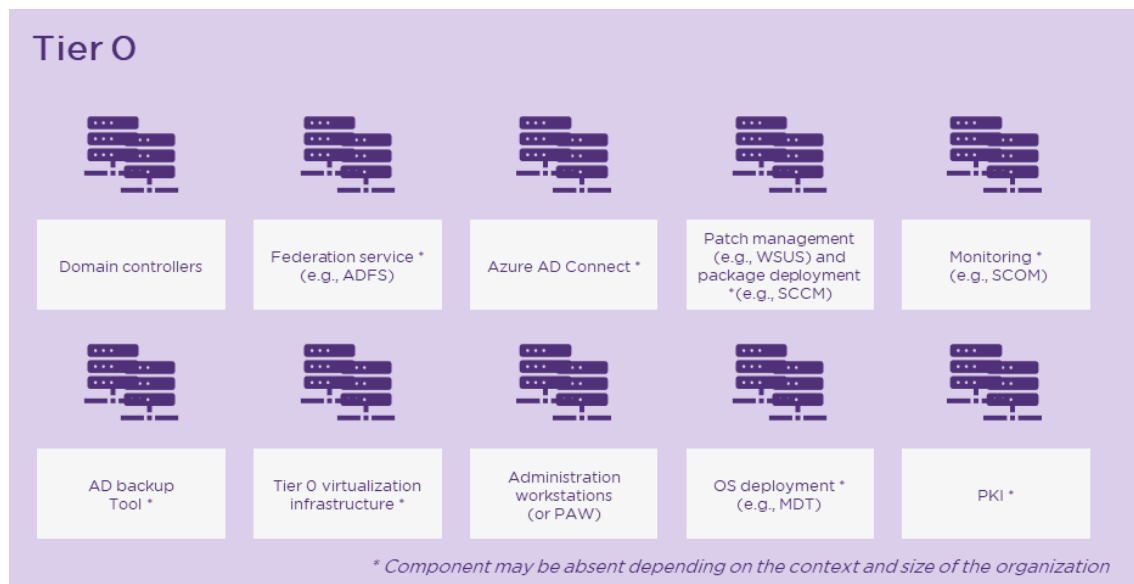
However, providing a dedicated infrastructure can take a long time in certain organizations. To mitigate the risk more rapidly, an existing shared infrastructure can be used tactically in order to start taking the security-related actions described in the next section without delay.

The residual risk mentioned previously can then be completely covered by a migration to the target dedicated infrastructure a few months later.

Step 2: implementing tier 0 - 6 to 24 months

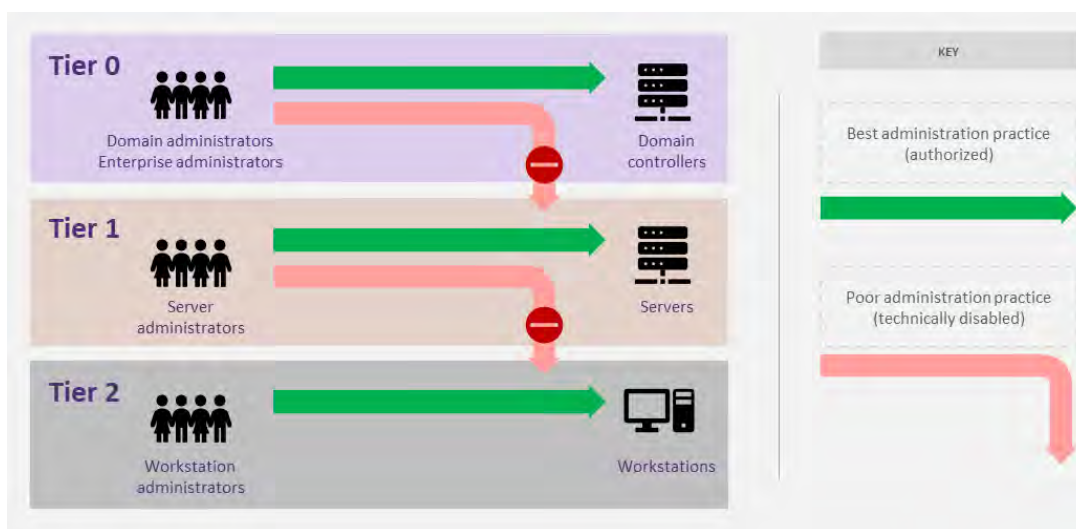
At the start of this phase, the final workshops take place to fine-tune the targets to be reached on the basis of the scoping, to bring the teams onboard and to present the sequence of actions.

Identification of the organization's various administration teams (e.g., IT support, AD teams, server teams, workstation teams, etc.), their responsibilities and activities and the rights they need to perform them (least privilege principle). This step defines a clear RACI and prepares the future accounts that will be used, in accordance with the delegation model of the tiering.



Then come the central steps: the application of the tiered structure to Active Directory, the creation of the administration accounts specific to each tier and the movement of the objects into the right Organizational Units (OU). The final step consists of

disabling connections using an account in a given tier to an asset in a lower tier (as shown in the diagram below). This can be done, at first, through "deny logon" GPOs, then in a more permanent way through Authentication Policy Silos.



Deploying a "deny logon" GPO technically prevents logins from Tier 0 accounts to the machines on which the GPO is applied (Tier 1 & 2). However, GPOs are client configuration items, which could be disabled locally by an attacker (in order to trick a Tier 0 administrator into logging in on a compromised machine for example). The "deny logon" GPOs therefore protect against administration errors but have flaws that can make them vulnerable in case of an attack. Authentication Policy Silos allow authentication of certain accounts (typically Tier 0 administration accounts) only from certain machines (typically administration workstations).

Compared to the GPO blocking mechanism, Authentication Policy Silos are no longer based on a client configuration but on a mechanism imposed by the Active Directory services. This mechanism also covers the risk of replaying Tier 0 credentials (Kerberos tickets) elsewhere than on an administration workstation (which must therefore also be compromised). However, the greater complexity of implementing this method compared to GPO's can lead to a step-by-step approach: in the short term, implement GPO's, in the medium term, implement Authentication Policy Silos to further strengthen the security level.

Finally, this step must also include the definition and deployment of the Privileged Access Workstation (PAW) for administration purposes. This workstation must only be able to log onto tier 0 assets to take administrative actions. Therefore, it can be hardened to the highest standards and only the software used to make remote connections should be installed on it.

Access to the Internet (apart from the Azure cloud portal in a hybrid infrastructure, for example) and to e-mail must not be possible, in order to limit the exposure of this workstation. If necessary, a zone for exchanges between the office and the administration environments should be provided.

Difficulties: membership of the administration bastion that already exists in the organization, with its established model, can sometimes be a source of complexity, despite the advantages that it offers (i.e., easy implementation of multi-factor authentication, records of actions).

Hardening

Hardening amounts to taking the security measures wherever possible, in order to raise the level of security.

Hardening tier 0 assets: formal definition and application of a guide to hardening, minimization of the number of agents on tier 0 assets, because they can

constitute a channel of breach. Ideally, only native Microsoft software (anti-virus, integrated backup, transmission of events by WEF⁽¹⁾, MDI⁽²⁾ supervision agents, Application Control, etc.) should be present.

From the strict perspective of risk reduction, using the tier-based model takes priority for hardening purposes, even if it is more complex, due to its impact on administration practices.

It is also necessary to modify the configuration of these assets in order to use the MCO/MCS systems dedicated to tier 0 (supervision, OS updates, software deployment, etc.).

The remediation actions on privileged accounts must also be taken in this step: replacement of the accounts belonging to the Built-In groups by accounts with rights that apply the least privilege principle, mapping of the service accounts, use of Managed Service Accounts (MSA/gMSA) wherever possible, identification and quarantining of inactive accounts, activation of LAPS, etc. The list of Active Directory checkpoints provided by ANSSI can be used to raise the level of security in steps.

(1) Windows Event Forwarding (2) Microsoft Defender for Identity

Difficulties: the service accounts of certain solutions that demand to be a member of the Domain Administrators or Enterprise Administrators groups, often for simplicity's sake. Accounts whose membership is unknown and whose impacts in the event of an outage are difficult to estimate. The implementation of a reliable process to deal with inactive accounts, apart from by mass manual processing as part of the project.

Backups

On the whole, organizations have a strong command of the backing up and restoring of machines. Nevertheless, it is necessary to review this question in view of the current threat and of the worst-case scenario, in which Active Directory and its backups are breached and totally destroyed. The backup infrastructure is often linked to the AD domain that it backs up. Despite this program to increase security, it is important to remain cautious and to always remember that a breach is possible (assume breach). The goal consists of completing the existing backup system, which will be maintained due to its performance and restore times, with an external backup system that is disconnected from AD.

Detection

Once the tier-based model is in place, and since the modes of administration are known and standardized, detection scenarios can be deployed to detect any

practices that deviate from the model (e.g., addition of an account in the built-in groups, modification of sensitive and normally stable configurations of AD, use of emergency accounts, etc.). Finally, the implementation of attack technique detection scenarios (e.g., hash retrieval, high numbers of Kerberos ticket requests in a short lapse of time, etc.). Microsoft provides a basic list of AD events to be collected. Alert and incident processes must obviously also be described and operational.

Difficulties: the collection method (use of WEF), which may be different from the standard defined by the SOC (collection by an agent installed on the asset) and require adaptations.

Rationalization

In parallel to these actions to reinforce security, it is also necessary to keep track of the decommissioning or migration plan defined in the scoping phase.

Any delays in the schedule would immediately incur a serious risk, since no measures would be taken to increase the security of these scopes. Especially if they have trust relationships with the other forests.

Difficulties: uncertainty about the precise impacts when the trust relationships are broken, or the impacts of decommissioning.

Program supervision, change management and reporting

It is essential to mention all the cross-functional activities that are inherent in the supervision of the program, in order to be exhaustive. Scheduling, the synchronization of the teams and communications with the stakeholders are all key factors, in view of the number of technical actions to be taken, their interdependency and their potential impacts.

Moreover, the significant changes in administration practices (e.g., use of separate administration accounts for each tier, use of a PAW) required by the implementation of the tiered model, means that close attention must be paid to change management, to avoid any disruption of the activities.

Finally, it is essential to make sure that the communications and reports on the state of progress of the program are sufficiently simple and clear to avoid any frustration and misunderstandings about the inevitable roadblocks and requests for arbitration.

Difficulties: the difficulty in making the link between the performance of technical actions and risk mitigation, the major effort that the teams tasked with the AD run are asked to make in order to take the technical actions.

Step 3: guaranteeing durability - 1 to 3 months

More than any other component of the IS, the security of AD cannot be reduced to a program that lasts for a finite duration, but it must be transformed into a regular process that checks the overall level of security. This process must be executed regularly, and any identified deviations must give rise to corrective actions in the short term, and preventive actions in the long term.

This control plan can call on a variety of resources: scripts, additional market tools, ANSSI's Active Directory Security (ADS) service or manual checks, when they cannot be easily automated. It is also possible to call on annual red team audits and exercises to check that tier 0 can no longer be breached.

Finally, as well as checking that the backups of the Active Directory infrastructure are successfully completed

systematically, the complete restore process of a backup should also be tested, just like in any continuity plan.

These tests also measure the time required to completely restore a backup. These measurements should be sent to the team responsible for business continuity and they also provide an indicator that needs to be optimized, test after test.

SUMMARY: Securing Tier 0



Rationalize and decommission

Focus efforts on the long-term scopes and decommission the rest.



Implement tier 0

Partition Active Directory against the risk of breach.



Keep the components in a secure condition

Install security patches and harden the configurations.



Back up and run restore tests

Put the backups out of reach and be prepared to rebuild.



Supervise the actions and manage change

Oversee and rigorously deploy this action plan.



Centralize the log files and implement detection

Monitor and detect weak signals in order to respond rapidly.

Securing Azure AD subscriptions

The Identity Secure Score: a first step

The Identity Secure Score is a native indicator that compares the situation of the organization with Microsoft's best identity security practices.

This score was introduced in 2020 on the basis of the Microsoft Secure Score presented at the end of 2017.

The following calculation principles apply:

- / Percentage of implementation of the checks proposed by Microsoft
- / The checks measured depend on the identity architecture and the available licenses (see table below)
- / A single security license activates the appearance of the check

If all the checks are available to the organization, the score is

currently broken down as follows:

- / 61 points for the cloud identities (Azure AD settings, conditional access, self-service password reset, Azure AD Identity Protection)
- / 58 points for the local identities (with Microsoft Defender for Identity)

Changes can be explained by two cases: a change in the organization's situation or changes made to the checks by Microsoft.

The reasons for these changes can be tracked by looking at the changes affecting the Microsoft Secure Score in the Microsoft 365 Security Center.

Note that the Microsoft Secure Score does not include all the Identity Secure Score checks.

In hybrid organizations without a Microsoft Defender for Identity license, the following checks are made:

License	Checks	Weight	Microsoft Secure Score
Free	Designate more than one global admin	1	x
Free	Do not allow users to grant consent	4	x
Free	Use limited administrative roles	1	x
Free	Enable Password Hash Sync	5	x
Free	Do not expire passwords	8	x
Free (P1)	Block basic authentication	8	x
Free (P2)	Enroll all users for strong authentication	9	x
Free (P2)	Require MFA for administrative roles	10	x
Azure AD P1	Enable <i>self-service password reset</i>	1	x
Azure AD P2	Activate the user risk policies	7	x

Free (P1 or P2): A Premium license is required to customize the security measures



Going further than Secure Score

The Identity Secure Score produces metrics of the organization's situation relative to Microsoft's best practices, which are indispensable but not exhaustive.

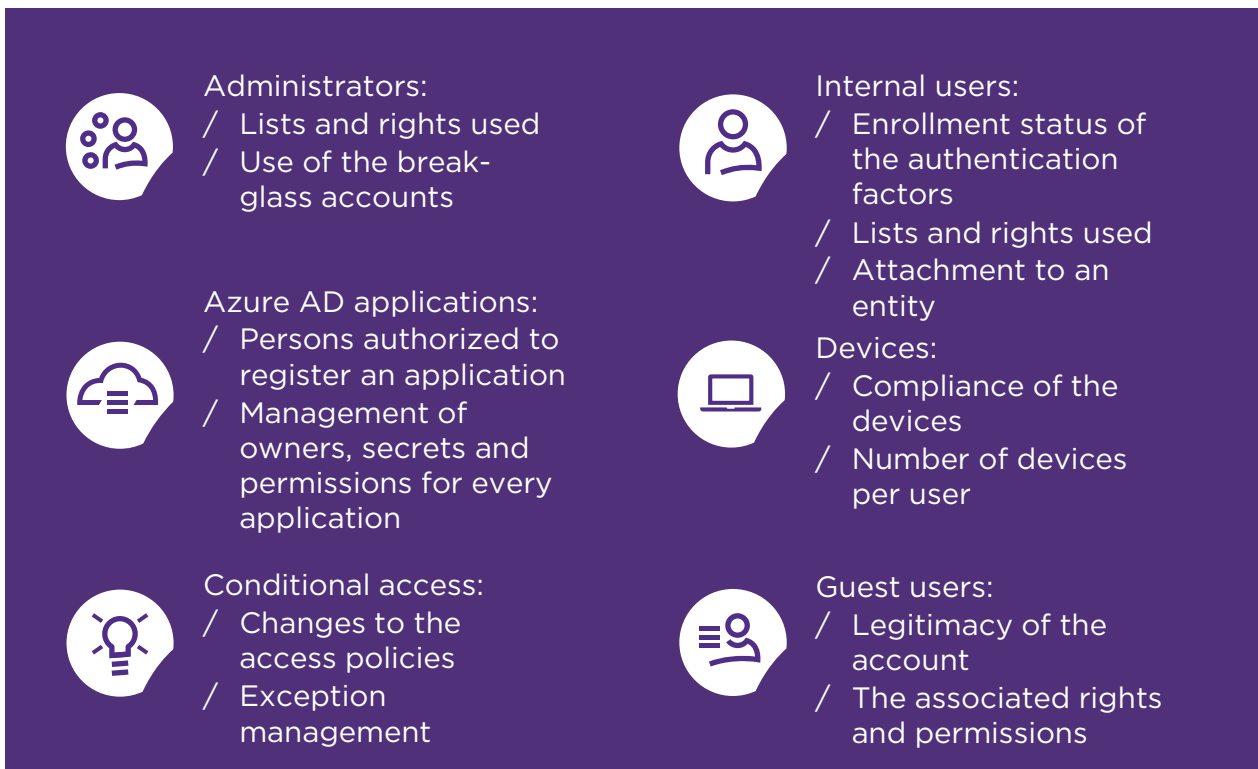
It is also essential to make sure that all the Azure AD settings are consistent with the issues facing the organization (e.g., domains authorized for guests). It only takes a few hours to review and define concrete means of making improvements.







Looking beyond the configurations of the platform, it is more than advisable to implement a permanent control

plan to keep track of the objects that are bound to multiply (cloud or on-premises applications, internal and external users, conditional access policies, etc.).

This control plan will keep Azure AD in an operational and secure condition, just like the existing plans for the local ADs.

These checks in the program must cover the following points:



-  Administrators:
 - / Lists and rights used
 - / Use of the break-glass accounts
-  Azure AD applications:
 - / Persons authorized to register an application
 - / Management of owners, secrets and permissions for every application
-  Conditional access:
 - / Changes to the access policies
 - / Exception management
-  Internal users:
 - / Enrollment status of the authentication factors
 - / Lists and rights used
 - / Attachment to an entity
-  Devices:
 - / Compliance of the devices
 - / Number of devices per user
-  Guest users:
 - / Legitimacy of the account
 - / The associated rights and permissions

Understanding the roles in Azure AD

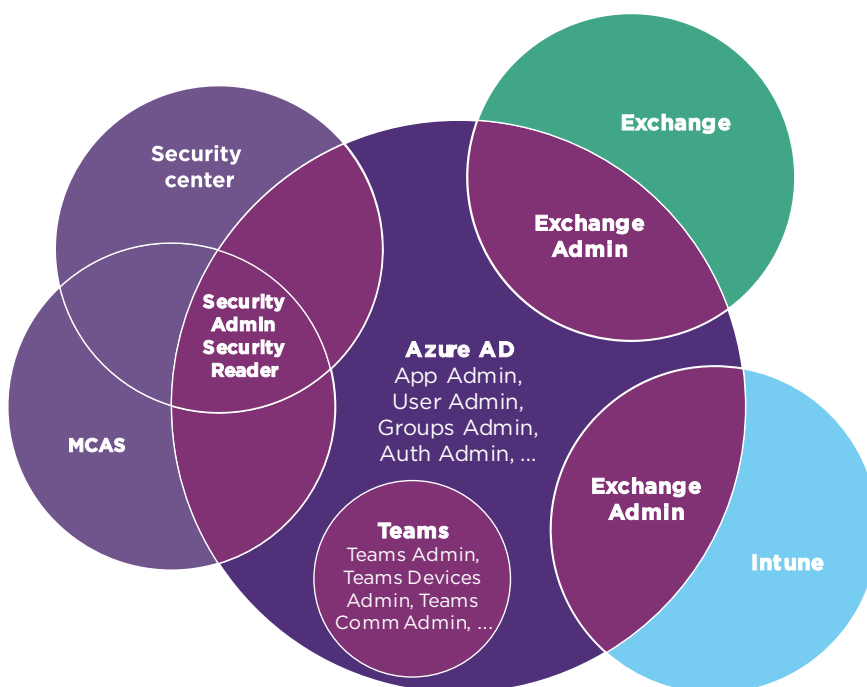
Azure AD roles are an important part of the Microsoft cloud security system. There are more than [70 integrated roles](#) that apply to both the management of the Azure AD resources, but also to certain Microsoft 365 (SharePoint Online, Exchange Online, Teams, AIP, etc.) and Azure (Azure DevOps, etc.) services. To take things further, it is possible to create your own roles using customized roles.

The following diagram shows the roles specific to Azure AD and the roles that apply to other Microsoft 365 services.

By default, the Azure AD and Azure resources are secured independently of one another. Nevertheless, Azure AD general administrators can elevate their access to manage all the subscriptions and resources in Azure.

Through the User Access Administrator role, this higher level of access allows for interaction with all the Azure subscriptions that trust their Azure AD tenant. Consequently, general administrators can use this role to grant access to the Azure resources to other users.

Understanding the roles with rights extending beyond Azure AD



Understanding the applications in Azure AD

Amongst the identities managed in Azure AD, it is essential to understand applications, which are very different from what existed on-premises.

Registration of an application

An application that wants to outsource authentication to Azure AD must be declared in Azure AD by *application registration*, which registers and uniquely identifies the application (AppId) in the directory using the notion of an *application object*.

This application registration is made in the tenant of the owner of the application. (e.g., Exchange Online is registered in the Microsoft tenant).

The owner of the application can then declare APIs, which will be used by the application (e.g., read and write an e-mail in Exchange Online).

The application can then be used, either in its tenant or in an external tenant, depending on its configuration.

To this end, the application object is used as a model to create one or more main service objects. A service principal is created in each tenant, in which the application is used.

It is important to note that the administrator of the tenant, in

which the application is registered, can add credentials (secrets or certificates).

Anyone in possession of the credentials can use the permissions of the application.

Applications are represented in Azure AD by two classes of objects:

- / *Application objects*, which store the information about the application,
- / *Service Principal objects*, which represent an instance of the application.

Service Principal

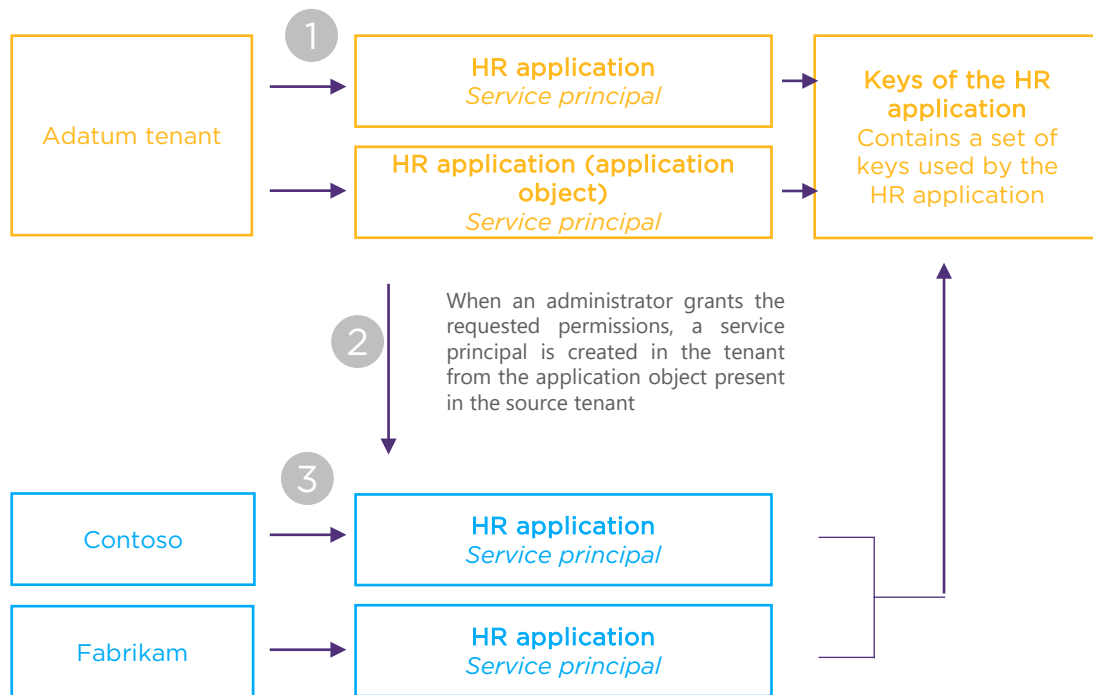
Authentication is necessary through either a user or an application, also called the service principal, in order to consume the resources protected by Azure AD.

Service principals are a type of object that exists in Azure AD to represent an application.

In particular, a service principal is created when an application or a code must access or modify resources, which can only be facilitated by an identity with the necessary authorizations. Creating an identity for an application enables the administrators to attribute roles and authorizations to it.



Applications and service principals



There are three types of service principal:

/ **Application:** An instance of an application created by an administrator or a user, when the requested permissions are granted.

When the application is declared as multi-tenant, a service principal is created in each tenant that adds the application (with a different specific ObjectID identifier).

In this way, specific policies and permissions can be applied, as well as an authentication and authorization specific to each tenant.

This SP is created in every other tenant after consent.

The owner of an instance of the application can add credentials.

/ **Managed identity:** An identity used by an Azure service to obtain an Azure AD token without having to use authentication secrets.

The registered enterprise applications, granted permissions and the corresponding credentials must be reviewed on a regular basis

/ **Inheritance:** A long-standing service principal similar to applications, but which can only be present in a tenant (not recommended).

Securing the service accounts

Knowing how to manage credentials in the code can be a challenge when creating cloud applications.

Ideally, they are never handled on the developer workstations and they are not present in the source code of the applications.

There are three types of service accounts in Azure AD:

- / **Managed Identity:** recommended for native Azure services
- / **Application:** recommended if the service is not native in Azure or is multi-tenant
- / **Dedicated user account:** not recommended if the application supports OAuth

Common recommendations

No matter which type is chosen, a life cycle should be defined (creation, review and deletion) and the least privilege principle should be applied:

- / Prefer OAuth2 permissions (e.g., read the files) rather than native Azure AD roles (e.g., SharePoint Online administrator)
- / Use different service accounts for different permissions
- / It is also advisable to limit the authorized applications by element for Exchange Online and SharePoint Online.

Protecting application credentials

Since service principals do not currently support conditional access, there is a risk that the credentials will fall into the wrong hands.

In order to protect oneself, it is necessary to:

- / Only use applications created in one's own tenant
- / Store secrets in Azure Key Vault

There are two types of Managed Identity.

A managed identity allocated by the system is directly activated on an Azure resource. When the resource is activated, Azure creates an identity for the resource in the Azure AD tenant. Once the identity has been created, identification information is sourced on the resource. The life cycle of an identity allocated by the system is directly linked to the Azure resource.

An identity managed and allocated by the user is created as an autonomous Azure resource. Azure creates an identity in the Azure AD tenant that is approved by the subscription with which the resource is associated. Once the identity has been created, it can be allocated to one or more Azure resources. The life cycle of an identity attributed by the user is managed separately from the life cycle of the associated Azure resources.

Understanding Azure AD licenses and their benefits in terms of security

It is impossible to talk about identity security in a Microsoft environment without mentioning the different levels of licenses.

Azure AD Premium adds advanced administration functions, conditional access control, dynamic groups, identity protection, self-service functions for users and a higher service level (99.99% guaranteed since 2021).

Application Proxy is another function of Azure AD Premium.. This service enables organizations to connect traditional intranet applications to Azure AD. This connection is made by combining modern protocols (based on identity claims) and older protocols, such as Kerberos or NTLM.

Users must possess the licenses shown below in order to benefit from the main security functions:

FREE AZURE AD / O365	AZURE AD PREMIUM P1	AZURE AD PREMIUM P2
<i>Security defaults</i>		
	Integration with MIP Self-service password reset Password protection Azure MFA Conditional access (1)	
		Identity protection (2) Identity governance (3)

(1) It is possible to use conditional access through a third-party identity supplier, but this will not allow for any granularity for the various Azure AD applications or the management of the duration of sessions.

(2) Identity protection: User or connection risk strategy and risk-based conditional access policies.

(3) Identity Governance: Azure AD PIM, access revisions, access package management.



Security defaults: a minimum security level accessible to all

In 2019, Microsoft introduced the *security defaults* for security policies that are predefined by Microsoft and guarantee a situation of minimal security, without any licensing conditions:

- 1 For administrators, activation of multi-factor authentication with every connection
- 2 For users, activation of multi-factor authentication for high-risk connections
- 3 For all users, registration of an MFA factor within 14 days of their first connection
- 4 Deactivation of inherited protocols
- 5 Access to the Azure AD portal for administrators only

The Azure AD Premium Plan 1 or 2 licenses are necessary to customize these policies.

Note that the security defaults cannot be activated at the same time as conditional access policies.



Which governance for Azure Active Directory?

In certain organizations, responsibility for Azure Active Directory is in a no man's land. This is largely due to the absence of a target:

- / a simple technical directory for Office 365/Teams or a future identity repository for the organization's applications?
- / an identity repository for cloud applications only, or for internally hosted applications too?

Who are the players today?

Three players are usually involved in the administration of Azure AD:



the Identity team, which configures synchronization and federation or connects applications.

These operators are used to having high-level rights on Active Directory and they inherit the Global Administrator's role, even if they do not yet have the skills required by this role.



The Digital Workplace team, which configures the collaborative services. The scope of this team is often extended beyond Exchange or SharePoint Online for convenience's sake. Even if they were not responsible for identities in the past, it is often the same team, which has the skills required to manage all subjects related to collaboration (e.g., guests, conditional access or third-party applications).



The Security team, which defines the policies for security and the control of the installed applications.

Which principles for the governance of Azure AD?

The logics used to govern Active Directory and Azure Active Directory are similar. The operational model must apply two fundamental principles: the segregation of rights and the least privilege principle. In a word, everything that can be delegated must be delegated.

On the other hand, Azure AD does not allow for the same level of delegation as Active Directory, in which everything can be delegated by changing the environment, despite the fact that Microsoft recently introduced some changes in this regard.

What are the possibilities to delegate?

First, it is necessary to differentiate the container (Azure AD and the underlying synchronization and federation infrastructure building blocks) from the content (the user and application identities and identities linked to terminals).

Microsoft proposes an RBAC model (Role-Based Access Control) with these native roles. These roles include the administration of the Azure AD tenant, of the Office 365 services and of the different identities. Note that only the global administrator role can edit certain general settings, such as the graphics of the login page.

In 2019, Microsoft introduced the possibility of creating **customized roles** using APIs. For the time being, only actions related to application administration can be used when defining one of these roles.

It is also essential to remember that Azure AD is currently designed as a centralized organization, even if Microsoft published the administration units in 2021, which are the equivalent of local organizational units (OU), in order to restrict the scope of action of an administrator.

What are the possible scenarios for the management of Azure AD?

In order to adhere to the above-mentioned principles, it is necessary to identify teams that each have their own specific responsibilities and rights.

This is not complicated for content management, if the centralized character of the platform is put to one side. For example, the Digital Workplace teams will logically be in charge of the Office 365 services and data. **Container management remains THE central question..**

A central team must be set up with global administrator's rights. This team performs all the administration tasks requiring very high privileges, which may extend beyond the strict scope of identity.

/ This team must be sufficiently **staffed** and **trained**.

/ Processes **with SLAs must be** implemented

Since the central team is not capable of defining all the functionality accessible by this level of administration, making the right teams responsible is a key factor (e.g., communications for the graphical guidelines).

Since these rights are rarely used, it is essential to control who can access them and when (through a bastion or Azure AD PIM).

One scenario could consist of assigning this responsibility to the identity team in order to remain consistent with identity management and data quality. Another could consist of training the central team, so that it acts as the guarantor of the applications on a group-wide scale.

Migrating from Active Directory to Azure Active Directory

In a modernization strategy, Microsoft ultimately recommends reducing Active Directory, as the internal applications and resources gradually become available in the cloud, either as SaaS applications or simply as existing applications.

While the identities will be migrated to Azure AD, the workstations will be taken out of Active Directory and managed from an MDM service in the cloud. This switchover will gradually limit exposure to attacks.

Identity management with a cloud service like Azure AD is easier to understand, because there are fewer concepts to learn and no infrastructure components to be updated.

By migrating to Azure Active Directory, it is easier to benefit from the centralized analysis of connection events, a fact that facilitates the detection of low-volume attacks and connection anomalies.

Why attach a device to Azure AD?

Attaching a device to Azure AD produces several benefits:

- / Reinforcing conditional access

- control on the basis of the state of health of the device or its geographic location;

- / Gaining simple and secure access to cloud applications with SSO by obtaining a Primary Refresh Token (PRT);

- / Managing devices with an MDM solution;

- / Deploying passwordless authentication, like in Windows Hello for Business or FIDO2.

What is Azure AD DS?

Azure Active Directory Domain Services (Azure AD DS) is the AD directory in the form of a cloud service proposed by Microsoft, and by other cloud service providers too. In simple terms, in this case, Microsoft manages tier 0 and the organization manages the other tiers. Azure AD DS provides a sub-set of Active Directory functionality, such as the domain junction, group strategies (GPO), the LDAP protocol and Kerberos/NTLM authentication.

This solution can be used for Windows servers that are migrated to the cloud by lift-and-shift or that are removed from the live AD forest.



Protecting the cloud against a breach of AD

1 Completely isolate the Microsoft 365 and Azure AD administrator accounts

The administrator accounts must be:

- / Created in Azure AD;
- / Authenticated by multi-factor authentication (MFA);
- / Controlled by Azure AD conditional access;
- / Only accessible from workstations managed in Azure;
- / Activated for a limited period of time using Azure AD PIM.

2 None of the Active Directory accounts must have high privileges on the cloud

Make sure that these accounts, including the service accounts, are not included in the privileged roles or groups in the cloud, and that any changes made to these accounts cannot have an impact on your entire cloud environment. The on-premises assets in tier 0 must not be capable of having an impact on the privileged Microsoft 365 accounts.

3 Manage administrators' devices from the cloud

Use Azure AD Join and MDM-type management of devices in the cloud to remove any dependencies on the management infrastructure of on-premises devices that may breach the security measures of the devices used to administer the cloud.

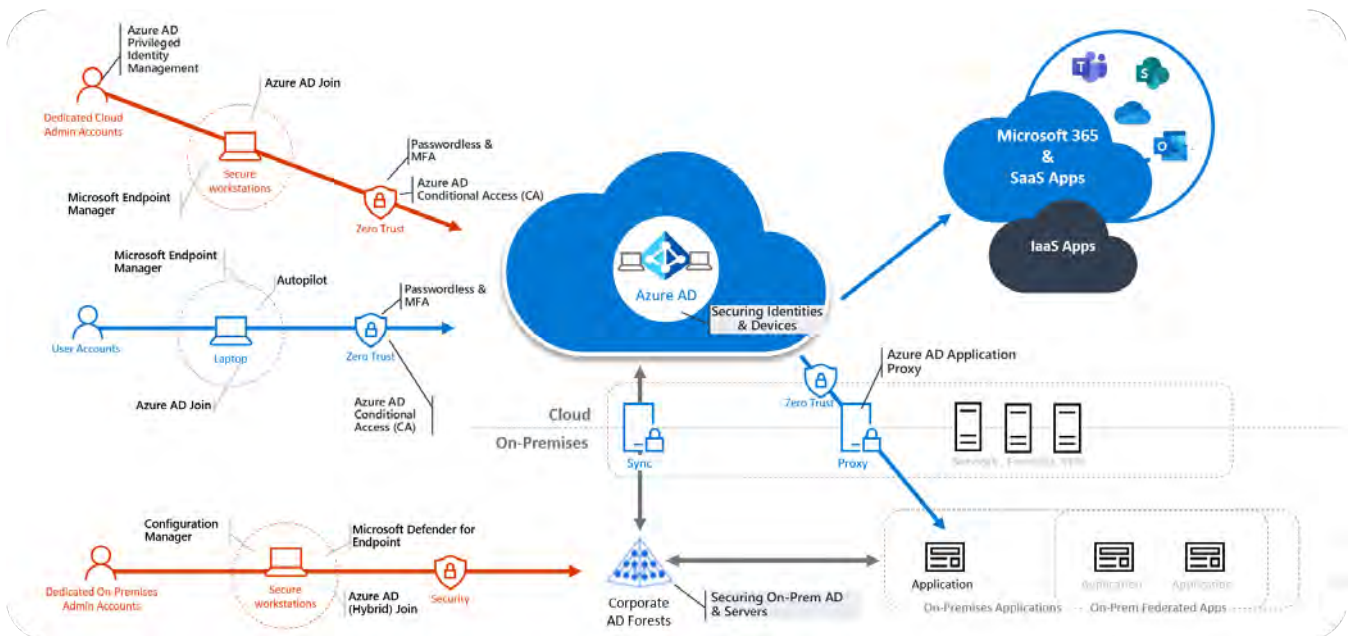
4 Use Azure AD authentication to remove the dependencies on AD

Always use a strong authentication method, such as Windows Hello for Business, FIDO2 or Microsoft Authenticator. Switch to a passwordless authentication method and consider removing the passwords from these accounts.

The journey to Azure AD

When a project to migrate to Azure AD is considered, it is only natural to ask the following question: when will Azure AD be capable of completely replacing Active Directory? Even if it may be tempting to look for a precise date, bear in mind that Azure AD is not “Active Directory in the cloud”, so there is no point in looking for the same functionality.

This transition will probably take several years, both for organizations and for Microsoft. Today, all the technologies required for a 100% cloud deployment of Azure AD are not yet available.



The best solution consists of approaching this journey to Azure AD immediately, and in a pragmatic manner that does not seek to cover every case, in order to avoid falling into the trap of waiting for the perfect solution that covers every eventuality.

An Active Directory contains devices, applications and users. The migration path must take these three elements into consideration, so that they can

gradually be shifted to Azure AD. The speed of the migration path of organizations with a strong legacy with Active Directory will depend on the number of objects to be migrated, but also on the complexity of the AD environment and the number of existing forests.

Modernization based on three pillars

Devices

Put the existing Windows 10 workstations in Hybrid Azure AD Join mode

Natively attach the new Windows 10/Windows 11 workstations to Azure AD using Windows Autopilot

Applications

If possible, migrate application authentication to Azure AD

Migrate the AD FS federation servers to Azure AD

Migrate to Azure AD DS any applications that are too old, are based on NTLM and cannot be migrated to more modern protocols

Users

Implement strong and passwordless authentication for all users

Deploy the new [Azure AD Connect Cloud Sync](#) tool, which simplifies management by centralizing the configuration in Azure AD, facilitates high-availability deployments and grants access to new scenarios, such as support for multi-forest AD environments in disconnected mode

Which path towards modernization?

For organizations that choose to modernize towards the cloud, the transition to Azure AD is gradual and takes place in steps: synchronize the users, migrate the workstations, integrate the applications in Azure AD, then migrate the eligible application servers to the cloud.

The content of Active Directory is gradually reduced and simplified. The respective migrations of tier 2 and tier 1 to Azure AD and Azure AD DS take priority.

The center of gravity shifts towards Azure AD, which becomes the Zero Trust control plane and the environment where the resources are created and then synchronized locally towards Active Directory, if necessary.

Active Directory is gradually isolated and is only used to service critical scenarios and applications that cannot be migrated to the cloud. The creation of tier 0 and the hardening of the Active Directory configuration are the only long-term investments.

In this transformation, the devices are no longer present in

Active Directory, but are directly integrated into Azure Active Directory and are managed by modern, MDM-type solutions.

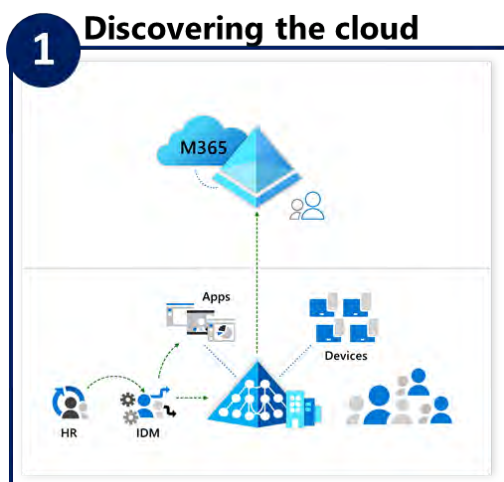
The approach to security becomes more modern by integrating the applications in Azure AD, and in particular by publishing the IaaS and on-premises applications that have not yet migrated to the cloud using Azure AD Application Proxy.

Granting secure access to applications that are accessible from the Internet, without using a Virtual Private Network (VPN), represents a major change for many organizations, even if split-tunneling has become an increasingly common approach.

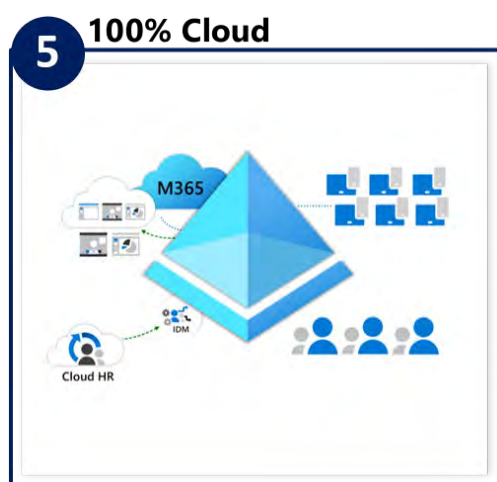
Migrating applications that are highly dependent on Active Directory to Azure AD Domain Services is a solution that should be preferred.

Azure AD becomes the baseline directory that synchronizes identities with third-party repositories, such as a local AD.

Over time, Active Directory will become a lesser cause for concern, due to the reduction in the number of managed assets.



...

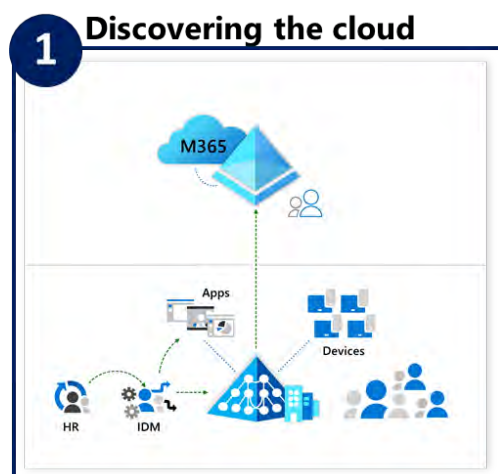


The steps in detail

The progress of a project to move from Active Directory to Azure AD can be measured by the following steps:

1. First steps towards the cloud and the use of Azure AD
2. Universal implementation of the hybrid mode
3. Cloud-centric investment
4. AD is isolated and reduced to a strict minimum
5. 100% cloud - Azure AD

The first steps

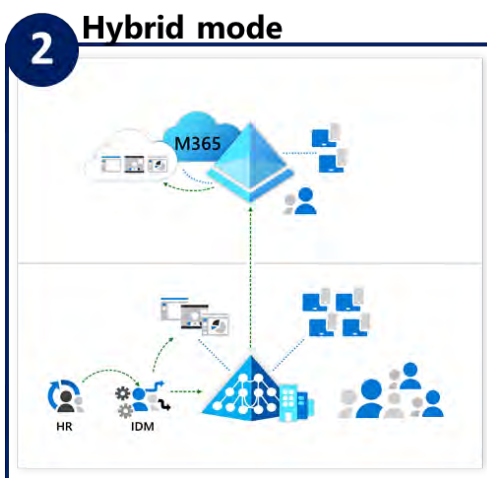


Organizations possess an Azure AD tenant containing at least the user objects, in order to access Office 365 and the Microsoft cloud services in general.

This is the situation of any organization that has started to use the Microsoft cloud: an on-premises Active Directory, devices joined to AD that are configured using group strategies (GPO), on-premises applications that use integrated AD

authentication to control user access, and finally, users who are created in AD with HR systems and are then synchronized from AD to Azure AD using Azure AD Connect.

Hybrid mode



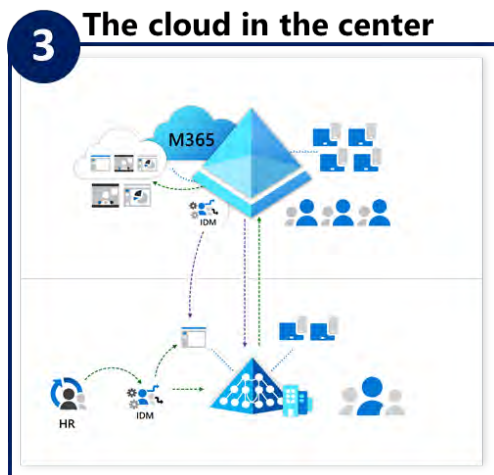
The next step, which numerous customers are currently exploring, consists of becoming hybrid and **using the security services available in the basic version of Azure AD**, but also in the premium versions, such as passwordless authentication, conditional access control, management of privileged identities or self-service password reset for users.

Existing Windows devices are declared in Azure AD using the Hybrid Azure AD Joined mode, which facilitates SSO.

Certain applications are migrated to the cloud on an IaaS infrastructure, by joining Azure AD Domain Services.

Other legacy applications are still hosted on-premises, but they are published for remote workers using Azure AD App Proxy. This grants external access without a VPN, while still protecting access with Azure Active Directory.

Becoming cloud-centric



During this phase, the decision is taken not to add any new devices or applications to Active Directory. Investment in the cloud now takes priority. Active Directory integration projects are wound down in order to prevent the spread of the technical debt.

The migration of tier 2 and tier 1 to Azure AD now takes priority.

In this phase, organizations stop integrating new devices in Active Directory and attach new

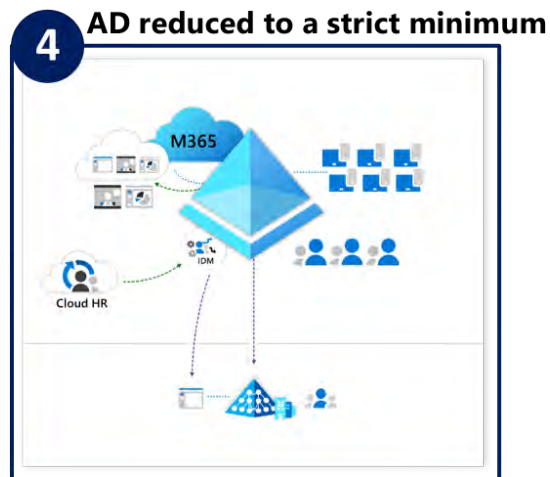
workstations directly to Azure AD using Autopilot, and perform management tasks with an MDM, such as a Microsoft Endpoint Manager. GPOs are no longer used to manage new devices.

The federated applications are gradually migrated to Azure AD. They are hosted in the cloud and use Azure AD for authentication purposes. Legacy applications are gradually published using Application Proxy and, if necessary, are migrated to Azure AD Domain Services. Applications that require user profiles to be synchronized use the [ECMA](#) connector.

Legacy applications use Azure AD DS and are published using Application Proxy.

File shares and printer servers are gradually migrated to the cloud. The Azure Files and Universal Print services are gradually used.

Isolating AD



After this step, Active Directory must be reduced to administering internal resources that will not be migrated to the cloud (e.g., industrial systems) with reinforced workstations, network segmentation and detection systems that guarantee the security of a scope that is now excessively restricted. If it has not already been done, setting up tier 0 and hardening the configuration of Active Directory take priority.

Users are gradually created and

managed in Azure AD only, according to a cloud-first strategy, and they are only synchronized with Active Directory using the write-back functionality, if necessary.

To progress to this step, it is first necessary to modernize the legacy applications, by updating the configuration and the code of the applications, or replacing them with an equivalent cloud version. When Azure AD becomes capable of issuing Kerberos service tickets, it will then be possible to continue to manage legacy applications that are Kerberos-compatible, without any need for user accounts in Active Directory. Moreover, managed AD services, such as Azure AD DS, help to support legacy applications on IaaS servers by offering the AD service in the cloud.



Full Cloud



Finally, step 5 is “100% Cloud Azure AD”, in which the organization no longer has an Active Directory footprint. At this point, there are no more Active Directory domain controllers and Azure AD provides all the identity management tools.

The applications authenticate users with Azure AD using modern protocols or with the support of Kerberos with Azure AD DS and Azure AD. At this point, all the devices are only attached to Azure AD and they are managed with a compatible MDM.

The main subjects to be addressed when switching to 100% Azure AD

OBSTACLE



NTLM



GPO



AD Join

TARGET

Kerberos /
OpenID Connect

MDM strategy

Azure AD Join

COMPLEXITY



Improving the security situation

Did you say “Zero Trust”?

Addressing cyber security from the perspective of identity is one of the underlying trends that organizations are trying to adopt through a Zero Trust approach. What is the principle? **Never trust, always check.**

Every time a login request is made, the context of the access is analyzed to assess the level of trust that can be granted to the user and their device, but also to the applications.

This approach to security protects organizations by granting access on the basis of a continuous verification of identities and the security situation of the devices.

Is this the end of VPNs due to Zero Trust?

One widespread misconception is that switching to a Zero Trust architecture means that any remote access to the enterprise’s resources over a VPN can be discontinued.

But the answer is not so simple. For example, if the workstations use Hybrid Azure AD Join, then these devices must occasionally connect to a domain controller,

because they are attached to Active Directory and Azure AD. Active Directory is the authority that signs the verifier of the authentication secrets on the device in order to open a session in disconnected mode. If the credentials cache is empty or desynchronized, the device must engage with a domain controller in order to add new credentials to the cache. This is possible over a VPN connection or over the organization’s network.

In another example, if a user forgets their password and asks to reset it or resets the password themselves using self-service, their device must be able to reach the Active Directory domain controller in order to use the new password and unlock the computer.

The same requirement for on-premises connectivity applies to the first configuration of Windows Hello for Business on a Hybrid Azure AD Join device. The device requires a connection with a domain controller in order to finalize the configuration of Windows Hello for Business (definition of the PIN code, opening the first session with Windows Hello for Business).

These restrictions do not apply to Azure AD Join devices, which do not need a connection with Active Directory, unlike in the preceding scenarios.



“Always on VPN” technology, which is natively present in Windows 10, is very useful in order to open a VPN tunnel, before the user even opens a session, and to meet this requirement for internal connectivity.

The modernization of VPNs, with a more flexible split-tunneling-type approach, increases the long-term chances of success when implementing a Zero Trust architecture.

Managing passwords

Password spraying is an attack that enables a third party to break into accounts. It consists of testing a few weak passwords on high numbers of user accounts, rather than using numerous common passwords. These attacks are dangerous, because the usual security measures (locking accounts or imposing time delays between successive attempts) are ineffective.

A Microsoft survey of CIOs in several countries who have started their journey to Zero Trust revealed that 76% of them initially implemented strong authentication and that 60% of them have implemented policy-based conditional access.

Azure AD detects password spraying-type attack models by examining failed login attempts for millions of organizations all over the world.

This protection can be extended to Active Directory by installing an agent on the domain controllers and following the instructions in the [deployment guide](#).

The move towards passwordless authentication

We have observed strong enthusiasm for passwordless authentication. This is hardly surprising, given that passwords are responsible for 80% of breaches by hackers and that the deployment of multi-factor authentication is thought to reduce the risk of breach by 99.9%.

Authentication is based on a biometric characteristic, such as a face, a fingerprint or a confidential code specific to a device that is not sent over the network. You can choose between using your Windows PC with biometrics and/or a PIN code, connection using a FIDO2 security code of the Microsoft Authenticator application for mobile devices.

How to combat a cyber attack

The background of the page features a blue gradient with silhouettes of several people holding hands in a circle, symbolizing teamwork and support. A bright light source in the center creates a lens flare effect.

This chapter explains the main difficulties encountered when rebuilding an Active Directory and proposes actions that should be taken in readiness for a cyber attack.

Understanding the difficulties of rebuilding Active Directory in order to better anticipate a crisis

Side effects that are difficult to predict and to resolve very quickly

Following a security breach, there are two methods that can be used to rebuild AD: rebuilding the domain and directory controllers from scratch, or rebuilding the domain controllers from a replication of the existing directory.

Reconstruction following a major cyber incident is often an opportunity to push through security measures, unlike the “small steps” approach that is adopted in standard projects in order to avoid any disruption of business activity.

The first method ensures that any possible means that the hacker can use to remain present will be eliminated, but it demands a significant reconstruction effort and involves a serious interruption of service.

The second method, on the other hand, allows AD to restart more quickly, but without guaranteeing for certain that the directory is intact.

Any operational side effects of a reconstruction will vary

significantly, depending on the environment. When rebuilding Active Directory after an attack, for example a ransomware attack, hardening and remediation actions are taken, frequently during the implementation, without an analysis and test phase. These actions may have operational side effects.

Notable side effects include:

- / breakdown of the Secure Channel between the machines and Active Directory due to the reset / restore of the computer accounts;
- / malfunction of applications using service accounts whose password has been reset;
- / denial of access rights (ACL) on network file shares (in the event of a restore *from scratch* that invalidates the SIDs defined on the resources);
- / the need to migrate obsolete servers that are unable to authenticate themselves because the *legacy* authentication protocols have been deactivated.

Moreover, the absence of an over-arching vision of the entire information system can make the reconstruction

much more complex (digital documents destroyed, obsolete physical copies, partial residual knowledge, etc.). The length of the downtime does not only depend on the complexity of the operations but also on the ability to mobilize enough experts and the priority given to the reconstruction of AD, which involves players who are also required to maintain a possible degraded mode.

“On average, it takes at least one week to rebuild the AD core without any preparations ”

Internal resources are often called on to implement a degraded mode, to the detriment of the reconstruction operations. In addition, the possibility of quickly calling on reinforcements from partner organizations cannot be guaranteed in a crisis. Today, there are only a few operators that possess the levels of expertise required to operationally support AD reconstructions.

Multiple discreet means of persistence

As the IS comes back on stream, it is also necessary to eliminate the hacker's means of persistence in order to prevent the environment from being breached again.

There are numerous techniques of Active Directory persistence that are difficult to verify exhaustively. While some persistence techniques are well known and tools exist to combat them, such as the renewal of the “krbtgt” account, the detection and correction of other techniques can be more complex. For example, this is the case of persistence that uses specific extended permissions and rights, or instances of local persistence on domain controllers.

The CERT-W lists numerous techniques that attackers can use to maintain AD persistence following a breach.

While the complete reconstruction of the Active Directory forest may not be compatible with a quick resumption of activity, it is often the only means of making sure that all the attacker's AD persistence capabilities are eliminated following the breach of a domain.

In addition to the domain controllers and the AD directories themselves, resources and services that domain controllers relies on, such as update servers or PKI, or Tier 1 servers, may also have been compromised by the attacker and used as persistence means.

An action plan must be prepared for these various elements, which form the trust core of the IS, that is adapted to the knowledge of the attack produced by digital investigations.

Can Azure AD be of help during the reconstruction?

Definitely, by using SaaS solutions to restore certain productivity services.

However, Azure AD will not be of any help in the following cases:

/ A large majority of the

applications used in the organization are linked to AD. Authorizations are usually granted to accounts on the basis of the on-premises Active Directory, and the applications do not even know what Azure AD is.

/ The organizations can only natively join workstations (Windows 10/Windows 11) to Azure AD, but not servers.

This is the reason why it is essential to make regular and disconnected backups of Active Directory for greater resilience in the event of a ransomware-type attack.



Preparing the reconstruction of Active Directory

Essential



A backup of AD that is resilient to the selected cyber attack scenarios and is protected

Encrypted Windows backup stored in an unalterable place



Access to the IS without any dependencies on AD

Healthy administration workstation, bypass of the NAC, VPN without AD



A trusted environment for the reconstruction

Independent infrastructure and network



Possibility to call in the right people

Audit, incident response, crisis management and operational teams

Recommended



Have procedures to clean up and rebuild Active Directory

Formalize, automate and practice



Access the applications of the IS without Active Directory

Standalone Azure AD for Office 365, local accounts

To be analyzed



Anticipate the opportunity of relying on Azure AD to rebuild

Workstations in Azure AD Join

Simply rebuilding AD is not enough

The reconstruction of an IS after a cyber attack is a sprint to which everyone must contribute. But there is a danger of assuming that the race stops there.

In this case, it is quite common to be attacked again a few months or years later.

debt that stretches back for years. For example, when managing a crisis for several weeks, it will not be possible to transform the security model or to address all the cases of obsolescence. The narrow path must be set out to restore the service for the business as quickly as possible.

Crisis management must be seen as the first leg of the race that allows the business to reach a milestone in terms of security maturity. It is then necessary to draw up a program to transform the IS in order to pay off the debt and, often, to change the security model, by aligning it with the needs of the business. This is a real marathon!


“80% of the companies that paid a ransom have been victims of a second cyber attack”

Cybereason

Look beyond the crisis to transform the IS

Rebuilding the IS solves the most urgent problems and reaches a first plateau in terms of security. However, the breach of the IS is often a symptom of a security

“A program lasting several years is necessary to transform a security model”



Conclusion

The security of identity repositories is often addressed in terms of numerous technical details that are discussed between experts.

However, it is possible to maintain a pragmatic approach by asking the right questions: How is Active Directory administered? From which workstations? Since the services have strong relationships with these servers, are these workstations really secure? The vision must be broadened to include the security of Azure AD, which is not just a simple extension of Active Directory into the cloud.

It is necessary to define a short- and medium-term target that is consistent with the transformation strategy of the organization.

Security is a question of arbitration. It is essential to know which vulnerabilities expose the organization, while also understanding the associated risks and benefits. This will help to take the right decision on the path to modernization and the priorities.

The administration model evolves and takes the new dimensions of the extended enterprise into account through these hybrid and multi-cloud aspects.

Throughout this transformation, organizations must remain focused on the essentials, with identity as the cornerstone of information system security.

The identity system of organizations is now hybrid and this new reality must be embraced.

Acknowledgments

AUTHORS



THOMAS DIOT
Senior Consultant,
Wavestone



THIBAUT JOUBERT
Manager,
Wavestone



ARNAUD JUMELET
National Security
Officer, Microsoft France



ÉTIENNE LAFORE
Senior Manager,
Wavestone



ALEXANDRE LUKAT
Manager,
Wavestone

CONTRIBUTORS

PIERRE AUDONNET
Principal Customer Engineer,
Microsoft Canada

FLORENT BENOIT
Partner Technology Strategist,
Microsoft France

RÉMI ESCOURROU
Manager, Wavestone

JEAN-YVES GRASSET
Chief Security Advisor, Microsoft
France

BENOÎT MARION
Senior Manager, Wavestone

GREGORY SCHIRO
Compromise Recovery Security
Practice, Microsoft

JULIEN ROUSSON
Manager, Wavestone

Useful links

ANSSI - ACTIVE DIRECTORY CONTROL PATHS

<https://github.com/ANSSI-FR/AD-control-paths>

ANSSI - THE ACTIVE DIRECTORY SECURITY (ADS) SERVICE

<https://www.ssi.gouv.fr/administration/actualite/le-service-active-directory-security-ads-accompagner-la-securisation-des-annuaires-active-directory-des-acteurs-critiques/>

ANSSI - ACTIVE DIRECTORY CHECKPOINTS

<https://www.cert.ssi.gouv.fr/uploads/guide-ad.html>

ANSSI - SECURITY RECOMMENDATIONS FOR ACTIVE DIRECTORY

<https://www.ssi.gouv.fr/guide/recommandations-de-securite-relatives-a-active-directory>

M365INTERNALS - INCIDENT RESPONSE IN A MICROSOFT CLOUD ENVIRONMENT (HUY KHA)

<https://m365internals.com/2021/04/17/incident-response-in-a-microsoft-Cloud-environment/>

MICROSOFT - APPENDIX L: EVENTS TO MONITOR

<https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/appendix-l-events-to-monitor>

MICROSOFT - AZURE ACTIVE DIRECTORY SECURITY OPERATIONS GUIDE

<https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/security-operations-introduction>

MICROSOFT - BEST PRACTICES FOR AZURE AD ROLES

<https://docs.microsoft.com/en-us/azure/active-directory/roles/best-practices>

MICROSOFT – DETAILS OF AZURE AD LICENSES

<https://www.microsoft.com/en-us/security/business/identity-access-management/azure-ad-pricing>

MICROSOFT - ENTERPRISE ACCESS MODEL

<https://docs.microsoft.com/en-us/security/compass/privileged-access-access-model>

MICROSOFT - SECURING AZURE ENVIRONMENTS WITH AZURE ACTIVE DIRECTORY

<https://aka.ms/AzureADSecuredAzure>

MICROSOFT – APPLICATION OBJECTS AND SERVICE PRINCIPALS

<https://docs.microsoft.com/en-us/azure/active-directory/develop/app-objects-and-service-principals>

MICROSOFT - AZURE DEFENSES FOR RANSOMWARE ATTACK

<https://azure.microsoft.com/en-us/resources/azure-defenses-for-ransomware-attack/>

MICROSOFT – WHAT IS THE IDENTITY SECURE SCORE?

<https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/identity-secure-score>

MICROSOFT – SECURING AZURE SERVICE ACCOUNTS

<https://docs.microsoft.com/en-us/azure/active-directory/fundamentals/service-accounts-introduction-azure>

MICROSOFT – DOCUMENTATION ON EXTERNAL IDENTITIES

<https://docs.microsoft.com/fr-fr/azure/active-directory/external-identities/>

MICROSOFT – MICROSOFT AZURE AD ASSESSMENT

<https://github.com/AzureAD/AzureADAssessment>





Microsoft is committed to trusted, inclusive and sustainable digital technology. Its mission consists of giving every individual and every organization the means to achieve their ambitions in the era of the smart cloud and the intelligent edge.

A catalyst of innovation in France for almost 40 years, Microsoft France has been chaired by Corine de Bilbao since July 2021. With more than 1,800 employees and 10,500 economic and technological partners, players in the public sector, researchers or start-ups, Microsoft France contributes to the development of the economy and digital skills all over the country.



In a world where the ability to transform oneself is the key to success, Wavestone has set itself the mission of informing and guiding major corporations and organizations in their most critical transformations, with a view to making them positive for all the stakeholders. This what we call “The Positive Way”.

As one of the leading independent consulting firms in Europe, Wavestone employs more than 3,000 people in eight countries, including more than 600 cyber security consultants. These consultants help organizations to address all the issues related to cyber security, from the most strategically important, to operational implementation, incident response and digital investigations.

Wavestone is listed on Euronext in Paris.

More information at www.wavestone.com/fr/
@Wavestone_
@RiskInsight





www.microsoft.com



www.wavestone.com