A SECURITY STRATEGY AGAINST STEAL-AND-PASS CREDENTIAL ATTACKS

Juan Ceballos
Security Consulting & Engineering, T-Systems International, Munich, Germany

ABSTRACT

Stealing and passing credentials is currently one of the preferred cyberattack techniques within the hacking community as shown by the increasing number of related incidents over the last years. Instead of targeting passwords, attackers focus on obtaining derived credentials like hashes and session tickets. This type of credentials facilitates taking advantage of omnipresent background mechanisms like Single Sign-On. A combination of malware and penetration tools is used in order to exploit architecture vulnerabilities and steal the credentials. Vulnerabilities also allow the attacker to get access to other systems and covertly take the control of central infrastructure like Active Directory. The ultimate goal is not creating damage that can be noticed but covertly and constantly leaking confidential information for profit or cyber spionage.
This paper proposes a comprehensive strategy of six points against steal-and-pass credential attacks and is intended to mitigate the risk significantly. Even if some points of the strategy can be considered security best-practices, other points require the establishment of technical and process controls that are not part of typical security management programs. Controls have to be regularly reviewed as part of security audits, since administrators and other privileged users have often the means to remove or bypass technical controls.

KEYWORDS

Security, Cyberattack, Hacking, Malware, Security Threats & Countermeasures

1. INTRODUCTION

Stealing and passing credentials is currently one of the most preferred cyberattack techniques within the hacking community when targeting enterprises with a Microsoft Active Directory infrastructure. Even if there is not an official statistic about how many companies have been hacked using this technique, the number of related incidents has significantly increased over the last years [1] and may continue to grow, since there is no definitive solution against it.

Some years ago the use of cracking tools and social engineering were the preferred methods of hacking to obtain credentials such as passwords. However, users are nowadays more conscious about choosing better passwords, which makes cracking them more difficult. Users are also more aware against clicking unknown links requesting them to log in and are more cautious about providing information to phone callers. In contrast, methods like Pass-the-Hash (PtH) and Pass-the-Ticket (PtT) are targeted to obtain other types of credentials like the hashes behind the password or access granting tickets. These types of credentials, which are as powerful as the password itself, are used by the operating system or applications to authenticate the user in the background. Since operating systems often keep hashes and tickets locally in memory or in a credential store, it is in theory as simple as getting access to such credentials and using them to authenticate and impersonate the user.

References [4], [5] and [7] provide sound reasons why hashes and tickets should be considered not secure. Most of arguments are not related to software flaws but to Microsoft's internal architecture and features:

1. A hash is a non-reversible set of data that is cryptographically tied to another set of data (e.g. password) from which the hash is calculated. If the set of data is changed, the hash also changes. Microsoft's hashing implementation stores passwords in two different ways, by default: as the LAN Manager one-way function (LM OWF) and as the NT one-way function (NT OWF). Neither the LM hash nor the NT hash is salted [9]. Salting is a process that combines the password with a random numeric value called the salt before computing the one-way function. Not salting the hash, as it is the case in Microsoft implementation, has the implication that the hash is the same on the client side and on the server side. This situation makes the hash susceptible to reutilization by an attacker [5].

2. Regardless of implicit security restrictions, Kerberos granting and service tickets can also be used for steal-and-pass credential attacks. Well know forms of exploits are the so called Golden and Silver tickets [4]. As explained in references [2] and [3], Kerberos relies on a central secret key. If the server that stores the secret key is hacked, it may be possible for an attacker to generate Kerberos credentials and impersonate other users everywhere in the network. This is called "The Golden Ticket" (TGT). Another variant of this attack, which is supposed to have been fixed by Microsoft since November 2014, are the "Silver Tickets" [7]. A Silver Ticket is a forged Kerberos Service Ticket (ST) valid only for a specific server. Some Linux and Unix versions are also likely to be impacted by this vulnerability.

3. One of the most convenient features in Windows is Single Sign-On (SSO). After an initial authentication SSO allows users to do things like access network file shares, connect to an Exchange server, a Print-Server or a SharePoint without having to authenticate again [5]. In order to do so, the hash or the token is stored locally and then passed along each time a service needs to authenticate. Unfortunately, if the operating system is able to access the credentials, then malware with sufficient rights can do this as well [4].

Even if Microsoft has done some architectural improvements as part of Windows Server 2012 R2, Windows 8.1 and Windows 10; a complete change of the described hash, token and SSO behaviour would require a complete re-design in Windows architecture, which is unlikely going to happen in the near future. However, and until a definitive solution for this problem is available, IT policies and security strategy have to change in order to meet these threats.

2. ATTACK VECTORS

Steal-and-pass credential attacks can be considered Advanced Persistent Threats (APT). An APT is a covert, planned and long term hacking activity targeting organizations for business or political motives. APTs are called advanced, since a complex combination of techniques and tools is used in order to exploit vulnerabilities in systems or to create new ones. Persistency refers to the ability of the attackers to continuously monitor the target looking for opportunities to expand their control and leak data without being discovered.

Typical attack vectors are as follows [8]:

- **Step 1:** An attacker attempts to compromise a Windows client PC with malware. The sources of such malware could be for instance opening an infected mail attachment, clicking on an Internet link or plugging a device in a USB port. In case that the malware is successfully installed, only limited access to the system resources is normally acomplished, since most

organizations don't provide users with administrative rights to desktop users. In order to be able to perform the next steps, the attacker has to communicate with the malware via a wired (e.g. Internet) or a wireless network (e.g. Wi-Fi) in order to gain control of the malware. In case that the user has administrative rights on his PC, the attacker can go directly to step 3 or even to step 5.

- **Step 2:** From the infected PC, the attacker tries to infect other PCs in the same network, for instance using a compromised E-Mail account or an existing vulnerability. This is called lateral movement. With some luck one of the users of the newly infected PCs may have local or domain administrative rights and the attacker can install his tools with higher privileges and jump directly to step 5.

- **Step 3:** After collecting enough information about the user and his system, the attacker tries to exploit other vulnerabilities in order to get higher local rights. When this is achieved tools like Mimikatz, Windows Credential Editor (WCE), gsecdump and Metaexploit can be used. These tools give an attacker the capability to capture the hashes from the credential storage or read passwords from the LSASS (Local Security Authority Subsystem Service) process through a code injection. Since the source code for interfacing with the process is in many cases freely available on the Internet, the tools can be recompiled and not identified by the antivirus and other antimalware software. In case that privileged hashes or tickets are already stored in the compromised system, the attacker can go directly to step 5.

- **Step 4:** The attacker creates an error condition or a fake critical message in a compromised and prepared PC so that the user has to call the help desk. A supporter with domain rights logs into the PC locally or remotely in order to troubleshoot the issue. With the help of the previously installed tools, the attacker may be able to compromise the supporter account.

- **Step 5:** With the already compromised administrator or supporter account, the attacker starts to take over other systems. With some luck and patience, the attacker may be able to capture the credentials of a privileged service account or a Domain Administrator and get into more critical systems like servers or even domain controllers. At this point, the complete Active Directory infrastructure and all subordinated systems be considered as compromised. Central Microsoft infrastructure can also be used as launch point of attacks to systems with Windows Interfaces but running on other operating systems like Linux and Unix.

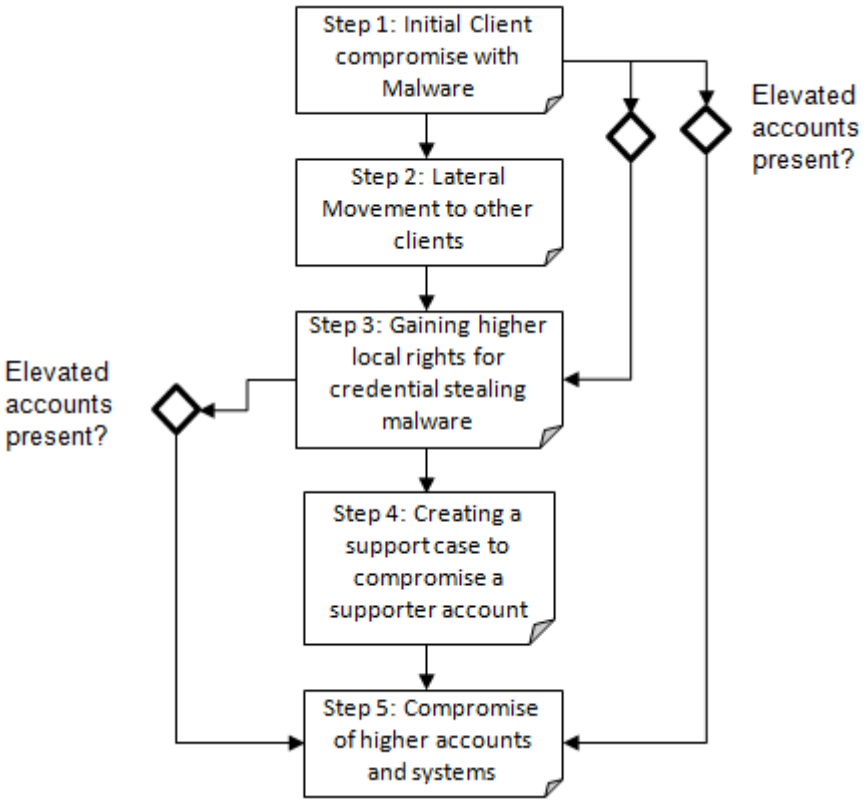

Figure 1. Typical attack vector

Once Active Directory has been compromised, there are not many options for recovery. Without good logs and forensics is nevertheless hard to determine the exact timeline of the attack and the path followed by the attackers. If the attack time and systems infected are precisely known, a recovery with backups to a point before the attack may be feasible. In the case that the incident investigation shows that the attack point was not recent; a recovery may not be practical due to outdated backups leaving often only two possibilities: a new setup or a cleanup of the Active Directory. Both alternatives are expensive, take time, require knowledge normally not widely available in most organizations and have significant impact to normal IT operations.

3. IMPLEMENTING A SECURITY STRATEGY AGAINST THE ATTACK

Documents from Microsoft [2], the Information Assurance Directorate (NSA/CSS) [5] and other sources propose a long list of mitigations and controls against this type of attack. However, it is not always easy to recognize what is really priority against the attack vector from a IT budget perspective. The cost of new security controls must have a clear benefits case, show the operational feasibility of the controls and consider potential user acceptance issues. In contrary to other types of attacks, and due to the nature of the vector involving different steps and the architectural gaps of Microsoft Windows, it is just simply not possible to patch against steal-and-pass techniques in order to fix the problem. This special situation may not be easy to explain to the decision makers approving security costs. For this reason, we have created a six-point comprehensive strategy focused on preventing and containing the attack vector. The controls outlined around each strategy point provide specific countermeasures against each step of the attack, are operationally feasible, acceptable to users and can be cost-effectively implemented, in most cases with existing security assets and resources.

3.1. Control potential malware entry points

Probably the most critical part of the attack vector is establishing countermeasures against step 1. Controls for this step are very generic and are likely to be already in place in most organizations with a security management program. A typical example of such practices is providing regular employee security awareness training. Awareness training should include information about the harm and potential consequences of visiting certain websites and opening links and emails of unknown sources, as well as plugging in unsolicited gifts to USB ports. Since it is not possible to guarantee that all users follow these practices, additional technical controls like web filters, malware scanners and USB firewalls significantly reduce the possibility of malware ingress through these channels. In case that a malicious file still lands on a machine via mail, the second line of defense is avoiding its installation. Traditional antivirus software remains the best protection against malware installation. A two vendor anti-malware strategy, for instance one product for desktops and different one for mail servers and gateways, may increase the probability to stop certain infections, since signatures and detection algorithms may differ. If malware gets installed anyway, i.e. when both antivirus products have no signatures for it, this doesn't automatically mean that it can be used for an attack. Regular security updates and critical security patches remains the last line of defense against the exploitation of known vulnerabilities.

3.2. Reduce the attack surface with system hardening

Best-practice system hardening can significantly help to avoid lateral movements as in steps 2 and 5 and also to reduce the attack surface against malware in steps 1 and 3. Excellent recommendations and baseline tools are provided for instance by Microsoft and the Center of Internet Security (CIS):

• The Security Compliance Manager (SCM) is a free tool from Microsoft that enables to check and manage system configurations according to the machine's assigned role. The configuration templates are based on Microsoft security recommendations. These are available for Windows-based operating systems only.

• A more independent approach is provided by the CIS Benchmarks, which are consensus-based, best-practice security configuration guides and tools developed and generally accepted by government, business, industry, and universities. There are also CIS Benchmarks for non-Windows operating systems, databases and common enterprise applications.

Other tools and features provided by Microsoft as part of their Operating Systems like Personal Firewalls, AppLocker and EMET (Enhanced Mitigation Experience Toolkit) can also be implemented as part of Windows Group Policy Objects (GPO):

• Host-based firewalls like the Windows-Firewall included by default as part of the Operation System can be configured to block all incoming and outgoing connections except those explicitly required for normal business purposes.

• AppLocker complements standard antivirus software by allowing administrators to create security whitelists based on file locations and signatures. AppLocker whitelists control which files users and processes are allowed to run or call including executables, scripts, Windows Installers and Dynamic Link Libraries (DLLs).

• EMET allows enabling or adjusting advanced security parameters individually for the execution of applications and processes. These settings are normally not directly accessible to

administrators and difficult to configure by hand. EMET may be able detect and prevent some common exploitation techniques employed by certain types of malware.

When properly configured, these tools and features will significantly reduce the possibilities for installation and execution of malicious software and lateral movement. Centrally managed GPOs and policy monitoring defend policies from bypass when a system is locally compromised.
In addition, newer Windows versions starting from Windows 8.1 and Windows Server 2012 R2 allow the execution of LSASS as a protected process and limits potential interactions with normal processes like penetration tools [9]. With this feature a LSASS code injection is much more difficult. Only code signed by Microsoft can be executed as a protected process regardless of administrative or debug rights.

3.3 Implement administration layers based on system purpose and criticality

The implementation of administration layers provides security isolation to systems of different purpose and criticality. An effective separation of systems and personnel in layers depends on a novel architectural paradigm proposed by Microsoft called the three-tier administration model. This model can be enhanced with additional countermeasures against steal-and-pass credential attacks.

3.3.1. Microsoft's three-tier administration model

Microsoft introduces in [2] a three-tier administration approach as adaptation of Biba and Bell-LaPadula hierarchical models with the objective of preventing escalation of privileges with stolen credentials. In this document Microsoft defines three administration layers called tiers. A tier is defined by the administrative roles accessing each type of system:

• **Tier 0:** Privileged administrators with control of Active Directory, domains, domain controllers and other central systems (e.g. Certificate Authority).

• **Tier 1:** Server administrators with control over a single or multiple servers and applications.

• **Tier 2**: Client administrators with control over a single or multiple end user devices (laptops and desktops).

The aim of the tier model is that a system, and in particular trusted systems like domain controllers, should never be accessed from a workstation that is not at the same level as the managed system.

Before implementing a layered administration model a classification of systems by the administration accounts accessing them has to be done. Even if labeling higher systems like domain controllers as Tier 0 is straightforward, there are other systems more difficult to classify. For instance, systems that at a first look may seem to belong to a lower tier but after a detailed examination contain for example service accounts or other objects being executed at a higher tier. In this case, the system tier should be classified with the level of the highest account accessing it. For example, if an application server runs a service with Domain Administrator rights, it has to be consequently classified also as a Tier 0 system. The general rule to be enforced after grouping the systems is that accounts controlling higher tiers are not allowed to log into lower tier systems (see Figure 2). This is done to avoid that administrator or service account credentials of a higher tier are stolen in case that the lower tier is compromised.

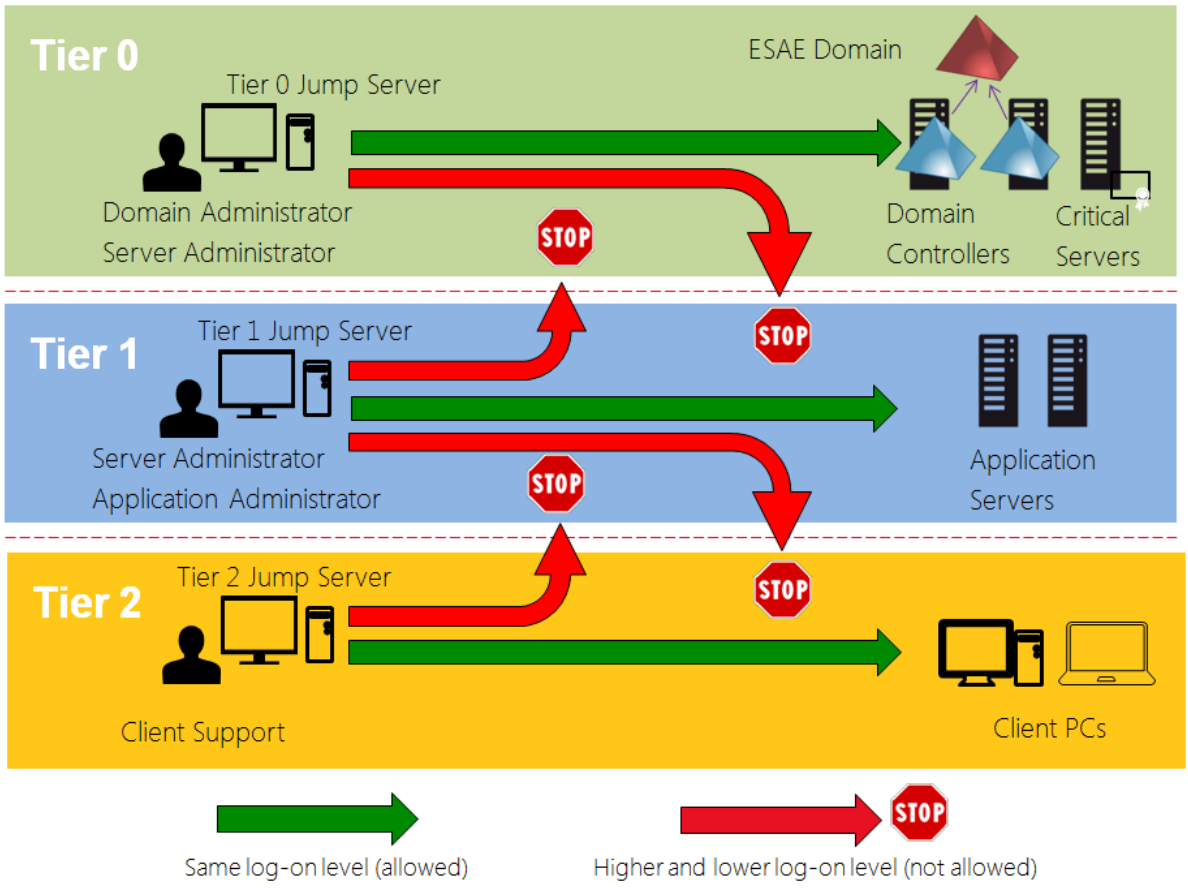

Figure 2. Modified Microsoft's administrative three-tier Model [2]

While this model is probably the key for preventing privileged credential theft in steps 4 and 5, Microsoft does not provide much detail about its practical implementation and in particular a methodology to perform system classification.

3.3.2. Implementation of additional restrictions

Microsoft's tier model can be further enhanced with a stronger restriction of not allowing logons to higher tier systems from lower tier systems in order to prevent that stolen accounts are used (see Figure 2). This restriction is also meaningful, when considering the possibility that a compromised lower tier system could also infect with malware a higher tier system.

3.3.3. Implementation of jump servers for administration

Another practical adaptation to the tier model is the utilization of Terminal Servers as administrative jump servers instead of using dedicated administration workstations as suggested by Microsoft (see Figure 2). Jump servers are hardened Terminal Servers running the Remote Desktop server role and providing virtual sessions to the administrators of each tier. Jump servers are not only meaningful from a security point of view but also for cost and organizational reasons, since many companies don't have dedicated administration workstations. In fact, many daily tasks of an administrator, including searching for technical information in the Internet, require only a normal PC and a user without administrative privileges. Performing this kind of activities with administrative accounts or from an administration workstation increases the risk of step 1 and a shortcut to step 5 (see Figure 1).

3.3.4. Implementation of an Enhanced Security Administrative Environment (ESAE)

In the more general context of Consulting Services, Microsoft introduces the ESAE concept to protect highly privileged accounts. Even if the details behind the concept are not openly available, one of the ideas that can be derived of Microsoft's ESAE flyer [6] is the creation of a new domain for forest-wide administration. This can be done by establishing an additional domain only for administrators, which is trusted by other domains in the forest (1-way trust). The goal of this approach is that the daily administration of the production domains is not performed anymore with production domain accounts but using the external ESAE forest accounts. However, administrative domain accounts should still exist for recovery purposes and for troubleshooting connectivity issues with the ESAE domain. Accessing the ESAE domain with production domain accounts is not possible due to the 1-way trust.

3.3.5. Implementation of smartcards for administration

ESAE accounts and hashes can be protected against credential passing attacks by using a combination of two-factor authentication (i.e. smartcard and PIN) and Tier 0 jump servers. A domain administrator can only access a target system in a managed domain (e.g. a domain controller) after a successful smartcard authentication with a non-privileged RDP account via the confined environment of each jump server (see Figure 3). The ESAE roles and accounts necessary to manage other Tier 0 systems are also contained in the administrator's smartcard. The combination of enforcing authentication with smartcards for administrators and a restrictively configured environment of jump servers decouples the ESAE domain of the managed domains limiting hash and token passing authentication mechanisms.

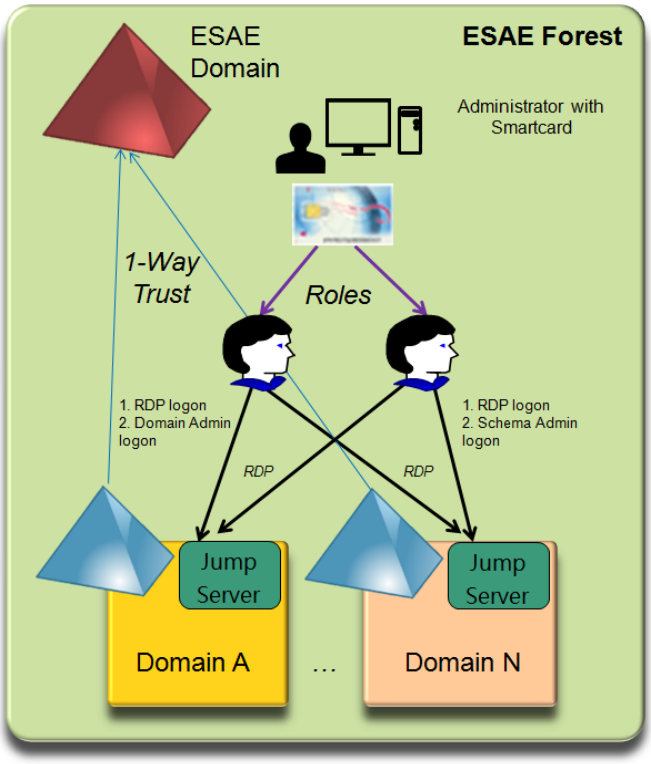

Figure 3. Domain Administration with ESAE Domain and smartcards

However, it is important to emphasize that passwords cannot be completely disabled and substituted by smartcards in a Windows environment. Even if smartcards have to be used by default, giving an administrator the exceptional possibility of using passwords remains important

for disaster recovery and troubleshooting purposes. For this reason, password complexity and change period policies for both privileged and normal accounts remain necessary. A password change also regenerates the hashes and may also help to delay an ongoing passing credential attack in the case that any hashes have been already compromised.

3.3.6. Implementation of roles with strict separation of duties based on tiers

Also controlling the number of privileged accounts and their assigned rights remains important in spite of ESAE restrictions. The more accounts with administrative access, the higher the probability of a compromise. Therefore, maintaining a policy for granting and auditing rights following the least-privilege principle is essential. A simplified example of such a privilege structure grouped by tier is shown in Figure 4.

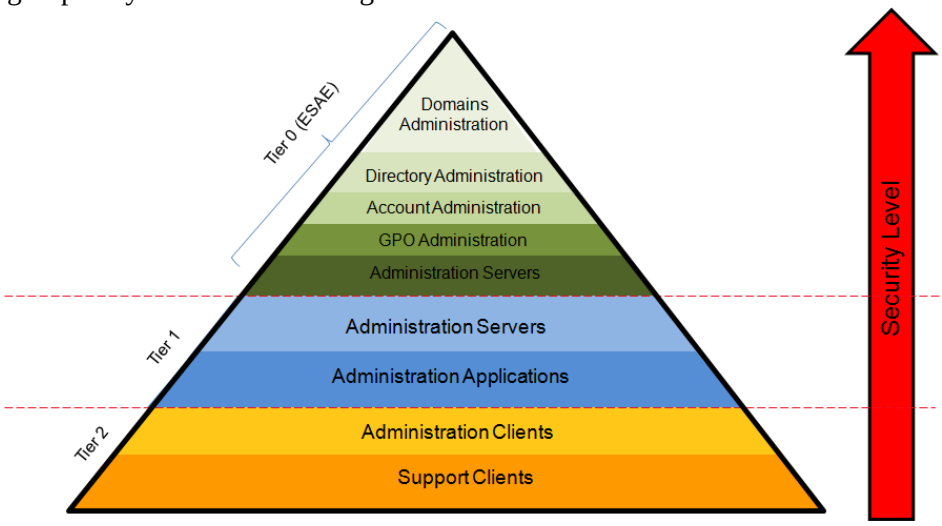

Figure 4. Role groupings following the tier model

Even if some administrative tasks could be still delegated to other administrator roles for convenience, the tier-boundaries for administrative tasks must always be enforced.

3.4. Segregate the network in Security Zones based on system purpose and criticality

In order to avoid lateral movements between systems of different tiers as described in step 5, it is fundamental to segment the network in at least four different security zones separated by network firewalls and routers (see Figure 5).

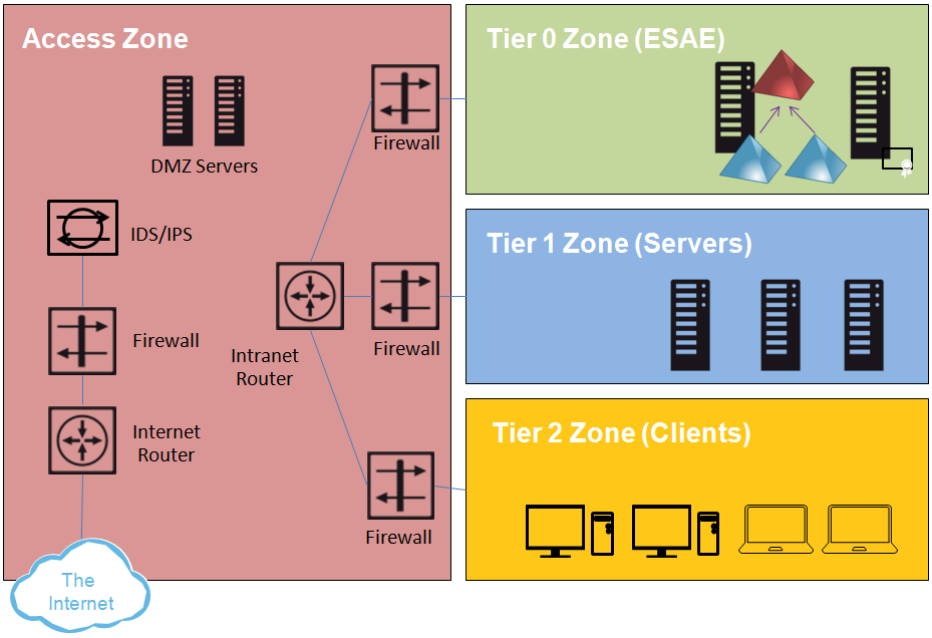

Figure 5. Segregation in Security Zones

• **Access Zone:** This zone contains systems for interconnection to external networks (e.g. the Internet) and hosts also the DMZ (e.g. for mail, VPN and public web servers). With respect to network protection, this zone contains security perimeter systems like firewalls, proxies, IDS/IPS, etc, which act as perimeter defense against malware and attackers.

• **Tier 2 Zone (Clients):** This zone contains end user desktops and laptops.

• **Tier 1 Zone (Application Servers):** This zone is intended for protecting servers for business applications. Sub-zones for systems with different data confidentiality requirements can be added as needed.

• **Tier 0 Zone (ESAE):** This zone is intended to protect critical assets like Domain Controllers, CA Servers, etc. It also contains the ESAE Domain infrastructure used for central administration. The idea behind security zones is not only to segregate the network but also to enforce with the firewalls traffic restrictions between systems of different tiers only to essential communication flows. Interactive traffic between zones for administration is prohibited. Privilege escalation should only be possible using the jump servers within each zone as entry point, as shown in the Modified Microsoft administrative Tier Model (see Figure 2).

3.5 Secure end user support

A key goal of steal-and-pass credential attacks is to gain any kind of domain-wide administrative privileges. This typically happens in step 4 after the attacker manages to create an error message serious enough that the user has no other option than calling the service desk for support. The best strategy against avoiding this happening is by creating the awareness in the support team that client PCs can always be infected with malware. Doing so, the supporter can troubleshoot a system having always in mind that a security breach could also be the potential cause of the failure. In case that there are signs that a failure was intentionally caused the system must be immediately isolated.

A second strategy is to prohibit direct logons from supporter's workstations to client PCs and also implement jump servers with smartcard authentication for support tasks. This approach also

facilitates remote elevation of rights with the smartcard instead of using “run as” on the PC. The jump server also acts as an internal barrier in this zone against the spreading of malware and credential theft. In addition, Windows Server 2012 R2, Windows 8.1 and versions above provide new features like the option “restrictedadmin” for remote administration. With this option encrypted authentication is enabled and the accessed machine will not cache supporter credentials.

3.6 Establish capabilities for Monitoring, Detection and fast Reaction

Establishing targeted monitoring, detection and fast reaction capabilities rounds the strategy against this type of attacks. Even if it is very difficult to detect credential passing techniques in a network, since most events look like typical Windows authentication workflows, it is possible to monitor and correlate some events that could be an indicator of a potential attack. References [2] and [5] provide a list of messages and processes that could be monitored for detection. Newer Windows versions starting from Windows 8.1 and Windows Server 2012 R2 offer new events for monitoring. Particularly useful are LSASS process messages. Another possibility for monitoring is creating the so called honey-hashes and honey-tokens (as an analogy to honeypots) in potentially interesting systems for an attacker. Honey-hashes and honey-tokens are fake credentials stored in the same location as real hashes and tokens and just waiting to lure an attacker. Their use by anyone in any circumstances should raise an immediate alarm, since no legitimate user or process should interact with these credentials. In addition, the implementation of the ESAE domain for administration and the Tier Model provides also the possibility to monitor non-ESAE domain accounts that should only be used for disaster recovery or exceptional troubleshooting.

The implementation of more complex monitoring scenarios may reduce the number of false positives and effort caused by having to verify every single suspicious event. Solutions like a SIEM (Security Information and Event Management), which is the deluxe variant to simple log monitoring can be used to centrally collect logs from heterogeneous relevant sources like operating systems, network devices and security devices. They correlate different events and raise alarms upon detecting complex event sequences. Implementing an Intrusion Detection Systems (IDS) in the network and feeding its data to the SIEM not only improves malware detection capabilities but also provides visibility to irregular network patterns.

As soon as it is suspected that an attack is in progress, there is generally little time to react. In the case that administrative accounts have been compromised, it may take the attackers just a couple of hours to gain control of critical systems. For this reason, it is crucial to develop in IT security and administration teams the necessary skills to identify and respond to steal-and-pass credential attacks in their early steps. The response team should be able to determine as early as possible, which systems are compromised and when exactly this happened. This capability helps to promptly isolate and clean up the affected objects with minimum impact to the organization.

4. CONCLUSIONS

A steal-and-pass credential attack is a current APT that has to be taken seriously by IT departments managing Active Directory infrastructures, especially when considering the cost of a full system recovery after a serious compromise. Where concrete countermeasures are not part of the architecture and IT processes, an attacker can after few steps take control of the complete infrastructure without anybody noticing it. The goal of most attackers, who may belong to organized criminal groups or are cyber spying agents, is not to create damages that can be noticed but covertly and constantly leaking valuable business data, intellectual property or even classified information.

The following points provide a comprehensive strategy against steal-and-pass credential attacks and help to mitigate the risk significantly:

1. Control potential malware entry points
2. Reduce the attack surface with system hardening
3. Implement administration layers based on system purpose and criticality
4. Segregate the network in Security Zones based on system purpose and criticality
5. Secure end user support
6. Establish capabilities for Monitoring, Detection and fast Reaction

A successful implementation of this strategy requires a change of mindset in the IT administration team, who may have to modify old administration habits from the time when administrators where allowed to do everything. Creating awareness about the existence of the attack and its potential consequences helps users to better accept constraints due to the implementation of the strategy. In particular IT administrators should understand the risks of losing sensitive information and the huge effort behind a complete infrastructure recovery in the case of a successful compromise. Finally, security auditing is a fundamental part of this strategy, since administrators and other privileged users often have the means to remove or bypass technical controls.

ACKNOWLEDGEMENTS

The author would like to thank Rory Brennand and Christian Stengel of T-Systems International for reviewing this article.

REFERENCES

[1] Verizon. Data Breach Investigation Report (2014). Online at http://www.verizonenterprise.com/de/DBIR/ (accessed 17/01/2016).
[2] Microsoft. Mitigating Pass-the-Hash (PtH) Attacks and Other Credential Theft Techniques (2014). Online at https://www.microsoft.com/en-us/download/details.aspx?id=36036 (accessed 11/07/2015).
[3] Microsoft. Kerberos Authentication Technical Reference. Online at http://technet.microsoft.com/en-us/library/cc739058(v=ws.10).aspx (accessed 17/01/2016).
[4] CERT-EU (2014). Security White Paper 2014-07 Pass The Golden Ticket v1.1 1. Online at http://cert.europa.eu/static/WhitePapers/CERT-EU-SWP_14_07_PassTheGolden_Ticket_v1_1.pdf (accessed 09/07/2015).
[5] Information Assurance Directorate (2013). Reducing the effectiveness of Pass-the-Hash.Online at https://www.nsa.gov/ia/_files/app/Reducing_the_Effectiveness_of_Pass-the-Hash.pdf (accessed 09/02/2015).
[6] Microsoft. Enhanced Security Administrative Environment. Online at http://download.microsoft.com/download/A/C/5/AC5D21A6-E04B-4DC4-B1F2-AE060319A4D7/Premier_Support_for_Security/Popis/Enhanced-Security-Admin-Environment-Solution-Datasheet-[EN].pdf (accessed 17/01/2016).
[7] No, Microsoft Hasn't "Fixed" Silver Tickets. Passing-the-Hash blogspot. Online at http://passing-the-hash.blogspot.de/ (accessed 17/01/2016).
[8] Schweizerische Eidgenossenschaft Informatiksteuerungsorgan Bund ISB (2013). Technologiebetrachtung „Pass the Hash"-Angriffe. pp2-3.
[9] Coates , Andrew and Sanders, Stephanie (2014). PasstheHash Defense: Analysis of Strategies to Mitigate Weaknesses in Microsoft NTLM Authentication. Online at http://insurehub.org/sites/default/files/reports/passhash-report.pdf (accessed 18/01/2016).

AUTHORS

JUAN CEBALLOS is Senior Consultant for IT and Security Architecture at T-Systems International. He is based in Germany. Mr. Ceballos graduated in Electronics Engineering from the Universidad Pontificia Bolivariana and is Software Development specialist from the Universidad de los Andes in Colombia. He received a M.Sc. in Computer Science from the University Erlangen-Nuernberg in Germany and a MBA from the University of Warwick in England. He is a Certified Information Security Professional (CISSP) and holds other certifications, including some of Microsoft.