White paper

Pass-the-Hash and Pass-the-Ticket

Is the Software Defined Perimeter really Zero Trust?

Pass-the-hash and pass-the-ticket are commonly used attacks which many traditional security products (i.e. firewalls, proxies, and multifactor authentication) may not stop. Verizon Software Defined Perimeter (SDP) inherently implements multifactor authentication after the user logs into the Windows domain, stopping these attacks from accessing servers protected by the Verizon SDP. Other implementations of the Zero Trust Architecture may be more challenged to do so.

What is pass-the-hash?

Pass-the-hash and its relative, pass-the-ticket, are attack techniques commonly used by adversaries to compromise enterprise servers after some other attack has created a foothold inside the enterprise. These attacks enable the adversary to impersonate other authorized users from the initially compromised computer – including a server administrator.

Both attacks rely on the underlying authentication and authorization methodology of Microsoft Windows Domain networking. Therefore, we need to understand how authentication and authorization work in a Windows Domain before studying the attacks themselves.

How common are these attacks?

By referencing the MITRE ATT&CK matrix[1], we can see that these techniques are used by APT1[I], APT28[II], APT32[III], and other state-sponsored adversaries. In addition, there are a number of open source hacking tools such as Mimikatz, Pass-The-Hash Toolkit, Metasploit, and others[IV] that enable both knowledgeable hackers and script kiddies alike to execute these attacks.

Why haven't I heard about these attacks?

Few security professionals know about these attacks, possibly because many traditional security products do not stop them.

Since these attacks enable the adversary to impersonate authorized users of the Windows Domain, security products like firewalls, proxies and multifactor authentication are not effective. And when security products do not stop a particular attack, vendors tend to not talk about those attacks.

How does a user log into Windows?

When a user enters his or her username and password to log into a Windows PC, the user's password gets hashed by a well-known formula[2]. Then, instead of using the plaintext of the user's password, Windows uses the hash of the password in a challenge/response handshake to log into the Domain Controller.

After logging into the PC, it is the hash of the user's password that acts like the user's password throughout the domain. So, the adversary does not need to crack the hash of the password because the hash acts like the password This is true for both NTLM and Kerberos authentication. Furthermore, the password hash is stored in the RAM of the PC to enable Single Sign On (SSO) throughout the Windows domain. And, by default, it is not just the present user whose password hash is stored in RAM, but rather the password hashes of the last 10 users who have logged into this PC. That includes logging in via Remote Desktop Protocol (RDP) – typically, the action of the desktop admin.

How does pass-the-hash work?

As mentioned, first the adversary must compromise a domain-connected PC (and there are many ways to do that). Now, the adversary can use the tools discussed to extract all of the password hashes from RAM. If the desktop admin's hash is on the PC, and if the desktop admin is also the server admin, then the adversary's job is almost done. More likely, however, the server admin uses a different username and password to access the servers.

Recall from a previous Verizon white paper[3] that once an adversary gets access to the internal network, there's almost always at least one server that has known vulnerabilities or configuration errors that would enable the adversary to compromise it. After compromising that server, the adversary uses the password dumping tools on it.

verizon

It's likely that the adversary will get the server admin's password hash this way. If not, the adversary simply crashes the server and forces the server admin to restart it – providing the adversary the server admin's hashed password at that time. At that point, in your typical flat network, the adversary can log into any server he or she wishes **as the admin**.

How does pass-the-ticket work?

The pass-the-hash attack was invented back in the late 1990's[4] and used when NTLM was the primary authentication and authorization protocol for Windows domains. But Kerberos has become the dominant authentication and authorization, so pass-the-hash should not be possible, right? As it turns out, if the adversary attempts to connect to a Windows server via the IP address of the server instead of the host name, then the server defaults to NTLM authentication. Therefore, it is often possible for the adversary to bypass Kerberos authentication with this method. But, if NTLM authentication has been fully disabled, there is a similar attack known as pass-the-ticket.

For Kerberos authentication and authorization, after the user logs into the PC and the PC logs into the domain controller, the PC gets a Ticket Granting Ticket (TGT) that represents the user's permissions to access other Windows computers. The TGT is also stored in RAM for Single Sign On purposes. Unfortunately, hacking tools can also extract TGTs, and the TGT can also be used by the adversary to impersonate the server admin via a similar attack scenario as above. The major difference between the two attacks is that the hash of pass-the-hash lasts as long as the user has the same password (often up to 90 days) whereas the TGT often expires within 10 hours – making the adversary repeat the attack to get the new TGT.

Verizon SDP can defeat pass-the-hash.

With Verizon SDP, multifactor authentication is enforced after the user logs into the Windows domain and ties the user to the user's PC. Therefore, the user's password hash cannot be used on another PC. For example, if Alice and Bob are both members of the Verizon SDP, then Alice can log into the Verizon SDP on her computer but not on Bob's. Therefore, the password hash of a server admin, exfiltrated from a compromised server, cannot be used on the victim's computer to access SDP-protected applications.

1. Mitre ATT&CK matrix: Use Alternate Authentication Material: Pass the Hash, attack.mitre.org/techniques/T1550/002/
 i. MITRE listing for attack group APT1: https://attack.mitre.org/groups/G0006/
 ii. MITRE listing for attack group APT28: https://attack.mitre.org/groups/G0007/
 iii. MITRE listing for attack group APT32: https://attack.mitre.org/groups/G0050/
 iv. MITRE pass-the-hash technique: https://attack.mitre.org/techniques/T1550/002/

2. Microsoft Forum. social.technet.microsoft.com/Forums/ie/en-US/98d03c12-0cc5-4248-ac45-3461aa035206/does-microsoft-implementation-of-kerberos-use-a-salt-in-its-stringtokey-function?forum=winserversecurity

3. Verizon Zero Trust Architecture white paper by Brent Bilger. 2020.

4. Pass-the-hash. www.securityfocus.com/bid/233/info

5. IBID

6. NIST Special Publication 800-207. Zero Trust Architecture by Scott Rose, Oliver Borchert, Stu Mitchell, Sean Connelly. https://doi.org/10.6028/NIST.SP.800-207

What about traditional multifactor authentication?

Traditional multifactor authentication does not defeat this attack. RSA tokens, phone-as-a-factor (e.g., Duo), text messages, Common Access Cards (CAC) used by the DoD, and other forms of traditional multifactor authentication – none of these methods stop either pass-the-hash or pass-the-ticket. This is because all of them occur before the hash of the user's password is created by the Windows PC and because the secrecy of the password hash is based solely on the user's password[5]. And, as stated before, traditional security tools do not stop authorized users either.

What about remote access VPN users?

Traditional remote access VPN products put the user "on the LAN". Therefore, those users are susceptible to the same attack scenario – the adversary compromises the PC of a remote access user, extracts the password hashes, scans the enterprise network for a vulnerable server, compromises the server, and "becomes" the server admin.

Defeats all pass-the-hash attacks?

Verizon SDP can defeat client-to-server pass-the-hash but not server-to-server. To defeat the latter, put different applications on different VLANs or different VPCs during cloud migration. Pass-the-hash can also occur from client-to-client, typically, using RDP and the hash of the desktop admin (as explained previously). Pass-the-hash via RDP can also be defeated using Verizon SDP by implementing a PC-based Windows firewall rule to limit RDP access to only the IP addresses of an SDP Gateway cluster. This enables static IP addressing in Windows firewall to defeat client-to-client pass-the-hash via RDP and forces the desktop admins to go through an additional layer of security.

Can a product be Zero Trust if it cannot stop it?

You decide. The Verizon SDP product implements multifactor authentication after the user logs into the domain – effectively defeating both attacks. Many products that claim to implement the Zero Trust Architecture[6] often depend on other solutions formultifactor authentication and do not implement multifactor authentication after Windows login.

Summary.

Pass-the-hash and pass-the-ticket are commonly used attacks which many traditional security products do not stop. Verizon SDP inherently implements multifactor authentication after the user logs into the Windows domain, and, thus, can stop these attacks from accessing servers protected by the Verizon SDP.

verizon