NESCO/NESCOR

Common TFE Analysis:

CIP-007 R5.3 Password Complexity

National Electric Sector Cybersecurity Organization (NESCO)/NESCO Resource (NESCOR)

DISCLAIMER OF WARRANTIES AND LIMITATION OF LIABILITIES

THIS DOCUMENT WAS PREPARED BY THE ORGANIZATION(S) NAMED BELOW AS AN ACCOUNT OF WORK SPONSORED OR COSPONSORED BY THE ELECTRIC POWER RESEARCH INSTITUTE, INC. (EPRI). NEITHER EPRI, NOR ANY MEMBER OF EPRI, ANY COSPONSOR, THE ORGANIZATION(S) BELOW, NOR ANY PERSON ACTING ON BEHALF OF ANY OF THEM:

> (A) MAKES ANY WARRANTY OR REPRESENTATION WHATSOEVER, EXPRESS OR IMPLIED, (I) WITH RESPECT TO THE USE OF ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT, INCLUDING MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, OR (II) THAT SUCH USE DOES NOT INFRINGE ON OR INTERFERE WITH PRIVATELY OWNED RIGHTS, INCLUDING ANY PARTY'S INTELLECTUAL PROPERTY, OR (III) THAT THIS DOCUMENT IS SUITABLE TO ANY PARTICULAR USER'S CIRCUMSTANCE; OR
>
> (B) ASSUMES RESPONSIBILITY FOR ANY DAMAGES OR OTHER LIABILITY WHATSOEVER (INCLUDING ANY CONSEQUENTIAL DAMAGES, EVEN IF EPRI OR ANY EPRI REPRESENTATIVE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES) RESULTING FROM YOUR SELECTION OR USE OF THIS DOCUMENT OR ANY INFORMATION, APPARATUS, METHOD, PROCESS, OR SIMILAR ITEM DISCLOSED IN THIS DOCUMENT.

REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY ITS TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE, DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY EPRI.

THIS REPORT WAS PREPARED AS AN ACCOUNT OF WORK SPONSORED BY AN AGENCY OF THE UNITED STATES GOVERNMENT. NEITHER THE UNITED STATES GOVERNMENT NOR ANY AGENCY THEREOF, NOR ANY OF THEIR EMPLOYEES, MAKES ANY WARRANTY, EXPRESS OR IMPLIED, OR ASSUMES ANY LEGAL LIABILITY OR RESPONSIBILITY FOR THE ACCURACY, COMPLETENESS, OR USEFULNESS OF ANY INFORMATION, APPARATUS, PRODUCT, OR PROCESS DISCLOSED, OR REPRESENTS THAT ITS USE WOULD NOT INFRINGE PRIVATELY OWNED RIGHTS. REFERENCE HEREIN TO ANY SPECIFIC COMMERCIAL PRODUCT, PROCESS, OR SERVICE BY TRADE NAME, TRADEMARK, MANUFACTURER, OR OTHERWISE DOES NOT NECESSARILY CONSTITUTE OR IMPLY ITS ENDORSEMENT, RECOMMENDATION, OR FAVORING BY THE UNITED STATES GOVERNMENT OR ANY AGENCY THEREOF. THE VIEWS AND OPINIONS OF AUTHORS EXPRESSED HEREIN DO NOT NECESSARILY STATE OR REFLECT THOSE OF THE UNITED STATES GOVERNMENT OR ANY AGENCY THEREOF.

THE FOLLOWING ORGANIZATION PREPARED THIS REPORT:

ELECTRIC POWER RESEARCH INSTITUTE (EPRI)

ACKNOWLEDGMENTS

The research was paid for by the Department of Energy (DOE) under the NESCOR grant.

Principal Investigator

A. Lee

Primary Author:

Andrew K. Wright, N-Dimension Solutions, Inc.

NESCO/NESCOR Common TFE Analysis:

CIP-007 R5.3 Password Complexity

Abstract

TFE Category NERC CIP-007-4 **R5.3** specifies password complexity requirements for critical cyber assets. Many utilities are filing Technical Feasibility Exceptions (TFEs) to this requirement for Microsoft Windows systems because these systems cannot enforce the requirements.

1. TFE Category

Category of Equipment: primarily Microsoft Windows systems, but similar issues arise for all equipment that uses passwords to authenticate users for access or configuration, including RTUs, PLCs, relays, reclosers, communications processors, radios, modems, routers, switches, firewalls, etc.

2. CIP Requirement

CIP-007-4 **R5.3** reads as follows:

> **R5.3.** At a minimum, the Responsible Entity shall require and use passwords, subject to the following, as technically feasible:
>
> **R5.3.1.** Each password shall be a minimum of six characters.
>
> **R5.3.2.** Each password shall consist of a combination of alpha, numeric, and “special” characters.
>
> **R5.3.3.** Each password shall be changed at least annually, or more frequently based on risk.

3. TFE Description

Many TFEs have been filed for Windows systems, including but not limited to Windows XP, Windows Server 2003, and Windows Server 2008, since these systems cannot enforce the **R5.3** password complexity requirements.

The password complexity requirements that Windows Server 2008 is able to enforce are as follows [1]:

- Minimum password length can range from 0 to 14 characters.
- Passwords must contain characters from three of the following four categories:
  - English uppercase characters (A through Z).

    - English lowercase characters (a through z).
    - Base 10 digits (0 through 9).
    - Non-alphabetic characters (for example, !, $, #, %).
- Password maximum age can range from 0 to 999 days.
- Passwords cannot be reused (up to 24 old passwords remembered).
- Password minimum age can range from 0 to 998 days (generally used in conjunction with preventing reuse).

Other versions of Windows implement similar constraints. The password length and age requirements of **R5.3** can be enforced by Windows. The “three of the following four categories” requirement allows the construction of a password that contains uppercase, lowercase, and base 10 digits but does not contain any special characters. Such a password violates **R5.3.2**.

4. Security Risks Addressed by CIP Requirement

The security risk that CIP-007-4 **R5.3** is intended to address is that of an adversary determining and using an authorized person’s password to compromise a critical cyber asset. There are a variety of attacks that an adversary might use to learn passwords:

- offline attacks
- online trials
- eavesdropping
- social engineering
- man-in-the-middle attacks
- key loggers
- vendor default passwords
- shoulder surfing, covert video cameras, dumpster diving, post-it notes, and other physical attacks
- hybrid attacks – combinations of the above

Password complexity principally addresses offline attacks and certain eavesdropping attacks. Other attacks are largely unaffected by password complexity.

To mitigate offline attacks, Windows, UNIX, and many other systems store some form of one-way hash of a user’s password, rather than storing the plaintext password itself. To verify a password, the string entered by a user is hashed and compared to the stored hash. An offline attack involves obtaining a copy of the password hash database stored on a system and analyzing that information to recover passwords. A brute force attack is a simple offline attack that tries all possible passwords up to a certain length. **R5.3.1** attempts to address brute force attacks by requiring a minimum password length to increase the brute force search space. A dictionary attack is an offline attack that limits the search space of attempted passwords to a dictionary of commonly used words, and can thereby execute faster or try longer passwords than a pure brute force attack. A password such as “megalomaniac” that contains 12 characters would take a long time to be found by a pure brute force attack, but would likely be found by a dictionary attack. **R5.3.2** attempts to address dictionary attacks by requiring passwords to contain characters not found in dictionary words. Modern password cracking tools based on dictionary attacks often implement transmutation rules that mutate the provided dictionary word in prescribed ways to generate other potential passwords.

A secondary type of attack mitigated by password complexity is eavesdropping on authentication exchanges, which often involves sending a password hash or some other form of hash over a communications channel. By capturing the hash from the communications channel, the attacker can run brute force and dictionary attacks against it as described above. Again, password complexity requirements make the search space larger for brute force and dictionary attacks against the authentication exchange.

While password complexity requirements can also make some of the other types of attacks more difficult, there are generally better defenses against those types of attacks. For example, enforcing a permanent or temporary lockout after a small number of failed login attempts is a better defense against online trials.

5. Technical Analysis of Feasibility of Mitigating Risks

As indicated above, password complexity requirements primarily address offline attacks. This analysis focuses on offline attacks that attempt to recover passwords from password hash databases or authentication algorithms.

Windows Specific Risks

Versions of Windows prior to Windows NT used a password hash known as the LAN Manager or LM hash that has serious flaws. Virtually all LM hashes can be cracked without resorting to brute force in a matter of seconds on a modern desktop computer, using a freely available open source software package known as Ophcrack [2], or through various online cracking websites such as OnlineHashCrack.com [10].

Later versions of Windows than NT switched to a password hash algorithm that uses an MD4 hash. MD4 is a stronger hash algorithm, but recently several preimage attacks that essentially allow reversing an MD4 hash have been discovered [3][4]. While these attacks do not make reversal of an MD4 hash trivial, they do significantly reduce its strength from 2^{128} to 2^{78} bits, and thereby make practical attacks feasible. IETF RFC 6150 [5] retires RFC 1320 that documents the MD4 algorithm, and recommends against continued use of MD4. Additionally, the Windows password hash algorithm does not use salt, and is therefore vulnerable to the use of rainbow tables [12] to speed up a dictionary search.

Versions of Windows prior to Windows Server 2008 by default store both MD4 and LM hashes of a password for backwards compatibility. With both hashes available, an attacker can use the weaker of the two hashes to recover the password. Consequently recovering passwords from any such system is trivial.

Use of LM hashes can be disabled through group policy configuration on Windows systems and Active Directory servers [11]. LM hashes are not stored for passwords of 15 or more characters in length, but note that a minimum password length of 15 characters cannot be enforced using group policy. Windows 7, Vista, and Server 2008 no longer store LM hashes by default, but still use a simple unsalted MD4 hash.

Dictionary Attacks

Modern password cracking programs such as “John the Ripper” [14] go beyond simple dictionary searches. Such programs will generate and try variations on dictionary words by capitalizing various letters, inserting numbers and special characters in various positions, making replacements such as ‘0’ for ‘o’ and ‘1’ for ‘i’, reversing words, etc. Consequently, many of the methods that people use to add numbers and special characters to otherwise simple passwords in order to meet complexity requirements like that of **R5.3.2** add a relatively small factor to the search space for modern dictionary attacks. Thus a password such as “M3GA10man1ac” is not much better than “megalomaniac”.

Rainbow tables are a form of time/memory tradeoff that can be used to speed up dictionary attacks significantly. Rainbow tables are defeated by the inclusion of salt in computing the hash of the password. Most configurations of Linux use salt, but Windows systems do not.

Brute Force Attacks

Recent advances in Graphics Processing Units (GPUs) and the development of general purpose programming interfaces for GPUs have made highly parallel computing effective and affordable. This development is important because brute force password cracking is a highly parallelizable problem and ideally suited for GPUs.

The nVidia GTX 580 GPU, announced November 2010, marketed for gaming, and retailing for less than $500, has 512 cores and a peak performance of 1.58 TFLOPs. For parallelizable problems, this GPU exceeds the performance of the ASCI Red Supercomputer built by Sandia National Laboratory in 1996 at a cost of $50 million. This mid-performance GPU is capable of performing 1.2 to 1.3 billion MD5 (or similar algorithm) hashes per second, according to research performed by Richard Boyd of Georgia Tech [6]. The Chinese Tianhe-1A supercomputer includes not only 14,336 general purpose CPUs, but also 7,168 nVidia M2050 GPUs, each of which has 448 cores, and is thus capable of well in excess of a trillion hashes per second. For passwords consisting of characters drawn from the ASCII printable character set, which includes 95 characters, there are 956 = 0.75 trillion possible 6-character passwords. A machine such as the Tianhe-1A can therefore try all possible 6-character passwords that are stored as hashes in less than one second. Table 1 shows worst-case brute force MD5 password cracking time on a GPU cluster as a function of password length and complexity [6]:

Worst-Case Brute Force MD5 Cracking at One Trillion Hashes per Second			
Password Length	**Lowercase Letters**	**Letters & Numbers**	**All 95 Characters**
8	< 1 second	3.7 minutes	1.9 hours
9	5.5 seconds	3.8 hours	7.3 days
10	2.4 minutes	9.8 days	1.9 years
11	1.1 hours	1.7 years	180 years
12	1.2 days	102.4 years	17,135 years

Table 1: Brute Force MD5 Hash Cracking

In interpreting this table, it is important to note that these worst-case times represent the time to try all possible passwords. A search for a single randomly generated password would on average take half of the worst-case time. Searching for any one of a small number N of passwords will reduce the search time by a factor of N. Various techniques based on the fact that most passwords are far from randomly chosen can reduce the search space and search time even further. Finally, GPU performance is continuing to improve, with recent performance advances by nVidia and AMD besting Moore's law and roughly doubling GPU performance every year.

The introduction of cloud computing has made cluster computing ala the Tianhe-1A available to anyone on the Internet. Amazon's EC2 cloud includes GPU machine instances [9], and anyone can therefore run password cracking algorithms on a rented cluster of GPUs without building a supercomputer. Using 64 rented GPU machine instances, cracking a 6-character password stored as a hash would take about 100 times longer than on the Tianhe-1A supercomputer, but could be performed in about 15 seconds for a cost of about $0.50. Thomas Roth described his development of such a brute force hash cracking software suite that runs on the EC2 cloud earlier this year at Black Hat 2011 [7], although the software does not yet appear to be generally available. A similar service for cracking passwords to WiFi networks that use WPA or WPA2 preshared keys became available in early 2011 [8]. This service provides "access to a 400 CPU cluster that will run your network capture against a 135 million word dictionary created specifically for WPA passwords", and indicates that "while this job would take over 5 days on a contemporary dual-core PC, on our cluster it takes an average of 20 minutes, for only $17."

On June 20 2011 Intel introduced a new Many Integrated Core (MIC) architecture that it believes will lead to the development of a new generation of supercomputers that should break the ExaFLOP-per-second level by 2020 [15]. For highly parallelizable problems such as password cracking, this is approximately 1000 times faster than the Tianhe-1A, meaning an Intel MIC-based supercomputer should be able to try all possible 8-character passwords in the same time as the Tianhe-1A can try all 6-character passwords, in less than one second.

Unlike a dictionary attack performed using rainbow tables, a brute force attack is not significantly affected by the use of salt, since it does not attempt any precomputation.

Brute force attacks are feasible in part because hash algorithms such as MD4, MD5, and the SHA variants were designed to be fast. The results in Table 1 apply to passwords stored as MD5 hashes, but would be similar for other hash algorithms. A technique known as *key stretching* involves using an algorithm to compute a hash or key from a password that requires some nontrivial time to execute on a typical processor – long enough to significantly impact brute force attacks, but not so long as to impact the user experience. For example, recent distributions of Linux can be configured to iterate 5000 rounds of SHA-512 to compute the stored hash [16]. This makes a 10-character Linux password nearly as strong as a 12-character password that does not use any key stretching, since the 5000 rounds of SHA-512 make brute force search against the Linux password take 5000 times longer, while, two additional characters of password length expand the search space by a factor of 95*95=9025. Windows does not use key stretching.

Eavesdropping Attacks

Eavesdropping attacks can be used to capture messages exchanged during the execution of an authentication protocol, such as signing on to a WiFi connection or establishing a VPN. For authentication protocols that utilize passwords or preshared keys, depending on the design of that authentication protocol, it may be possible to perform offline brute force password cracking attacks against the messages obtained during a valid authentication that are similar to offline password database attacks.

The Internet Key Exchange (IKE) protocol is used in establishing authentication for an IPSEC VPN. When preshared keys are used, if the server can be forced to use aggressive mode rather than main mode, then a simple authentication hash can be intercepted [17]. Consequently IPSEC VPNs that use preshared keys and permit aggressive mode require effort to crack using brute force search similar to that shown in Table 1.

Most TLS servers on the Internet authenticate clients using public key certificates, and consequently these authentication exchanges are not vulnerable to offline password attacks. However, RFC 4279 [18] and RFC 4785 [19] are proposed standards that add several preshared authentication modes to TLS. Authentication exchanges negotiated using the PSK modes of these RFCs appear to require effort to crack comparable to that shown in Table 1.

For WiFi networks, WEP and WPA are known to have significant weaknesses. WPA2 is currently the only algorithm for WiFi network encryption that does not have significant known weaknesses. With preshared keys, WPA2 uses a *key derivation function* known as PBKDF2 that uses 4096 rounds of SHA1 and the SSID as salt during the authentication exchange. Consequently WiFi networks using WPA2 and preshared keys require about 4096 times more work to crack than Table 1.

Issues Specific to the Power Grid

Offline attacks are a high risk for many systems deployed in the power grid due to the proliferation of devices that are deployed in unmanned locations with limited physical security. Examples include devices in substations and on pole tops. Extraction of password hash databases from these systems may in some cases be possible without knowledge of the utility. Many older devices are likely to store passwords in the clear, obviating even the need for password cracking.

Passwords are frequently reused across unattended devices in the field, such as substation and pole-top equipment. Such reuse means that compromise of the password from a relatively unimportant device, perhaps one located in a substation, could lead to compromise of a more critical device, perhaps one located in a control center.

Many SCADA and control systems protocols send passwords in plaintext across unencrypted communications channels. In such cases, eavesdropping attacks to capture passwords are trivial.

6. Summary

Windows systems that store LM hashes are highly vulnerable to offline attacks, regardless of password complexity. There are several methods for preventing Windows from storing LM hashes as specified in group policy [11]. Some Windows authentication protocols (e.g. PPTP) send password hashes, and these protocols are thus highly vulnerable to eavesdropping attacks when LM hashes are in use.

Practical cracking tools for Windows MD4 hashes using preimage attacks are likely to become generally available in the near future. Alternatively, or in addition, these techniques will be used to speed up existing cracking tools and techniques, such as rainbow tables. Once this happens, Windows passwords will become highly vulnerable regardless of password complexity and length.

With current GPU performance and cluster computing available via cloud services, passwords shorter than 12 characters stored in Windows systems and shorter than 10 characters stored in Linux systems or used for IPSEC or WPA2 preshared keys are vulnerable to practical offline attacks. Even 12 character and longer passwords may become vulnerable within the 10+ year expected lifetime for power systems equipment.

R5.3 does not mitigate the risk that it is intended to address. Furthermore, **R5.3** cannot be complied with through technical means on Windows systems without installing 3rd party tools.

References

[1] http://technet.microsoft.com/en-us/library/cc264456.aspx

[2] http://ophcrack.sourceforge.net

[3] G. Leurent. MD4 is Not One-Way. Fast Software Encryption 2008, Lausanne, Switzerland, February 10-13, 2008, LNCS 5086. Springer, 2008.

[4] Guo, J., Ling, S., Rechberger, C., and H. Wang, "Advanced Meet-in-the-Middle Preimage Attacks: First Results on Full Tiger, and Improved Results on MD4 and SHA-2". http://eprint.iacr.org/2010/016.pdf

[5] http://tools.ietf.org/html/rfc6150

[6] Richard Boyd, Senior Research Scientist, Georgia Tech Research Institute, personal communication.

[7] Thomas Roth, Black Hat 2011. http://www.blackhat.com/html/bh-dc-11/bh-dc-11-briefings.html

[8] http://www.wpacracker.com/

[9] http://aws.amazon.com/ec2/hpc-applications/

[10] http://www.onlinehashcrack.com/

[11] http://support.microsoft.com/kb/299656/

[12] http://en.wikipedia.org/wiki/Rainbow_table

[13] http://technet.microsoft.com/en-us/library/dd277399.aspx

[14] http://www.openwall.com/john/

[15] http://newsroom.intel.com/community/intel_newsroom/blog/2011/06/20/intel-equipped-to-lead-industry-to-era-of-exascale-computing

[16] http://www.kernel.org/pub/linux/libs/pam/Linux-PAM-html/sag-pam_unix.html

[17] http://www.symantec.com/connect/articles/penetration-testing-ipsec-vpns

[18] http://www.ietf.org/rfc/rfc4279.txt

[19] http://www.ietf.org/rfc/rfc4785.txt