



CYBERARK®

CyberArk Labs

Analyzing Real-World Exposure to Windows Credential Theft Attacks

RESEARCH

Introduction

The term “credentials theft attack” refers to an attacker gaining unauthorized access to a user account’s credentials, such as a password or a password hash, and then using this stolen credential to authenticate as the associated user, gaining access to whatever resources the user can access. If the user account has access to any of the organization’s sensitive data, such as credit card numbers, bank account details or medical records, then unauthorized reuse of that user’s credentials alone could directly lead to a data breach. Examples of major breaches that began with credential theft include the Target breach in 2013 and the Home Depot breach in 2014.

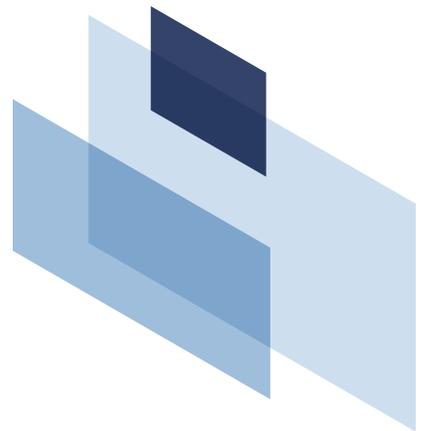
Even more dire, however, is the potential impact of credential theft if the stolen credential is *privileged*, meaning that it offers greater operating system permissions than a standard user account does. With credentials for a privileged account, an attacker can gain complete control over the host or hosts that accept those credentials. This allows the attacker not only to access and breach all the sensitive data on those hosts, but also to perform other malicious actions such as installing malware and disabling or reconfiguring security controls. These malicious actions help to establish a long-term presence for the attacker on the host, which in many cases may go undetected for months, even years. Rather than performing a one-time breach, an attacker could steal data or conduct other malicious actions repeatedly over the course of time.

This paper documents research recently conducted by CyberArk Labs to investigate real-world exposure to successful credential theft attacks against privileged accounts in Microsoft Windows networks composed of Windows servers and workstations. This research involves not only the direct impact of unauthorized credential reuse, but also how such a credential could be used to collect other credentials, such as passwords for privileged accounts on other Windows hosts, thus giving the attacker unauthorized privileged access to those hosts as well. Based on the findings of this research, in most Windows environments a savvy attacker can often build on a single compromise of a single Windows host to eventually compromise most or all of the Windows hosts on that organization’s network, as well as exposing all of the sensitive data that they store or can access.

Credential Theft Methods

There are many ways to perform credential theft. Some methods, such as social engineering and phishing, rely on a person being tricked and inadvertently revealing their credentials to an attacker. Other methods, such as keystroke loggers, use malware to infect a machine and monitor all of its activities, including a person typing in passwords. There are also methods that rely on searching a Windows host for unprotected credentials, including those hard coded into scripts or applications.

Finally, some forms of credential theft attack focus on *password hashes*, which are cryptographically secure representations of passwords that are used to prevent the disclosure of the actual password to unauthorized parties. Password hashes may inadvertently be left in memory after being used until a host is restarted. An attacker who gains access to password hashes may be able to reuse them to perform a *pass-the-hash attack*. Similarly, an *over-pass-the-hash attack* involves password hashes that are being used in a Kerberos-enabled network for Kerberos authentication to Windows hosts.



Research Methodology

CyberArk needed real-world data on privileged Windows accounts to conduct this research, so the first step in the research methodology was to ask a sampling of organizations around the world to scan the Windows hosts on their enterprise networks using the CyberArk DNA tool. These networks ranged from small networks with only a few dozen Windows hosts to large enterprise networks with thousands of Windows hosts each. CyberArk received results from 51 networks.

The next major step in the research was to analyze the results across all of the participating organizations. Major questions to be answered through this analysis included the following:

- If an attacker compromised a single critical Windows host, could the attacker leverage that compromise to gain access to other Windows hosts on the network? If so, what percentage of the network's other Windows hosts are accessible this way?
- For an attacker who wants to repeatedly use one machine to compromise another and chain these attacks together to reach an ultimate target, what is the relative value of gaining access to a Windows server as opposed to a Windows workstation?
- Which types of privileged Windows accounts pose the most danger to organizations?
- How many privileged accounts does an organization have that each offer widespread access to its Windows hosts?
- How can the credential reuse vulnerabilities in privileged user accounts be mitigated?
- How effective is each mitigation strategy in lowering the susceptibility to attacks and the impact that a single attack can cause?

CyberArk DNA

This research used the CyberArk Discovery & Audit (DNA) tool. This tool is freely available for organizations to use to gather information on their privileged accounts. An organization can run the scanning tool on one or more of its networks to identify the privileged accounts on the networks' hosts and the status of each privileged account's password. An organization can then analyze this collected information to assess the scope of its exposure due to weakly secured privileged accounts.

For more information on CyberArk DNA and to download it, visit <http://www.cyberark.com/discovery-audit-cyberark-dna/>.

The final step in the research methodology was to draw conclusions and recommendations from the analysis, then document those findings in this paper, so that the lessons learned from this research could help organizations to improve their security practices.

Measuring Network Exposure

One of the major goals of this research effort was to determine how much of a Windows network is typically exposed through compromise of a single privileged account on a single Windows host. Figure 1 shows a simple notional Windows network architecture. Each arrow in the figure indicates that there is a privileged account on the source host that can grant an attacker privileged access to the target host (the host the arrow is pointing to). So an attacker that compromises the source host can readily steal additional credentials from it and use them to gain access to the target hosts.

In Figure 1, imagine that a privileged account on the host on the far left has been compromised by an attacker, making the compromised account and host a threat against other hosts. This host can act directly and immediately to compromise two targets. The compromised host also poses an indirect threat to the other four hosts in the sample network because those hosts can be compromised using credentials acquired by the attacker jumping from one host to another through the privileged accounts. Each of these hosts that can be directly or indirectly compromised via privileged accounts from the original host is known as a *victim*. Figure 1 shows a total of six victims available from the original host.

During our analysis, we determined how many Windows hosts in each network could be directly and indirectly reached by compromising privileged accounts. We looked at each possible combination of a Windows host and a privileged account accessible from that host to identify the percentage of Windows hosts on that network that could be compromised through that pairing, with the ultimate goal of identifying the Windows hosts that, if compromised, would enable the greatest degree of access to other Windows hosts on the same network.

Our analysis showed that in some networks, nearly all of the Windows hosts could be used to access nearly all the other Windows hosts through privileged accounts. In these networks, attackers that compromised just about any host would gain access to a huge range of resources and data. In other networks, only a small number of Windows hosts could be used to directly or indirectly access other Windows hosts through privileged accounts. An attacker targeting such an environment would have to put forth significant effort to find the "right" Windows host and account to compromise so that access to most other hosts could be acquired.

Figure 1. **Potential Victims of a Compromised Windows Host**

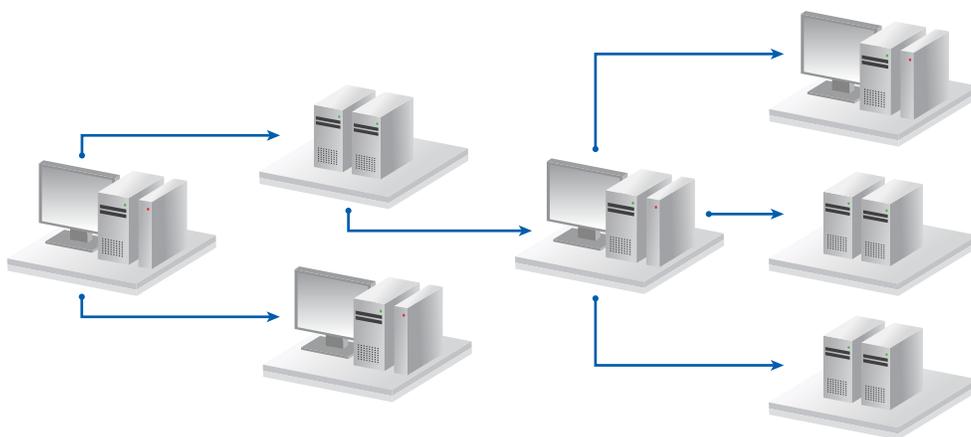
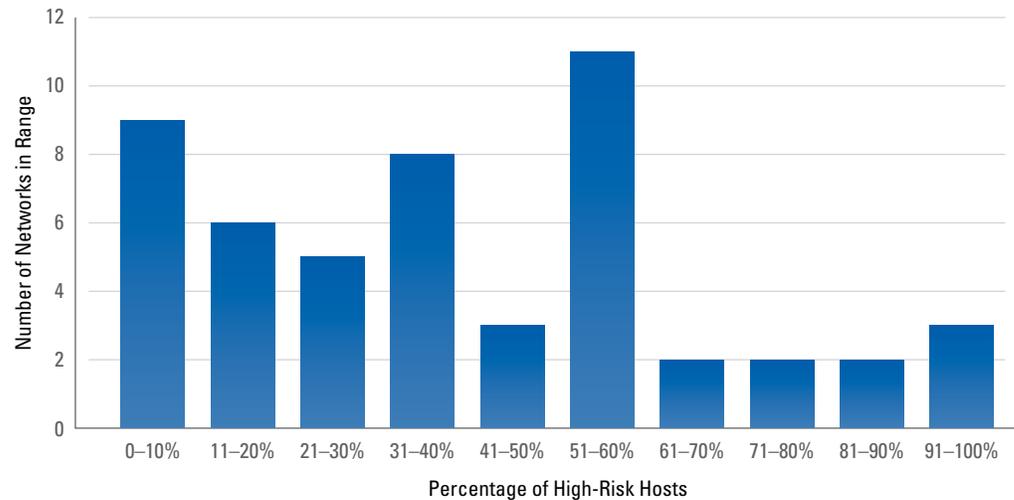


Figure 2. **Number of Hosts Grouped by Percentage of High-Risk Hosts**

Given the complexity of the data and its sheer size in terms of the number of hosts, we decided to simplify the presentation of these analysis results by developing a metric for them, which may also be of value to organizations in benchmarking the scope of their own privileged account issues as compared to other organizations.

Instead of reporting the percentage of the Windows hosts on the network that can be accessed directly or indirectly via privileged credentials from each host, we instead rolled these numbers up into a single value. First, we defined a *high-risk host* as one that can enable access to more than 80 percent of the network's other Windows hosts via privileged credentials. In other words, an attacker who can compromise a high-risk host can then use it to directly or indirectly gain privileged access to over 80 percent of the other Windows hosts on that network. Although an argument could be made for setting this threshold at a different value, there is no debate that a single host enables privileged access to such a high percentage of other hosts poses a high risk.

The metric we created indicates the percentage of hosts on a given network that are high-risk hosts. This makes it possible to measure host exposure to the theft and reuse of privileged credentials. What makes this simple metric so interesting is that it reflects the process that attackers follow to reach their ultimate targets. An attacker often starts an attack through

phishing or other means to gain access to a user account on a single host. If the attacker can leverage that access to gain privileged access, then the attacker can use those privileges to jump from host to host, using other privileged accounts acquired along the way as needed, before finally reaching the ultimate target. Networks with a high value for their percentage of high-risk hosts metric are making it much easier for attackers to achieve their goals.

When this metric was calculated for the networks scanned for the research, a great deal of variability was found. Figure 2 shows the distribution for the metric. The horizontal axis depicts the values of the high-risk hosts metric, grouped into ten percent ranges, and the vertical axis indicates how many networks fell into each range. Across the 51 scanned networks, the average value for the metric was 40%, and the median value was 36%.

As Figure 2 indicates, values for this metric are distributed across all the percentage ranges. The extreme values were 97 percent for one network and just over two percent for a few other networks. It is notable that while two percent is quite small, it does indicate that every scanned network has at least some hosts that can be leveraged to compromise most of the other hosts. This means that no organization with Windows hosts is "safe" from widespread compromise of those hosts through credential theft attacks.

Figure 3. **Windows Network Size Versus Percentage of High-Risk Hosts**

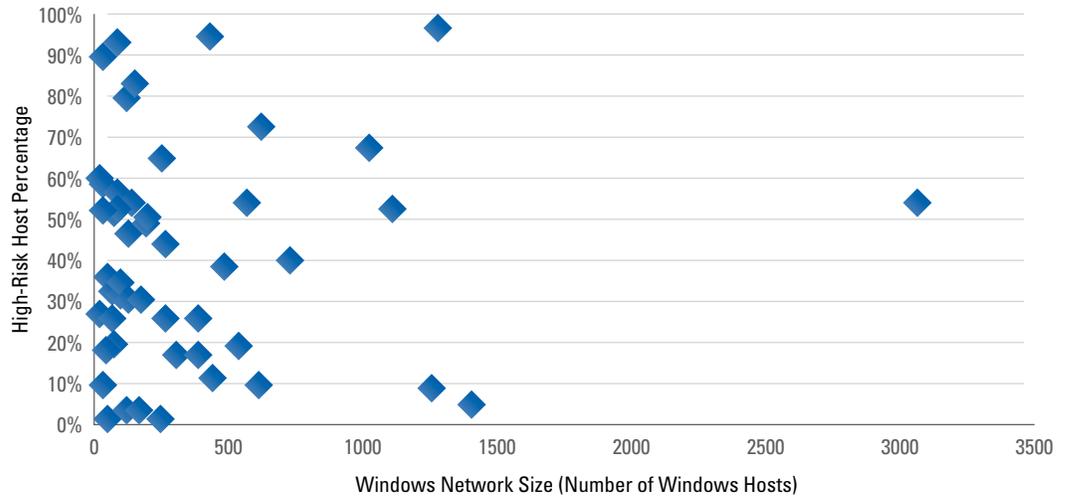


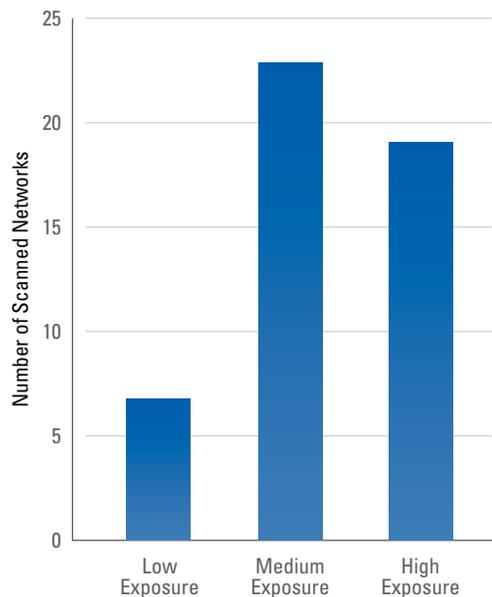
Figure 3 presents a different view of the metric values, which adds in a second factor: the number of Windows hosts scanned on each network, which is indicated on the horizontal axis. Figure 3 answers an interesting question: is there a correlation between the size of the network (the number of Windows hosts) and the percentage of high-risk hosts that the network has? Based on an analysis of Figure 3, the answer to that question is no. Networks of all sizes have a wide range of percentages. This means that every Windows network, no matter how large or small, could potentially be compromised by attackers through theft of privileged credentials.

To further simplify our representation of this data, we have defined the following groupings for networks:

- **Low-exposure:** a network where fewer than 10 percent of the hosts are high risk
- **Medium-exposure:** a network where 10 to 50 percent of the hosts are high risk
- **High-exposure:** a network where over 50 percent of the hosts are high risk

Figure 4 depicts the frequency of each of these groupings. We can see that most of the scanned networks—88 percent, to be precise—are either medium or high-exposure networks. This means that they are quite susceptible to being compromised through theft of privileged credentials.

Figure 4. **Number of Scanned Networks Per Exposure Level**



Identifying Key Pivot Points

A phenomenon discovered through this analysis and metric is that of key pivot points. A *pivot point* is a host that, once compromised, provides direct access to other hosts, and a *key pivot point* provides such access to many other hosts. One organization scanned 750 hosts and found that only three of those hosts had privileged account credentials that could grant direct access to most other hosts on the network. This would seem to indicate that an attacker would have to be lucky or perform extensive reconnaissance and exploitation activities to identify and gain access to one of the three affected hosts.

However, upon further analysis, this assumption is incorrect. The metric value for this network is actually 31 percent, with

232 hosts posing a high risk to the others. The reason is that the three hosts previously identified act as key pivot points on the network, with many hosts able to access them directly and many more hosts directly accessible from them. Figure 5 depicts a scaled-down example of this architecture. An attacker who can compromise one of the hosts on the left can use it to pivot to the central host, which then provides direct access to many hosts. This example illustrates the power of using our metric to calculate the percentage of high-risk hosts instead of just counting how many hosts have high-risk account credentials. See the “Measuring High-Risk Accounts” section for more insights into metrics related to high-risk account credentials.

Figure 5. **Example of a Key Pivot Point**

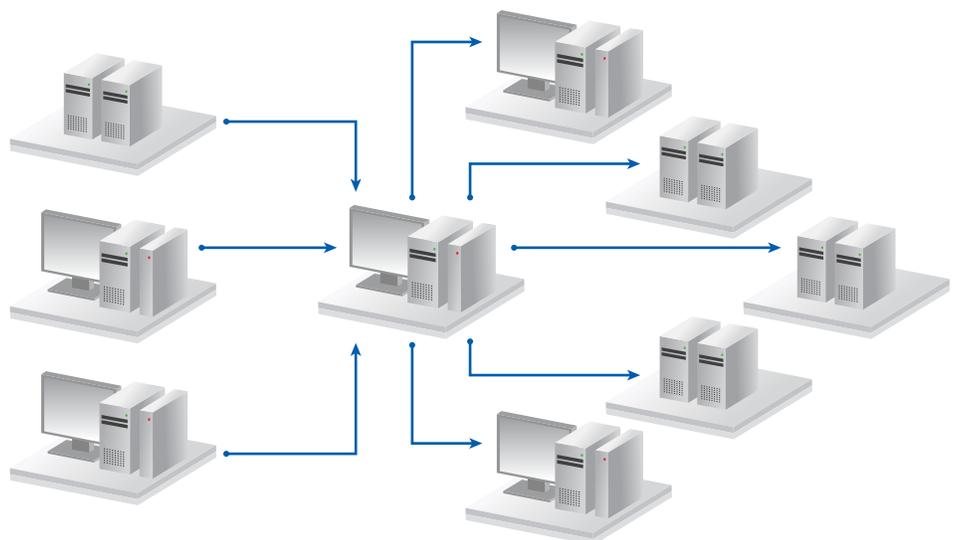
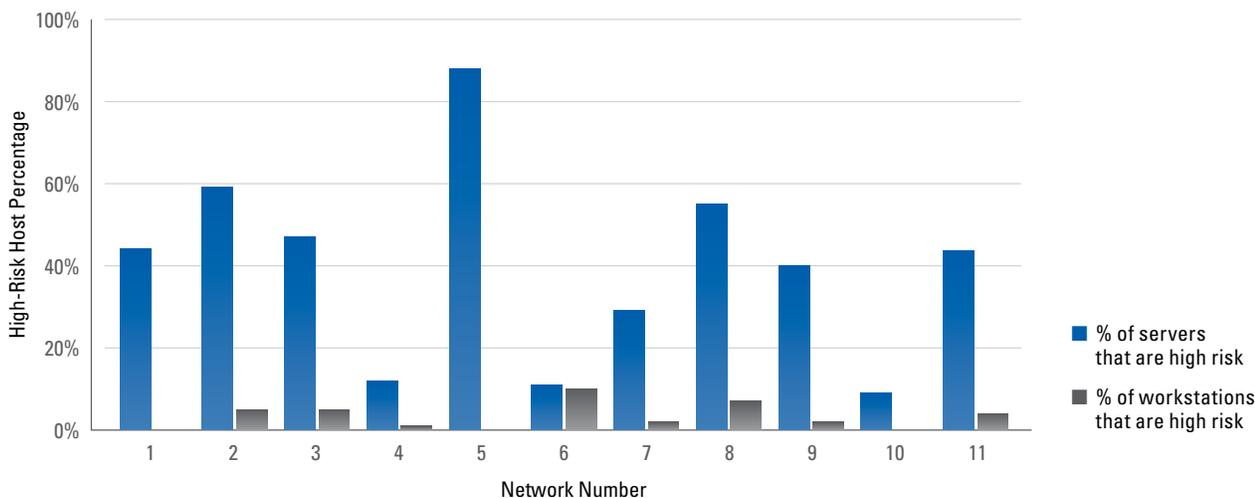


Figure 6. Comparing High-Risk Server and Workstation Metrics



Comparing Servers and Workstations

Another significant finding of this research is that compromised servers tend to pose a bigger threat to other Windows hosts than compromised workstations. Intuitively, system administrators, penetration testers, attackers, and others feel that servers are more valuable targets than workstations because they typically provide much greater access to sensitive data than individual workstations can. It also makes sense that servers tend to be more frequently accessed by domain-privileged users than workstations, and that servers are used to execute more automated processes with privileges than workstations. What may not have been clear, though, is the value of servers in acquiring privileged user credentials that grant access to most of the other Windows hosts on the network.

To illustrate our findings, we have selected a subset of the scanned networks. These 11 networks had both workstations and servers scanned. The percentage of high-risk hosts metric was calculated twice for each of these networks: once for the percentage of servers that are high risk, and once for the percentage of workstations that are high risk. Figure 6 shows the values for these metrics.

We can clearly see from Figure 6 that for all of these networks, the percentage of servers that are high risk is higher than the percentage of workstations that are high risk, and a great deal higher in several cases. Three of the networks had no high-risk workstations at all. For the other networks, the percentage of servers that are high risk was more than 10 times greater than the percentage of workstations that are high risk. In terms of the actual host counts instead of percentages, there are more than six times as many high-risk servers as there are high-risk workstations.

These findings mean that if an attacker is successful in compromising a server on a mixed server-workstation network instead of a workstation, that attacker has a much better chance of being able to steal credentials from it that enable the attack to be continued across many hosts, eventually compromising the majority of the network. The findings also show that regardless of the initial breach point, an attacker has a great incentive to move laterally to a server, as the chances of finding privileged credentials there are much higher than on workstations.



Measuring High-Risk Accounts

So far we have focused our analysis on the percentage of high-risk hosts, those hosts that pose the greatest risk to the network. We now turn our attention to a second metric we created: the *percentage of high-risk accounts*. This metric was constructed similarly to the high-risk hosts metric, but with an important distinction. We defined a *high-risk account* as a privileged account that can enable **direct** access to more than 80 percent of the network’s other Windows hosts. Indirect access – jumping through intermediate hosts – is not taken into consideration for high-risk accounts.

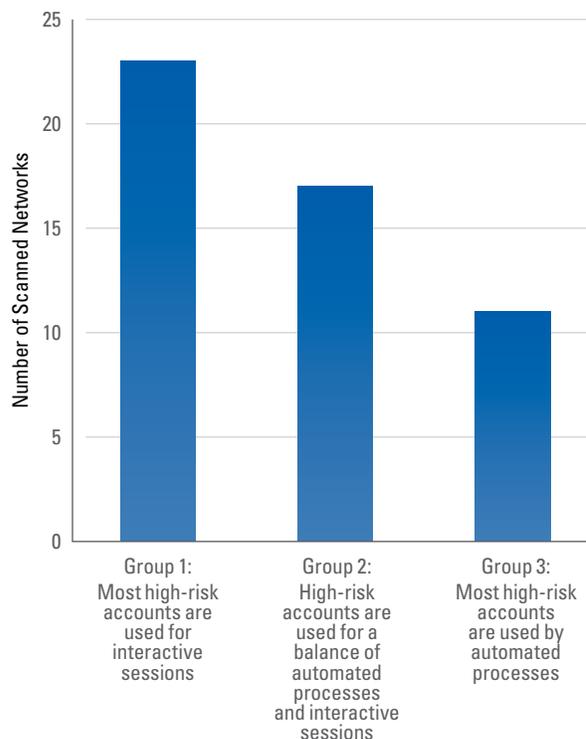
We can analyze the data on high-risk accounts by considering the different functions these accounts may perform. Each high-risk account falls into one of two mutually exclusive categories:

- Used by automated processes (service accounts)
- Used for interactive sessions (used by people to access hosts in the network)

Based on counting the number of high-risk accounts in each category, we can place each network into one of three groups, as depicted in Figure 7:

- Group 1: Most high-risk accounts are used for interactive sessions (0-33% of the high-risk accounts are for automated processes)
- Group 2: High-risk accounts are used for both automated processes and interactive sessions (34-66% of the accounts are for automated processes)
- Group 3: Most high-risk accounts are used by automated processes (67-100% of the high-risk accounts are for automated processes)

Figure 7. **Distribution of High-Risk Accounts**



Selecting Mitigation Strategies

In the real world, virtually every network has Windows hosts and privileged accounts that can pose serious risk to most of their other Windows hosts. It is critically important, particularly for those organizations with high values for the defined metrics, to lower their exposure to theft of privileged Windows credentials so as to lower their overall enterprise risk from attacks.

There are various mitigation options, thus it can be challenging to determine which options are most appropriate for a given environment. By using the CyberArk DNA tool and studying its results, an organization can identify those mitigations that are best suited for its privileged accounts.

There are mitigation options suitable for lowering the risk from privileged accounts used for interactive sessions (mostly applicable to high-risk account groups 1 and 2):

- Organizations can use privileged local accounts instead of privileged domain accounts. This prevents an attacker who compromises a host from using the credentials to gain privileged access to other hosts in the domain.
- Another option is to implement the use of one-time passwords, by using an automated tool that changes the password after every use of a privileged account. This frequent password rotation is typically performed behind the scenes, unbeknownst to the user, who accesses the host through an intermediate authentication service; the user doesn't even have direct access to the password. This option ensures that even if an attacker compromises a host and steals credentials, those credentials will only be valid on the compromised host for a short time and will not be valid for any other hosts.

Other mitigations can lower the risk from both types of privileged accounts – interactive session and automated process accounts (applicable to high-risk account groups 1, 2, and 3):

- Many organizations choose to use “*zoning*” – the organization grants each privileged account access to a small subset of the hosts in the network. This generally means having a larger number of privileged accounts, while minimizing the negative impact of a compromise of any single account.
- Organizations can also choose to implement *limited privileged domain* accounts. Such accounts have high privileges on a specific host, but only have lower privileges (e.g., standard user privileges, guest privileges) for other hosts.

And, finally, there are mitigation options for lowering risk from privileged accounts used for automated processes (mostly applicable to high-risk account groups 2 and 3):

- Embedding usernames, passwords, and/or other credentials in processes or in their proximity has never been a good security practice. It is highly recommended that organizations avoid doing this, and instead use a centralized system that can provide credentials to automated processes on demand. This greatly lowers the threat of an attacker retrieving the embedded credentials, and it also enables the organization to regularly rotate the passwords for these accounts instead of hard-coding passwords that are each used for years because of the overhead involved in modifying the software just to change a password.
- In addition, organizations can avoid the use of domain accounts for automated processes and use one or more of the following account types instead:
 - Local accounts – accounts that do not have privileges to access other hosts
 - Virtual accounts¹ – local accounts that can optionally access network resources using the HOST account of the host on which they are running
 - Managed service accounts – service accounts that have their credentials frequently replaced according to a predefined schedule

¹ [https://technet.microsoft.com/en-us/library/ff641731\(v=ws.10\).aspx](https://technet.microsoft.com/en-us/library/ff641731(v=ws.10).aspx)

Figure 8. **Reduction in High-Risk Hosts from Service Account Mitigation**

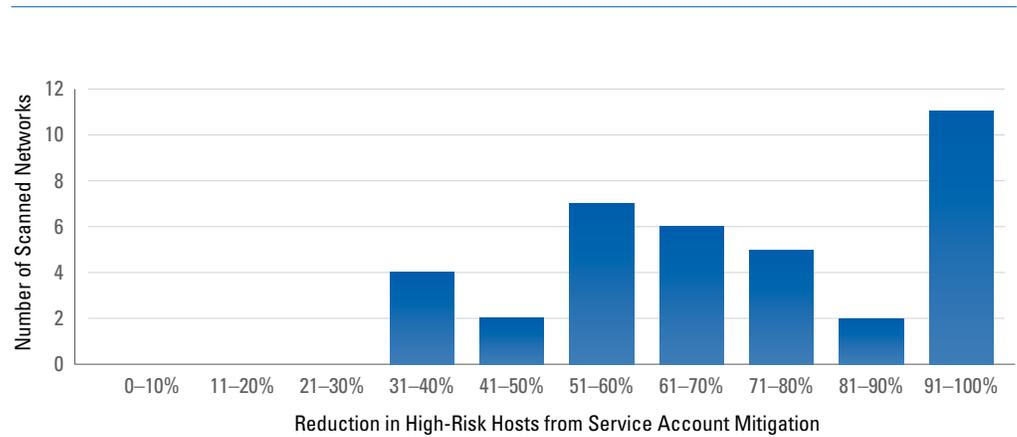
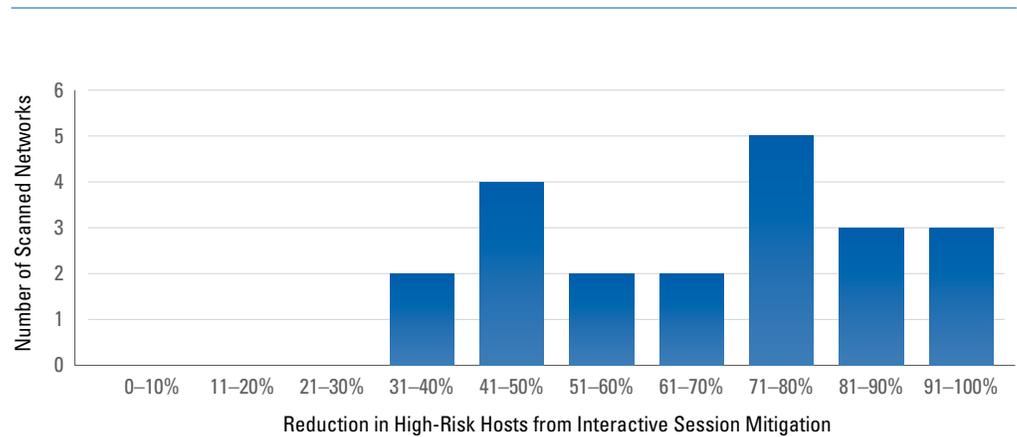


Figure 9. **Reduction in High-Risk Hosts from Interactive Session Mitigation**



To quantify the effectiveness of these mitigation options, we simulated their application to the analyzed networks and compared the new results to the original ones. Figure 8 shows the reduction in high-risk hosts achieved by applying the suggested mitigations for high-risk automated process accounts to the 37 networks at greatest risk from these accounts. It is important to note that Figure 8 represents those networks at greatest risk from service accounts, which are not necessarily the networks with the largest percentages of service accounts because one high-risk account may pose considerably more risk than another. Among the 37 networks represented in Figure 8, the mitigations reduced the number of high-risk hosts on each network between 33 and 100 percent, with an average decline of 71 percent.

Figure 9 shows the reduction in high-risk hosts achieved by applying the suggested mitigations for high-risk interactive session accounts to the 21 networks at greatest risk from these accounts. The average decline in the high-risk hosts per network was 67 percent, with a range of 33 to 100 percent.

Because of the dramatic reductions depicted in Figures 8 and 9, many high-exposure networks became medium or low-exposure networks, indicating a major improvement in their overall security. Applying all of the mitigations for both interactive sessions and automated processes throughout a network would effectively lower the exposure significantly below the 10 percent threshold, making all networks low-risk.

Summary

Our research has clearly demonstrated that nearly every organization is at significant risk of compromise through Windows privileged account credential theft and reuse. On average, 40 percent of the Windows hosts on a given network, if compromised, would provide an attacker credentials that would facilitate complete compromise of the vast majority of the other Windows hosts on that network – whether directly or through a series of compromises. This metric, the percentage of high-risk machines, is useful in comparing the exposure of networks and in evaluating the effectiveness of mitigations.

Every network, regardless of how many hosts it has or how many domain accounts it has, has hosts and accounts that can be used to compromise most of its Windows hosts through privileged account credential theft and reuse. In our study, 88 percent of scanned networks were significantly susceptible to this form of attack. Fortunately, there are several mitigation options available that can reduce the likelihood and potential impact of widespread privileged account credential theft and reuse.

Organizations that are concerned about the risks posed by their Windows privileged accounts are encouraged to adopt the approach we took in this research to identify and assess their own organization's potential vulnerabilities, then reduce the associated risks through one or more mitigations.



CYBERARK®

ACKNOWLEDGMENTS

Our thanks to all of the organizations that shared the data from their networks and made this research possible.

We would like to thank our colleagues who have assisted in collecting and analyzing the data and provided remarks on the analysis.

ABOUT CYBERARK LABS

CyberArk Labs is a team of cyber security experts who conduct research focused on targeted attacks against organizational networks – the methods, tools and techniques employed by targeted attackers, as well as methods and techniques to detect and mitigate such attacks.

ABOUT CYBERARK

CyberArk (NASDAQ: CYBR) is the only security company focused on eliminating the most advanced cyber threats; those that use insider privileges to attack the heart of the enterprise. Dedicated to stopping attacks before they stop business, CyberArk proactively secures against cyber threats before attacks can escalate and do irreparable damage. The company is trusted by the world's leading companies – including 40 percent of the Fortune 100 and 17 of the world's top 20 banks – to protect their highest value information assets, infrastructure and applications. A global company, CyberArk is headquartered in Petach Tikvah, Israel, with U.S. headquarters located in Newton, Mass. The company also has offices throughout EMEA and Asia-Pacific. To learn more about CyberArk, visit www.cyberark.com

US HEADQUARTERS

CyberArk

60 Wells Avenue
Newton, MA 02459
1-888-808-9005
or (617) 965-1544

All rights reserved. This document contains information and ideas, which are proprietary to CyberArk Software Ltd.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, without the prior written permission of CyberArk Software Ltd.

©2015 CyberArk Software Ltd. | www.CyberArk.com