WHITE PAPER | NOVEMBER 2015

Breaking the Kill Chain

Stopping Data Breaches with Privileged Access Management

ca technologies

Executive Summary

Challenge

It's impossible for a day to pass in which we don't hear news of yet another data breach, with its resulting loss of proprietary secrets, financial records or personal information. These incidents span all sectors of the economy: commerce, education and government. Already an annual drag on the worldwide economy accounting for hundreds of billions of dollars a year in costs,[1] without immediate and aggressive action it's projected the bill for cybercrime will mount to the trillions of dollars in less than a decade.[2] Beyond reckoning is the devastating impact to individuals who have suffered the compromise of the most intimate details of their personal lives.

Security specialists have striven to establish perimeter-based defenses that, in the most simplistic of terms, keep the bad guys out and let the good guys in. The never-ending string of breaches we're witnessing offers prima facie evidence these perimeters have failed at their primary goal. As a consequence organizations are coming to grips with the realization an essential new layer of security, focused specifically on the protection and management of identities, is a critical new requirement in efforts to stem the tide of breaches. Of these identities, none are so critical as those belonging to privileged users. By providing the "keys to the kingdom," the theft and exploit of these credentials increasingly serves as the principal attack vector in breach after breach.

Opportunity

Security teams have at their disposal a mature set of technologies and processes, broadly termed "privileged access management," providing the means to defeat and deter attackers. Malicious users, both internal and external, predictably follow a logical series of steps in order to successfully carry out their attacks. These sequences, originally identified and articulated by cybersecurity teams at Lockheed Martin,[3] have come to be called "kill chains" based on the fact that if the sequence of steps attackers follow can be interrupted, or "killed" at any point, the ultimate attack can be prevented or mitigated. Privileged access management provides the means to thwart attackers at multiple steps in the attack lifecycle. In this paper, we'll examine a somewhat simplified version of a kill chain and provide concrete example of how privileged access management can help stop attacks and protect organizations from breaches.

Benefits

The financial benefits of preventing these breaches are obvious. Harder to measure, but often more impactful, are the "soft" costs associated with damage to brand and reputation, loss of trust among partners and customers and impacts on company market valuations. However, as significant as these costs are, they pale in comparison to the devastating impact the theft of detailed personal information can have on unsuspecting, trusting individuals. Clearly, privileged access management, with its ability to mitigate these wide-ranging damages, is of immense benefit.

Challenge–Data Breaches: Escalating Risks and Incalculable Damages

When we consider the current stream of security incidents, it's common to refer back to the Target incident, which began in late 2013. With some 70 million payment card records stolen, Target wasn't the first—or even largest—data breach to take place in history, let alone in 2013. However, for a variety of factors, the Target breach did serve to galvanize the interest of multiple critical constituencies in the intensely damaging nature of these ongoing attacks. Since the Target breach, we've witnessed countless smaller, less publicized breaches, along with a continuing string of larger-scale incidents, such as the Home Depot and JP Morgan Chase incidents about one year later and most recently, as of this writing, the Experian breach of the extremely sensitive personal data of some 15 million T-Mobile customers.

> "For digital businesses, privileged identity management becomes both incredibly important and challenging. It's important because one administrator with malicious intent or the theft of administrator credentials can have a disastrous effect on your customers, revenues and long-term reputation."
>
> —Forrester Research[4]

The bill for such cybercrime in 2014 was estimated at around $400 billion according to data from Intel Security and the Center for Strategic and International Studies. It can be difficult to truly grasp such gargantuan numbers, so for perspective, contrast that $400 billion price tag with the estimated take from worldwide drug trafficking, which is somewhat dwarfed in comparison, clocking in at an estimated $300 billion annually. Cybercrime's impact is even greater than the GDP of many prosperous countries, for example, Singapore, which also, coincidentally, stands at about $300 billion a year. This is clearly a huge financial issue, and the data suggest, that absent prompt action, it's only going to get worse, with McKinsey projecting a global annual impact from cybercrime of $3 trillion in 10 years orders of magnitude beyond today's costs.

Now clearly, these are damaging incidents. The organizations that suffered the breaches and others like them lost market capitalization, lost sales, lost customer goodwill and lost profits. And that's all before the financial and emotional damage to the individuals affected by these breaches, from the ravages of crimes such as identity theft. However, as troubling as all of that is, there is even worse news if you examine other incidents that have begun to occur more recently.

First, we've begun to see attacks that are clearly aimed at causing a material impact to the operations of targeted organizations. You may not be familiar with Code Spaces, but it was a smaller, UK-based business focused on providing cloud-based version control and backup services to developers. In June of 2014, an attacker was able to gain administrative credentials to Code Space's Amazon Web Services (AWS) management console. After creating multiple accounts and backdoors, the attacker presented Code Spaces management with a ransom demand. By the time authorized administrators attempted to roust the attacker from their system, it was too late. The attacker, with full administrative access to Code Space's entire management system, began to retaliate by rapidly destroying the company's complete information technology infrastructure—servers, applications, and critically, system and data backups. The attack was

complete in a matter of hours. The company was forced to cease operations in mere days.[5] The Code Spaces breach is a dramatic example, but there are a number of other examples (such as the Sony Pictures Entertainment and Saudi Aramco incidents) that illustrate this trend.

From there, it's only a short step to the latest trend, which we're told is cyber-espionage. Among the first signs of these attacks were breaches at insurance companies, including Anthem, Premera and CareFirst in early 2015. While the breaches were never formally blamed on nation-states, widespread speculation holds the thefts of millions of records of personal data were part of a larger campaign to assemble dossiers on individuals holding sensitive positions within government, defense contractors, finance and telecommunications, along with geopolitical policy makers and others. The timing of the breaches aligned with the distribution of a confidential U.S. FBI alert that Chinese hackers were targeting personally identifiable information from U.S. commercial and government networks.[6] Since then, of course, we've learned of the breach of the U.S. Office of Personnel Management (OPM), in which personal data, including the extensive biographical, financial, employment and personal histories of individuals seeking security clearances, were stolen.

Opportunity—Privileged Accounts: The Emerging Front Line

Let's pause for a moment and summarize. Data breaches are a big problem and getting bigger. The stakes are higher and higher and we face ever more sophisticated—and well-financed—adversaries. Even the most optimistic reader would be excused for feeling at least a trace of pessimism at this point. In the face of such a significant challenge, what can we possibly be expected to do to combat it?

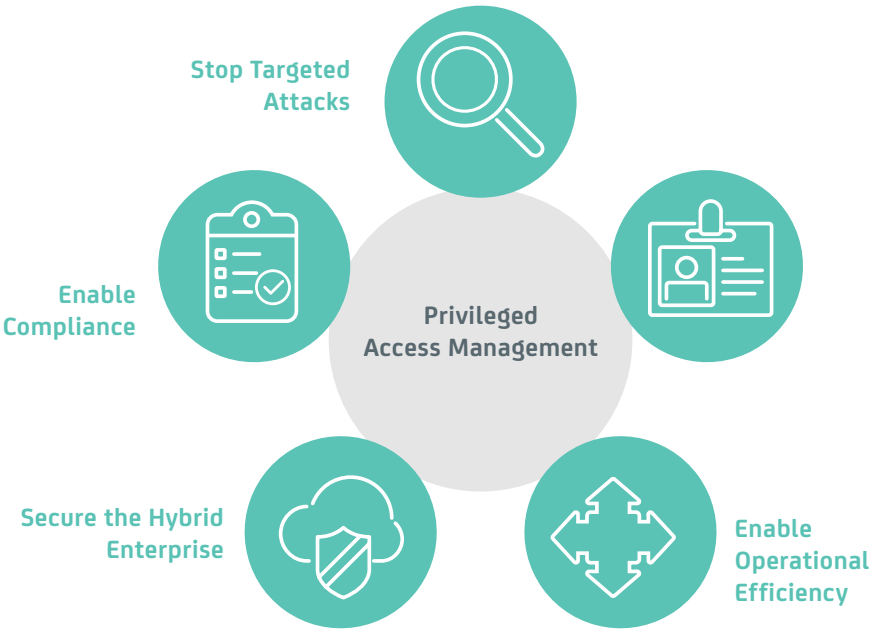

Figure A.

Privileged Access Management aids organizations in achieving five high level goals.

The good news is, there is significant reason for hope because there is a common thread observed in virtually all of these attacks. That common thread is privileged users and, more specifically, the privileged accounts and credentials those individuals use to configure, maintain and operate our information technology infrastructure. Stealing and exploiting these credentials, providing privileged access to IT infrastructure, has been shown to be a critical success factor and a primary attack vector for attackers in all of the breaches we've discussed so far.

> "By 2018, the inability of organizations to properly scope and contain privileged access will be responsible for up to 60 percent of insider misuse and data theft incidents, up from more than 40 percent today."
>
> —Gartner[7]

Before examining the central role that privileged access plays in successful breaches, it's useful to briefly examine just who these privileged users are, since the population of individuals with privileged access, as well as the number of actual accounts and credentials used to provide that access, is much larger than is commonly recognized.

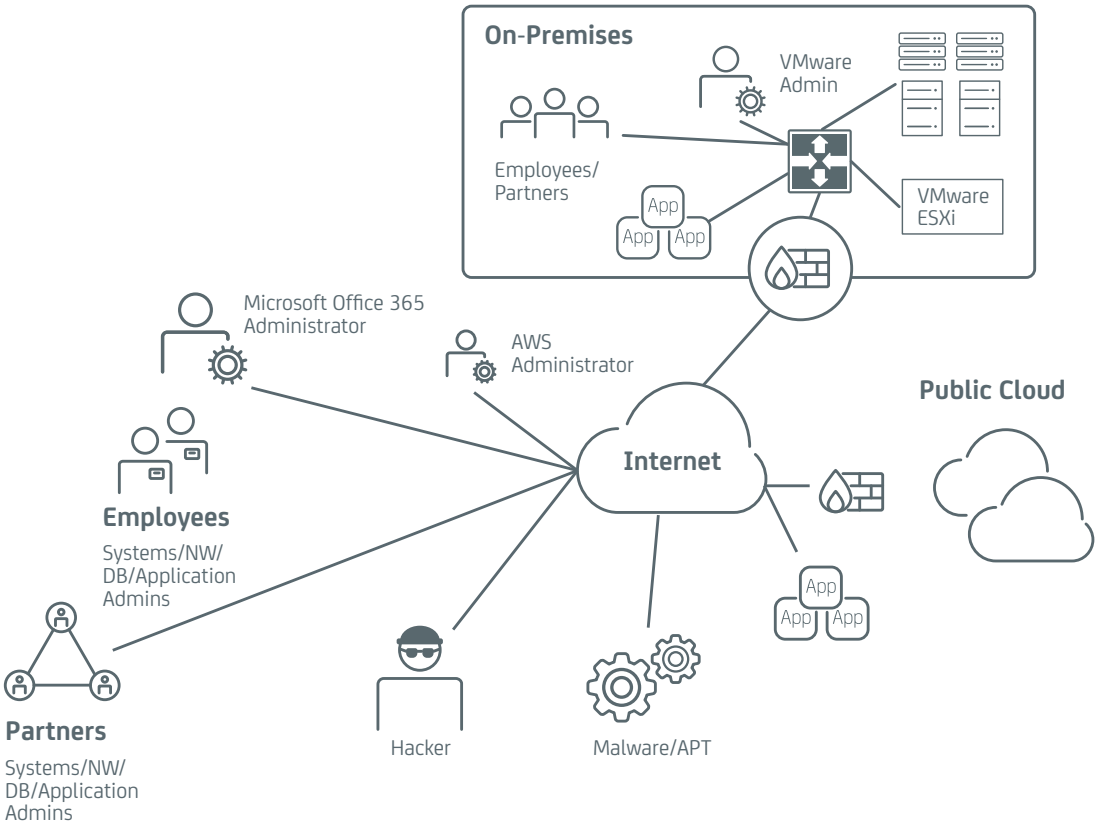

Figure B.

Privileged Accounts: The emerging front line

For years, when we've thought about privileged users, we've generally only considered those people inside the organization with direct, hands-on responsibility for system and network administration. That leads many to minimize the risk and view the challenge of privileged access management as one of controlling the so-called "insider threat." And while it's true malicious insiders can cause outsized damages, such incidents are relatively rare and account for a small number of breaches.

The reality is that many privileged users aren't insiders, they're vendors, contractors, business partners and others who have been granted privileged access to systems within the organization. In many companies, the number of such third-party users may well outnumber traditional "insider" privileged users. Experience also suggests these third parties are a greater source of risk. Consider the breaches we've mentioned—including Target, Home Depot and the OPM incidents, among others—where an authorized third-party user's credentials were compromised and then used for illicit access to the broader network and its resources

In addition, the number of privileged users has been increasing as we move to the cloud and embrace technologies like virtualization. Particularly when looking at the cloud, many of these privileged users may not actually be members of the traditional IT staff. As an example, consider the case of line-of-business representatives purchasing service-based offerings where, in the worst case, the traditional IT and security organizations may be completely unaware of the exposure.

And let's not forget that increasingly, many privileged users aren't actually users—or at the very least, they're not people. In cloud and virtualized environments, the emergence of automated configuration and provisioning tools driven by scripts and programs has introduced even more "users" with significant access to and authority over large tracts of infrastructure. A close corollary to these automated systems are the often uncounted number of scripts and programs assembled over years of operations which require administrative or sensitive access to resources like databases or other applications and systems. In both cases, this access and these operations are, quite properly, controlled by authentication. Unfortunately, the required credentials are typically hard-coded into applications or configuration files where they're easy targets for malicious users—insiders and outsiders alike.

Finally and perhaps most importantly, we have to remember we're not just talking about privileged users but, rather, all the privileged accounts and the credentials that go with them that exist in the typical organization. It's those credentials that pose the most significant threat because exploiting them is critical to how attackers carry out their breaches.

Introducing the Kill Chain—And Why It Works

The breach kill chain is comprised of a consistent and predictable series of steps an attacker must accomplish to successfully achieve their goal. While the articulation of some kill chains can be quite complex, it's possible to summarize the key steps associated with the typical data breach kill chain in a simplified manner.

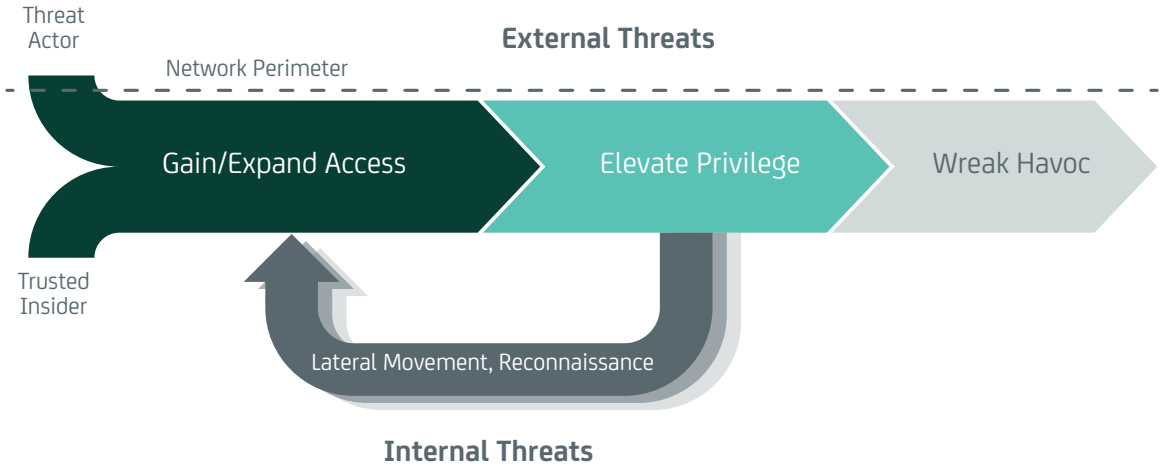

Figure C.

An example of a simplified kill chain

There are four key steps:

- **Gain Access:** First, it's necessary to gain access to the network. If you're a true insider or perhaps a trusted third-party, it's easy—you simply exploit the credentials and access you already enjoy. But it's not that much harder for an attacker to do the same thing. The growing popularity of social sites such as LinkedIn makes it relatively easy to identify and target specific individuals within an organization who likely have privileged access to systems. The growing sophistication of spearphishing means that it's easier than ever for an attacker to fool even the most experienced and skilled individual into handing over credentials—especially relatively unsophisticated ones like user ids and passwords.
- **Elevate Privileges:** Once an attacker has gained access, one of the first steps is to elevate privileges, typically by compromising other privileged credentials. This step supports two essential activities. First, it allows the attacker the ability to take steps—like altering or disabling logging or installing malware—that help prevent the discovery of their existence and activity. Secondly, it sets the stage for the next step in the kill chain, reconnaissance and lateral movement.
- **Perform Lateral Movement and Reconnaissance:** Unless you're a remarkably lucky attacker, the first system you've gained access to is unlikely to be your ultimate target. Your goal—payment card processing systems, proprietary data, personnel records and the like—are almost certainly located elsewhere in the network on other systems. So the next step in the kill chain is to conduct reconnaissance of the network and move to systems and servers closer to your ultimate goal.
- **Repeat The Process As Needed:** From here, it's simple—just keep repeating the process until you reach your ultimate goal, whatever it might be. Again, experience has shown that attackers can be remarkably patient, taking time to research and navigate networks to carry out their mission. Public breach reports routinely indicate attackers have been at work inside a victim's network for months or years in some cases. Once they finally reach their target, they can then carry out their attack—disrupting systems, stealing data or more.

Unfortunately and especially in the absence of even rudimentary privileged access management processes and tools, there are a number of things organizations do that make it easier for attackers to execute this kill chain. Common mistakes include:

- **Using weak authentication** techniques for access to the network or specific resources, including the failure to eliminate default administrative accounts and passwords and relying on unsophisticated credentials, like simple user id/password combinations, that are easily stolen or compromised.
- **Poor password and key management**, where credentials are not changed on a frequent and routine basis. In organizations with resources that may number into the thousands, this can be extremely problematic, since it's tempting to avoid operational problems and reduce overhead by engaging in poor practices like credential re-use and failure to rotate credential on a routine basis.
- **Allowing for the use of shared accounts**, especially powerful privileged accounts like root or admin. This practice introduces multiple risks, since it's easy for a user to share a credential with others, and the fact that many individuals have access to a given credential makes attribution proving that a specific individual performed a given task on a system virtually impossible, complicating forensics and troubleshooting.
- **Equating authentication with access control.** Many networks are poorly segmented and as a result, once an individual is in the network, they have visibility to many more resources than are necessary or prudent. That facilitates reconnaissance and lateral movement efforts, making it easier for an attacker to reach their ultimate goal.
- **Lack of monitoring and analysis of privileged user activity** can foster multiple problems. If activity isn't monitored or routinely analyzed, suspicious or suspect behavior can be missed, allowing attackers free reign. And it's human nature for people to bend or even break the rules if they know there is little likelihood their behavior will be detected.

Recommendations: Breaking the Kill Chain

Grouped at a broad level into three key steps, privileged access management provides multiple means of breaking the kill chain, stopping attackers and preventing breaches.

Step One—Preventing Unauthorized Access

Forcing privileged access to resources through a network-based gateway provides a simple way to enforce strong authentication. It's a given that such a system should integrate with existing identity management infrastructure. So the system should support links to existing identity stores, like Active Directory or LDAP directories or even RADIUS or TACACS+ in some environments. While the system can and should support local authentication, usually your organization will already have a well established identity store in place. Since those systems already define both authorized users as well as roles and permissions, you'll want to leverage that data as the basis for privileged access.

But, this is merely a baseline. With the relative ease of stealing an authorized user's credentials, passing through such a gateway can be a relatively simple step for an attacker. To prevent this, it's critical to mandate the use of multi-factor authentication (MFA) for privileged access. The addition of MFA significantly increases the level of difficulty for an attacker seeking to gain access to the network. At one time, MFA was an expensive, administratively cumbersome technology. However, advancements

in technology have drastically changed the economics of implementing multi-factor authentication technologies and, given the high level of risk associated with privileged access, even a rudimentary cost-benefit analysis will support its implementation.

The use of multi-factor authentication has also become a compliance and audit issue. The U.S. Federal Government took an early lead in this with the creation of standards mandating the use of so-called PIV/CAC cards for administrative access to systems. Short for privileged identity verification (for civilian agencies) and common access card (a similar device used within the military). These cards provide for PKI-based identification of an individual which, when combined with identity-proofing, provide a high level of assurance of a user's identity. Similar standards have also been added to a variety of compliance mandates, including, for example, the most recent revision to the Payment Card Industry Data Security Standard (PCI-DSS).

Other common sense measures that can be used to mitigate the risk of unauthorized access include restrictions on access to systems, based on a user's login source IP address or the time of day. These kinds of controls can be implemented both through a privileged access management gateway as well as through agent-based controls on specific servers or resources. If a given user routinely logs in during a certain time period or from a given set of locations, there's no reason to allow unrestricted access. And you may wish to completely block logins from a range of addresses from which access would never be expected or acceptable.

A second aspect of this problem is protecting the credentials used to actually access managed systems. As we've already outlined, it's all too routine for these credentials to be poorly protected, shared indiscriminately or managed poorly—presenting obvious risks. Ideally, a privileged access management system will provide a credential safe where passwords and key pairs can be stored, encrypted, away from prying eyes and malicious users. The credential safe must support the ability to actually actively manage credentials, interacting with systems to change passwords based on standards appropriate to an organization's or resource's level of risk. Automating this process decreases both security risks (since it's possible to routinely update credentials on thousands or even hundreds of thousands of resources while keeping the credentials out of harms way), as well as operational risks since automated password and key updates are less error prone. When combined with privileged user single sign-on, a high level of security can be achieved since it's possible to provide a user with access to a system without the need to actually provide them with access to the relevant credentials. And if a user doesn't have a credential, he can't steal it, share it or be tricked into giving it up to an attacker.

Step Two—Limit Privilege Escalation, Reconnaissance and Lateral Movement

This provides a transition to the next step in breaking the kill chain: limiting the ability of a user to conduct reconnaissance of the network and to move about. Unfortunately, in many networks, authentication ends up being essentially the same thing as access control—once you've logged into the network, you frequently have access to resources across the entire network. If you're an attacker, that's obviously great news; you've got the time and frequently the means, to move from system to system, moving closer to your target.

Capabilities like privileged user single sign-on help prevent all of these challenges. The approach is fundamentally based on least privilege access control, called zero trust access control. By separating authentication and access to the privileged access management system and actual access to managed resources, users only have visibility to those systems and resources as defined and permitted by policy. If a given user's work responsibilities require access to a single server or a class of resources, that's all they should be able to see in the network. And by proxying or brokering sessions between the privileged access management system and managed resources, it's possible to limit the authority they have over a system and control the commands they're able to issue, further restricting the ability to escalate privileges or move laterally within the network.

For example, with a proxied session, it's possible to log a user onto a system using a standard account, even a powerful one like root. Since the system can enforce command filters, it's possible to limit the individual to specific commands or to ban unauthorized ones. For example, a user may be assigned the job of updating software on a set of servers and it may be necessary to be logged in as root to perform the job. With command filters, it's possible to log the user in and permit just those commands needed to perform the job. Others, such as trying to kill a process or reboot the system, can be prevented.

Additional controls allow for variable responses to attempts to violate the policies. Say a user issues an unauthorized command—your policies might assume the action was the result of some innocent need or merely a mistake. In such a case, a warning could be issued to the user and the command prevented. Repeated attempts or more serious infractions, might prompt terminating a session or even deactivating a user's account until an administrator has the opportunity to review the incident in greater detail.

The addition of host-based agents allow for similar capabilities, but often with much finer-grained controls, such as the ability to tightly restrict access to files and directories or monitor files for modifications. It's also possible to prevent attempts to move laterally within the network. For example, having gained access to a system, an attacker might attempt to issue an SSH or TELNET command or open a remote RDP session to a target system. Once again, the privileged access management system can examine policies and determine whether the activity is permitted. If not, the command is prevented and the attempted violation is logged.

Step Three—Monitor, Record and Audit Activity

Ideally, our attacker will never get to the point where he's able to achieve his final goal; the many controls and checks established and enforced by a privileged access management system provide ample opportunities to break the breach kill chain. The final step of monitoring, recording and auditing activity acts as an additional deterrent to breaches as well as providing significant benefits in the event a breach is ultimately successful.

As we noted, knowing your activity is being recorded and analyzed can act as a powerful deterrent to misbehavior or seemingly innocent, but potentially dangerous, exploration and examination of systems. And extensive logging, alerting, recording and reporting capabilities provide for an "early warning system" alerting other administrators, managers and auditors of suspect or unusual behavior. Alerts and events provide immediate warnings of policy violations and attempts at a breach, enabling a rapid response. Logs can analyzed, either individually or via a log management or SIEM system in the context of other system activity to provide further clues to suspect events, enabling investigation before a breach occurs.

Because shared administrative accounts are so commonly used, the ability to support attribution of actions taken using such an account back to a specific individual is a critical requirement for compliance requirements.

Finally, session recording offers a number of benefits. Administrators sometimes make mistakes. Session recordings can be helpful in such cases, since they allow the ability to review activity and determine precisely what actions were taken during an interaction. That can speed troubleshooting, such as when a problem is noted with a system. If an update was performed or a configuration change was made on a previous shift, it can be difficult and time-consuming to determine exactly what happened. Session recordings provide immediate playback, speeding recovery. And they can also be used for training purposes, making it easier to point out where a mistake was made and the preferred course of action.

Of course, in the very worst case, in the event of a successful breach, such recordings and logs can be crucial in determining exactly what was done to a system, what information was taken and how the resource was compromised. All of which speeds forensics, helps in damage assessment and provides valuable information that can be used to mitigate the risk of future breaches.

Benefits

Unfortunately, data breaches—with all their attendant costs and damages—are a fact of life. However, as we've demonstrated here, attackers routinely follow a defined, predictable course of action in attempting to carry out these attacks. Privileged access management provides a host of capabilities and controls that actively prevent attackers from carrying out key components of their attacks—breaking the breach kill chain—as well as delivering additional support for reducing risks, minimizing damage and speeding recovery in the event of a successful attack. Implementing a comprehensive privileged access management solution, provides the following benefits:

- **Reduce risk.** Prevent unauthorized access and limit access to resources once entry is granted to the network. Protect passwords and other credentials from unauthorized use and compromise. Limit the actions users can perform on systems and prevent the execution of unauthorized commands and prevent lateral movement within the network.
- **Increase accountability.** Observe full attribution of user activity, even when using shared accounts. Comprehensive logging, session recording and user warnings capture activity and provide a deterrent to unauthorized behavior.
- **Improve auditing and facilitate compliance.** Simplify compliance by providing support for emerging authentication and access control requirements and limit the scope of compliance requirements through logical segmentation of the network.
- **Reduce complexity and boost operator productivity.** Privileged single sign-on not only helps limit risk, but it can boost the productivity of individual administrators by making it easier and faster for them to access the systems and resources the need to manage. Centralized policy definition and enforcement simplify the creation and enforcement of security controls.

Conclusions

- Privileged identity, accounts and credentials are core, critical assets for enterprises that must be highly protected through a combination of technology and processes which are enabled by privileged access management.
- Delivering that protection is instrumental in breaking the data breach kill chain, helping to prevent attacks and mitigating the impact of those that do occur.
- A zero trust access control model is essential for all types of privileged access, both human and programmatic.
- While perimeter-based approaches to security have been revealed to have serious shortcomings, defense-in-depth is still an essential strategy for protecting resources. Privileged access management is capable of providing multiple additional layers of defense around privileged users, accounts and credentials—both at the network and host layers.
- Given the prevalence of breaches and the sophistication of attackers, it's extremely tempting—and we're often encouraged—to focus solely on detection and response to breaches. That's a mistake. While they're essential activities, it's critical to remember that privileged access management can help organizations substantially improve their ability to prevent breaches in the first place.

About the Author

Dale R. Gardner has over two decades of experience with enterprise software, spanning network and systems management and multiple segments of security, including identity management, application security, vulnerability management, compliance and network security. A former research analyst and writer, he has defined, built and marketed multiple management and security solutions that enhance operations and help ensure the integrity and reliability of enterprise information technology infrastructure. He currently is responsible for the worldwide marketing of CA Technologies privileged access management product portfolio.

Connect with CA Technologies at ca.com

CA Technologies (NASDAQ: CA) creates software that fuels transformation for companies and enables them to seize the opportunities of the application economy. Software is at the heart of every business, in every industry. From planning to development to management and security, CA is working with companies worldwide to change the way we live, transact and communicate – across mobile, private and public cloud, distributed and mainframe environments. Learn more at **ca.com**.

1 Intel Security and the Center for Strategic and International Studies, "Net Losses: Estimating the Global Loss of Cybercrime, Economic Impact of Cybercrime II," June 2014, http://www.mcafee.com/us/resources/reports/rp-economic-impact-cybercrime2.pdf

2 World Economic Forum and McKinsey & Company, "Risk and Responsibility in a Hyper-connected World," January 2014, http://www3.weforum.org/docs/WEF_RiskResponsibility_HyperconnectedWorld_Report_2014.pdf

3 Eric M. Hutchins, Michael J. Cloppert, Rohan M. Amin, Ph.D., Lockheed Martin Corporation, "Intelligence-Driven Computer Network Defense Informed by Analysis of Adversary Campaigns and Intrusion Kill Chains," http://www.lockheedmartin.com/content/dam/lockheed/data/corporate/documents/LM-White-Paper-Intel-Driven-Defense.pdf

4 Andras Cser, Forrester Research, "Critical Questions to Ask Your Privileged Identity Management Solution Provider," 10 September 2014.

5 Ars Technica, "AWS console breach leads to demise of service with 'proven' backup plan," June 18, 2014, http://arstechnica.com/security/2014/06/aws-console-breach-leads-to-demise-of-service-with-proven-backup-plan/

6 Brian Krebs, "China To Blame in Anthem Hack?," February 15, 2015, http://krebsonsecurity.com/2015/02/china-to-blame-in-anthem-hack/

7 Anmol Singh and Felix Gaehtgens, "Twelve Best Practices for Privileged Access Management, Gartner," October 8, 2015, G00277332

 CS200-162063_1115