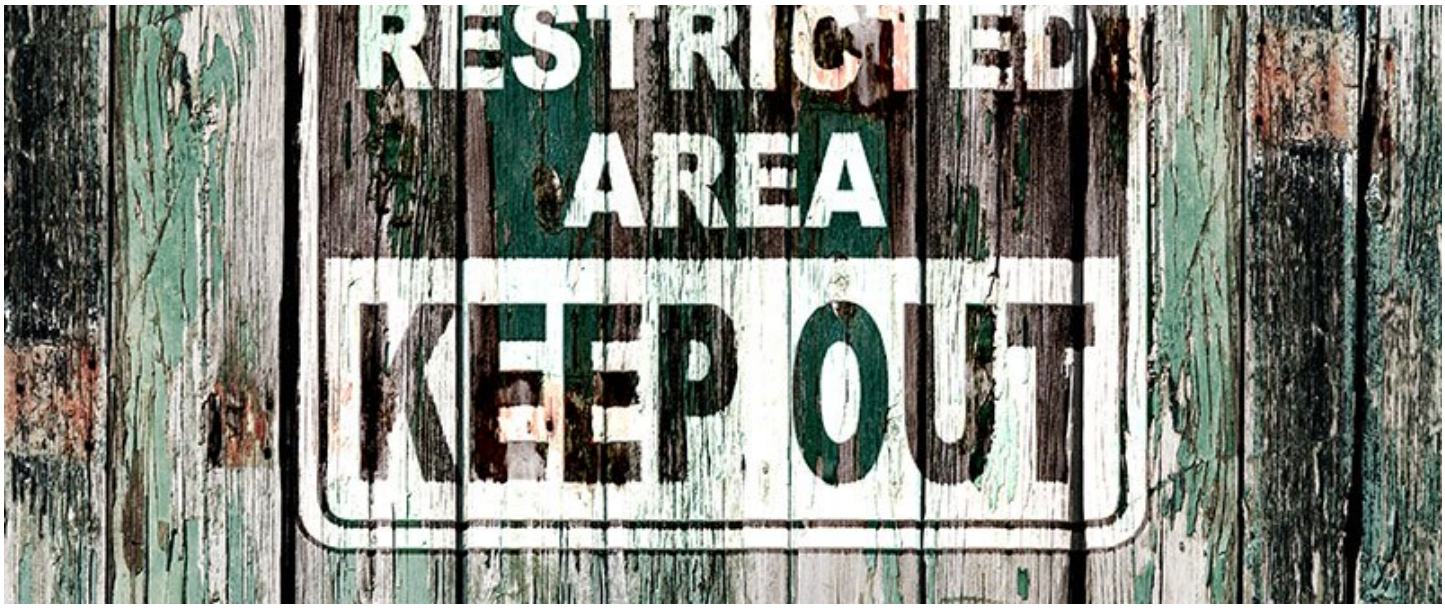

Least privilege is the concept and practice of restricting access rights for users, accounts, and computing processes to only those resources absolutely required to perform routine, legitimate activities. Privilege itself refers to the authorization to bypass certain security restraints. Applied to people, least privilege, sometimes called the principle of least privilege (POLP), means enforcing the minimal level of user rights, or lowest clearance level, that allows the user to perform his/her role. However, least privilege also applies to processes, applications, systems, and devices (such as IoT), in that each should have only those permissions required to perform an authorized activity.

While this blog will focus on the cyber security context of least privilege, no doubt you're familiar with analogous concepts, such as "need to know" popularized amongst military and governmental circles. In fact, adoption of "least privilege" was advanced by the publication of the "Department of Defense Trusted Computer System Evaluation Criteria" in 1985, following the recommendations of a task force dedicated to safeguarding classified data.

In the face of fast-emerging and evolving technology areas, like IoT, shadow IT applications run from the cloud, and more, least privilege remains highly relevant. In fact, Forrester Research estimates 80% of today's security breaches involve privileged credentials. Enforcing least privilege is a best practice that is instrumental in reducing security risk and minimizing business disruption that may result from errors or malicious intent.

While the concept of least privilege is relatively simple to understand, it can be complex to effectively implement, especially, when you consider the many variables, such as:

- heterogeneous systems (Windows, Mac, Unix, Linux, etc.)
- the expanding number and types of applications and endpoints (desktops, laptops, tablets smartphones, IoT, etc.)
- diverse computing environments (cloud, virtual, on-prem, hybrid)
- the many different types of user roles

- third-party / vendor access.

This blog aims to provide an overview of least privilege, and will cover: some types of computing privileges, privileged and non-privileged accounts, privileged threat vectors and attacks, challenges to applying a least privilege model, best practices and strategies for implementing least privilege, and the cornerstone technologies for enabling a least- privilege computing environment.

Privileged & Non-Privileged Accounts: An Abridged Overview

So, how do users acquire privileges? Depending on the system, some privilege assignment, or delegation, to people may be based on attributes that are role-based, such as business unit, (i.e. marketing, HR, or IT) as well as a variety of other parameters (seniority, time of day, special circumstance, etc.). Additionally, various operating systems provide different default privilege settings for different types of user accounts.

Superuser accounts, primarily used for administration by specialized IT employees, may have virtually unlimited privileges, or *carte blanche*, over a system. Superuser account privileges can include full read / write / execute privileges, and the power to render systemic changes across a network, such as creating or installing files or software, modifying files and settings, and deleting users and data.

Standard user accounts, sometimes called least privileged user accounts (LUA) or non-privileged accounts, have a limited set of privileges. In a least privilege environment, these are the type of accounts that most users should be operating in 90 – 100% of the time.

While most non-IT users should, as a best practice, only have standard user account access, some IT roles (such as a network admin) may possess multiple accounts, logging in as a standard user for routine tasks, while logging into a superuser account to perform administrative activities. Because administrative accounts possess more privileges, and thus, pose a heightened risk compared to standard user accounts, a best practice is to only use these administrator accounts when absolutely necessary, and for the shortest time needed.

Privileged Accounts in Unix, Linux, Windows, and OS X platforms

In Linux and Unix-like systems, the superuser account, called ‘root’, is virtually omnipotent over the system, with unrestricted access to all commands, files, directories, and resources. Root can even grant and revoke any permissions for other users! If misused, either in error (such as accidentally deleting an important file or mistyping a powerful command), or with malicious intent, these root accounts can easily wreak catastrophic damage to a system / organization.

Because of the immense potential for destruction, if misused or abused, inherent to root privileges, IT admins should only log into and assume root account privileges when

absolutely necessary. Sometimes, this is done by leveraging the *sudo* command, which allows the user to temporarily elevate privileges to root-level, but without having direct access to the root account and password.

Standard, "non-privileged" Unix and Linux accounts lack access to *sudo*, but still retain minimal default privileges, allowing for basic customizations and software installations.

In Windows systems, the Administrator account holds superuser privileges. Each Windows computer has at least one administrator account. The Administrator account allows the user to perform such activities as installing software and changing local configurations and settings. Standard users have substantially curtailed privileges, while guest user accounts are generally limited even further, to just basic application access and internet browsing.

Mac OS X, on the other hand is Unix-like, but unlike Unix and Linux, is rarely deployed as a server. However, Mac OS X endpoints are increasingly prevalent at enterprises. As a default, Mac users run with root access, though, as a best security practice, a non-privileged account should be created and used for routine computing to limit the likelihood and scope of privileged threats.

Privileged Accounts and the Cloud

While cloud and virtualized environments provide many benefits, chief among them, rapid scalability, many traditional security tools are architected for on-premise environments, and when extended to the cloud or across hybrid (on-prem to cloud, public to private cloud, etc.) environments, leave gaps that allow for excessive privileged access and permissions. Cloud and virtualization have also ushered in administrator consoles (such as with AWS and Office 365) that confer substantial super user capabilities, enabling users to easily provision, configure, and delete servers at incredible scale. While it's just a matter of a few simple clicks to spin-up and manage thousands of virtual machines (each with its own set of privileges and privileged accounts) from a single console, it can be complex to onboard and manage all of these instantly created privileged accounts.

Common Privileged Threat Vectors

Hackers, malware, partners, insiders gone rogue, and simple user errors—especially in the case of superuser accounts, round-out the most common privileged threat vectors as shown in this infographic.

Let's delve into how some of these vectors play out and also review a couple privileged threat exploits that rocked the world and underscored the need for least privilege.

Poor Computing Hygiene + Unmanaged Privileges = Opportunities for Exploits

Routine computing for employees or personal PC users might entail internet browsing, watching streaming video, use of MS Office and other basic applications, including SaaS (SalesForce, GoogleDocs, etc.). In the case of Windows PCs, users often log in with

administrative account privileges—far broader than what is needed. These excessive privileges massively increase the risk that malware or hackers may steal passwords or install malicious code that could be delivered via web surfing or email attachments. The malware or hacker could then leverage the entire set of privileges of the account, accessing data of the infected computer, and even launching an attack against other networked computers or servers. For example, should a standard user click on an attachment or link within a phishing email that consequently loads malware onto his system, the impact of that act alone would be fairly isolated to the user's own systems and the limited resources to which he has access. However, if that user was logged in as a superuser with broad admin privileges, the malware could leverage domain account privileges to modify settings, corrupt, and access sensitive data from other endpoints and servers across the network, with deleterious implications.

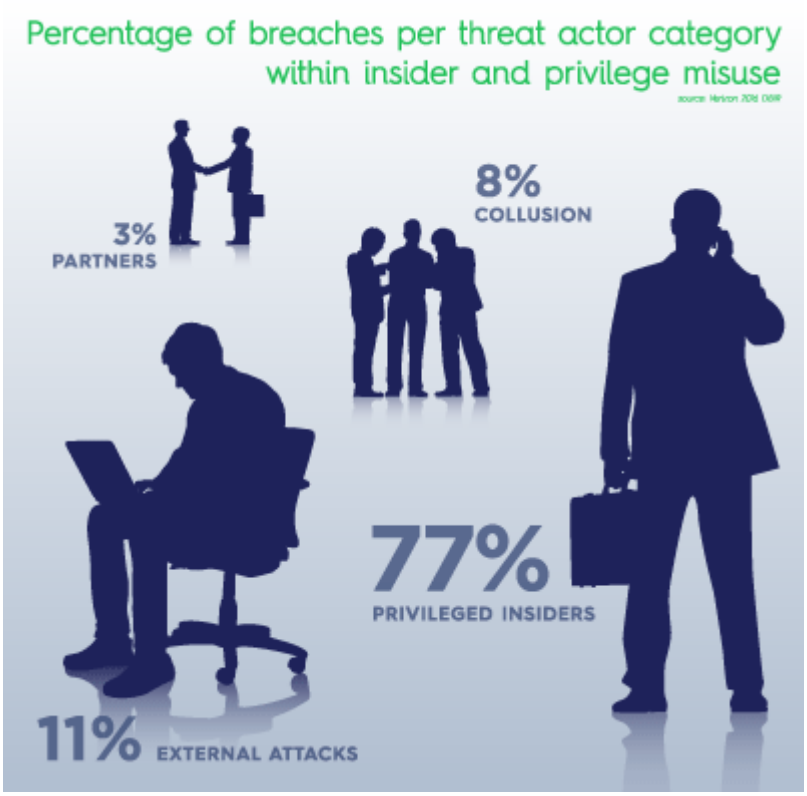

External Hackers

Hackers covet privileged accounts and credentials, knowing that, once obtained, they provide a fast track to an organization's most critical systems and sensitive data. With privileged credentials in hand, a hacker essentially becomes an "insider"—and that's a dangerous scenario.

One tactic hackers commonly use is to gain an initial foot-hold through a low-level exploit, such as through a phishing attack on a standard user account, and then surreptitiously zig-zag laterally through the network until they find a dormant or orphaned account that allows them to escalate their privileges. Organizations that lack robust enterprise password management capabilities, such as automated password rotation, are also more susceptible to pass-the-hash (PtH) attacks. In these attacks a hacker who has already gained low-level credentials, can steal the password hash from an admin account if revealed—such as during a helpdesk session with the infiltrated account—and then reuse the hash to unlock administrative access rights.

Hackers are also adept at obtaining deep privileges on a single computer and then expanding their privileges to other devices across a network. With the proper privileges, and inadequate technology controls, a hacker can easily erase their tracks to avoid detection while they traverse the environment and get closer to achieving their objectives.

Insider Privilege Misuse or Abuse

Allowing a user, or perhaps even multiple users, to utilize the all-powerful root in Unix / Linux environments means a simple error, such as a mistyped command or an accidental deletion, could have far-reaching consequences, causing down-time of Tier-1 systems, opening up gigantic vulnerabilities that let in rootkits and other exploits, or worse. In Windows environments, misused admin accounts also have potential to cause outsized damage compared to non-privileged accounts.

However, rogue insiders likely pose the most dangerous threat. Insider threats take the longest to uncover—as employees, and other insiders, generally benefit from some level of trust by default, which may help them avoid detection. The protracted time-to-discovery also translates into more potential for damage. Certainly, many of the most devastating breaches over the past 10 years have been perpetrated by insiders. Unlike external hackers, insiders already start within the perimeter, while also benefitting from know-how of where sensitive assets and data lie and how to zero in on them.

Cautionary Tales of Untamed Privileges: NSA & Target Breaches

As a technology contractor for the NSA, Edward Snowden had administrative access rights, ostensibly to perform such activities as backing up computer systems and migrating data to local servers. However, by infamously abusing his admin privileges, and utilizing some simple and widely available software tools, including an automated web crawler, Snowden illegally copied, accessed, and then leaked an estimated 1.7 million NSA files. In response to the Snowden breach, the NSA announced the drastic action of eliminating 90% of system administrators, to limit access and improve its least-privilege posture.

As for the 2013 Target breach, which impacted roughly 70 million customers, hackers gained unauthorized access into Target's systems through pilfered credentials of a third-party vendor, a heating and air conditioning contractor. The HVAC contractor had access to Target's network, including permissions to upload executables, that were far more than required for it to perform its maintenance work. By restricting access permissions to the fewest resources, functions, and areas necessary for the HVAC company, Target likely would have avoided the breach, and the subsequent fallout.

Emerging Privileged Threat Vector – Internet of Things

A rapidly expanding universe of connected "things"—from health monitoring and delivery devices to industrial controls to wearables, smart appliances, and more–presents enormous challenges for IT with regard to identifying and securely onboarding legitimate devices at scale. IoT devices frequently suffer serious security limitations, such as the inability to have the software hardened or firmware updated, and hard-coded passwords. IoT is also vulnerable to exploits such as man in the middle attacks (MiTM), when an attacker intercepts and / or possibly modifies data communicated between two systems.

In the latter months of 2016, large-scale IoT hacks finally made the leap from theory to reality. DDOS attacks leveraging IoT botnets comprised of as many as a million "things"

(such as cameras, thermostats, DVRs, and even lightbulbs) knocked many U.S. East Coast businesses, and the nation of Liberia, offline, in separate incidents.

While the term "IoT" and all the "smart" things are relatively recent phenomena, and certainly pose some unique security challenges, when it comes down to it, organizations need to be able to enforce least privilege and / or application control over any endpoint with an IP—traditional or IoT. Network segmentation for IoT devices is one way to broadly restrict the permissions of IoT devices and the associated systems and operations, while role-based access permissions should also be enforced as a best practice.

Challenges to Applying Least Privilege

If excessive privileges pose such a vexing threat, why don't organizations rein them in? The short answer is that companies do—or, try to via policies and technology solutions, but a number of factors contribute to make dialing in the optimal amount of privilege rights and access a challenge.

Lack of visibility and awareness

Lack of visibility and awareness of all of the privileged accounts, assets, and credentials across an enterprise stands as one pervasive stumbling block for companies in effectively managing privileges. Independent research, surveys, and audits shed light on the reality that permanent accounts, and orphaned accounts with high degrees of privilege, are sprawled across the physical, virtual, and cloud environments of most organizations. To complicate matters, manifold systems and applications have easily guessable default credentials. Until these applications (such as shadow IT) and systems are discovered, properly configured, and brought under management, they pose a high risk for exploit by a hacker or through malware.

Whimsically termed FoMP (Fear of Missing Privileges), – is "a pervasive apprehension that user accounts, privileged accounts, or assets might not be securely found or managed, creating a sense of foreboding and anxiety about insider threats and potential data breaches."

Cultural Challenges

Employee resistance is another challenge that often rears its head in the face of least-privilege policies. If privileged access controls are overly restrictive, they can disrupt user workflows, causing frustration and hindering productivity. To obviate helpdesk requests and end-user headaches (users rarely complain about having too many privileges), IT admins

traditionally provide end users with broad sets of privileges. This superfluous access translates into a broadened attack surface.

Role-based access, administered through Active Directory or another rights management solution, can help enforce general rules around a role, a group, a team, or an individual's set of privileges. However, an individual's role is often fluid and can evolve over the course of employment such that they accumulate new responsibilities and corresponding privileges—while still retaining privileges that they no longer use or that are irrelevant to their role. Moreover, a drawback to role-based access is that it lacks the contextual granularity to only provide access when required for a specific use.

Benefits of Least Privilege for Security & Productivity

According to the Verizon Data Breach Digest, in 2015, 85% of vulnerabilities on Windows systems could have been mitigated by removing admin rights. In fact, almost every vulnerability (99.5%) that could have resulted from users surfing the web using Internet Explorer could be mitigated by not running as an administrator!

Unfettered privileged rights and access essentially equates to uncapped potential for damage. The more privileges a user, account, or process amasses, the greater the potential for abuse, exploit, or error. Implementing least privilege not only reduces the likelihood of a breach occurring in the first place, but it helps limit the scope of a breach should one happen.

Least privilege confers at least several, high-level benefits, including:

1) A condensed attack surface: Limiting privileges for people, processes, and applications means the pathways and ingresses for exploit are also diminished.

2) Reduced malware infection and propagation: Least privilege helps dramatically reduce malware infection and propagation, as the malware (such as SQL injections, which rely on lack of least privilege) should be denied the ability to elevate processes that allow it to install or execute.

3) Improved operational performance: When it comes to applications and systems, restricting privileges to the minimal range of processes to perform an authorized activity reduces the chance of incompatibility issues cropping up between other applications or systems, and helps reduce the risk of downtime.

4) Easier to achieve, and prove, compliance: By constraining the activities that can possibly be performed, least privilege enforcement helps create a less complex, and thus, a more audit-friendly, environment. Moreover, many compliance regulations (including HIPAA, PCI DSS, FDDC, Government Connect, FISMA, and SOX) require that organizations apply leastprivilege access policies to ensure proper data stewardship and systems security. For instance:

- The US federal government's FDCC mandate states that federal employees must log in to PCs with standard user privileges.
- The HIPAA Privacy Rule provides guidelines for the establishment of least privilege, such as restricting access to data (i.e. a subset of a patient record as opposed to the entire record) base on the "minimum necessary use" to accomplish a specific purpose.
- PCI DSS states that organizations that process or store credit card data must restrict access to cardholder data by business need to know, and specifically invokes the use of least privilege user accounts [7.1.1 Restriction of access rights to privileged user IDs to least privileges necessary to perform job responsibilities; 7.2.2 Assignment of privileges to individuals based on job classification and function].

Least Privilege Strategies & Best Practices

So, what does a wholly-adopted least privilege computing environment look like? How does your organization get there?

Here are some best practices and strategies to help you bake in least privilege across your organization:

1) Perform a privilege audit to discover, and bring under policy management, all privileged accounts and credentials. This should include all user and local accounts, SSH keys, Windows and Linux groups, and default and hard-coded passwords. Here's a free tool.

2) Remove admin rights on endpoints – As opposed to provisioning default access, default all users to standard privileges, while enabling elevated privileges for applications and to perform specific tasks. If access is not initially provided and is required, the user can submit a help desk request for approval. (Diminish user frustration and help desk calls by enforcing least privilege through granular policies, such as with a PAM solution).

3) Remove all root and admin access rights to servers and reduce every user to a standard user. This dramatically reduces the attack surface and helps safeguard your most sensitive assets.

4) Segment systems and networks to broadly separate users and processes based on different levels of trust, needs, and privilege sets.

5) Enforce separation of privileges, this includes separating administrative account functions from standard account requirements, separating auditing/logging capabilities within the administrative accounts, and separating system functions (read, edit, write, execute, etc.).

6) Implement just-in-time privileges (JIT): Elevate privileges on an as-needed basis for specific applications and tasks only for the moment of time they are needed, without requiring administrative credentials or exposing passwords. This is sometimes called "privilege bracketing".

7) Expiration of privileged access: Access should be time-bound and/or bound by completion of a specific task or process.

8) Implement one-time-use credentials: For instance, use password "safes," where a single-use password for privileged accounts is "checked out" until an activity is completed, immediately after which time, it is checked back in. The used password is then retired.

9) Replace hard-coded credentials with APIs allowing credentials to be retrieved from password safes.

10) Limit membership for the superuser role to the minimal number of people necessary

11) Individual actions should be traceable: This can be accomplished through user IDs as well as auditing and other tools, to provide oversight and accountability.

12) Enforce vulnerability-based least-privilege access: Incorporate real-time vulnerability and threat data about an asset or user to make dynamic risk-based access decisions. For instance, restrict privileges of an account suspected of compromise or for an application with a known vulnerability.

13) Analyze and report on access to all privileged accounts: This should encompass shared admin accounts, application accounts (A2A), local admin accounts, service accounts, database accounts, cloud and social media accounts, devices (including IoT), and SSH keys. Auditing activities can also include capturing keystrokes and screens.

How to Implement Least Privilege

So, now that you have some best practices and strategies, what are the best tools or technologies to implement least privilege?

Network segmentation, such as the creation of different zones through firewall configuration and rules, is one key way to enforce least privilege in broad strokes. By controlling access and movement between zones, which may have a different mix of applications and services, firewalls can restrict users broadly based on privileges. For instance, firewalls are often used to create a DMZ (demilitarized zone) between a corporate network and the public network. Firewalls can also easily block unauthorized privilege elevation activity (such as from service requests) based on rules applicable to the zone.

Privilege Access Management (PAM), also called Privileged Identity Management (PIM) or just Privilege Management, involves the creation and deployment of solutions and strategies to manage privilege accounts across an environment. Holistic PAM solutions discover and bring under management all privileged accounts. These solutions remove admin rights from users, and instead, elevate privileges for authorized applications or tasks as-needed. While identity and access management (IAM) controls provide authentication of identities, PAM allows organizations to control authorization over ability to perform a granular activity. IAM

and PAM solutions working together can help provide fined-grain control, visibility, and auditability over all credentials and privileges.

Systems hardening, entailing the removal of superfluous programs, accounts, and services (such as with a server that connects to the internet), and the closing of un-needed firewall ports, is another common mechanism for applying least privilege. This practice not only markedly improves security posture by reducing the attack surface, but it also reduces complexity and simplifies the environment.

Effective implementation of least privilege will involve policies, procedures, technologies—and proper configuration of those technologies. Formalizing a policy should also help you get a better handle on where your sensitive data resides, and who can access it.

Most likely, various aspects of least privilege will need to be phased in piecemeal over time, rather than implemented all at once. While many organizations tackle privilege management challenges in a similar order, which you can learn from in this privilege management maturity model, the best path forward for any organization will always be tailored to its unique needs and resources The more mature a least-privilege policy implementation, the more effective an organization will be in condensing the attack surface, mitigating the impact of attacks (by hackers, malware, or insiders), enhancing operational performance, and in reducing the risk from and impact of user errors.

Thanks for tuning in! Get the full PDF for easy sharing and future reference.. In future blogs, we will tackle related topics, such as privilege access management, privileged password credential management, and more. In the meantime, I've share some related resources below.

Least Privilege Management Resources

- Privilege Discovery & Reporting Tool (DART) (free)
- Privilege Management Overview (video)
- Endpoint Least Privilege (video)
- Least Privilege for Servers (video)
- Privileged Password Management (video)
- Privilege Benchmarking Study (survey revealing best—and worst—privilege management practices)
- Removing Users From The Local Administrators Group (blog)
- 2016 Gartner Market Guide for Privileged Access Management (analyst research)
- The Forrester Wave™: Privileged Identity Management, Q3 2016 (analyst research)
- Delegating Privileges to Domain Controllers and Active Directory without the Security Risk (blog)
- Removing Users From The Local Administrators Group (blog)
- Best Practices for Managing Domain Admin Accounts (blog)
- Seven Steps to Complete Privileged Account Management (white paper)

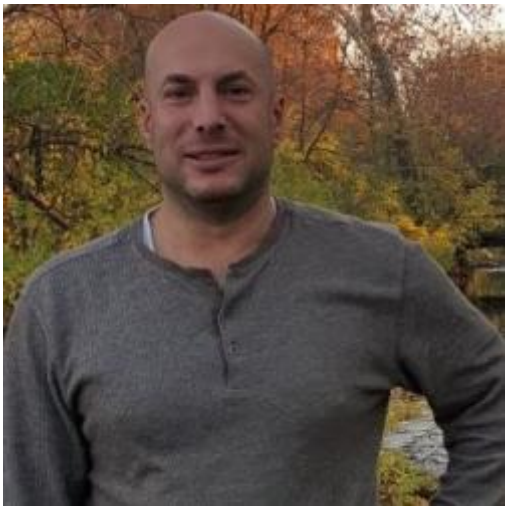

Matt Miller, Content Marketing Manager

Matt Miller is a Content Marketing Manager at BeyondTrust. Prior to BeyondTrust, he developed and executed marketing strategies on cyber security and cloud technologies in roles at Accelerite (a business unit of Persistent Systems), WatchGuard Technologies, and Microsoft. Earlier in his career Matt held various roles in IR, marketing, and corporate communications in the biotech / biopharmaceutical industry. His experience and interests traverse cyber security, cloud / virtualization, IoT, economics, information governance, and risk management. He is also an avid homebrewer (working toward his Black Belt in beer) and writer.

Other resources you might like