



Advancing Zero Trust with Privileged Access Management (PAM)

Thank you for downloading this BeyondTrust Resource. Carahsoft is the Distributor for BeyondTrust solutions available via GSA 2GIT, NASA SEWP V, ITES-SW2, and other contract vehicles.

To learn how to take the next step toward acquiring BeyondTrust's solutions, please check out the following resources and information:



For additional resources:
carah.io/beyondtrustresources



For upcoming events:
carah.io/beyondtrustevents



For additional BeyondTrust solutions:
carah.io/beyondtrustsolutions



For additional Cyber solutions:
carah.io/cybersolutions



To set up a meeting:
beyondtrust@carahsoft.com
(866)-421-4683



To purchase, check out the contract vehicles available for procurement:
carah.io/beyondtrustcontracts



>>> Taking Principles
of Zero Trust Into
Real-World
Privileged Access
Management (PAM)

Advancing Zero Trust

with **Privileged Access
Management (PAM)**





TABLE OF CONTENTS

Introduction

Defining Zero Trust

The Role of PAM

Design Considerations

Partnering with BeyondTrust

Introduction: The Zero Trust Journey

Zero trust principles and architectures are being widely embraced and adopted by the public and private sectors. Legacy security architectures and network defenses are simply ineffective at managing a world more reliant on the cloud and remote workers. In this world, zero trust principles have become one of the most effective approaches to mitigating the heightened risks to highly sensitive identities, assets, and resources.

NIST SP 800-207¹ has been widely embraced as a framework for zero trust. NIST refers to this SP as “a definitive source of ZTA concepts and principles.” This initial publication (circa August 2020) primarily covered zero trust concepts at a high level. It did not provide a foundation for architecting and implementing zero trust at a practical level.

To implement zero trust in practical terms, an organization must grasp which technologies and configurations can actually be implemented with tenets that meet theoretical requirements. To that end, NIST has published SP NIST 1800-35² on implementing Zero Trust (circa December 2022) to bridge that gap. Today, solutions exist that can address both the theoretical and practical requirements of zero trust. For BeyondTrust technology, that’s where this paper comes in.



This paper is for those IT and security professionals who seek to bridge the principles of zero trust, as scoped by NIST, into real-world privileged access management (PAM) and secure remote access product capabilities that can enable zero trust across the public or private enterprise.

No identities are more imperative to secure than those with privileged access to systems, data, applications, and other sensitive resources. Almost every attack today requires privilege for the initial exploit or to laterally move within a network.

PAM protects privileged accounts and credentials, granularly enforces least privilege, and monitors and manages every session involving privileged access -- whether human, machine, employee or vendor. **PAM is also essential to enabling a zero trust architecture (ZTA) and to addressing the 7 tenets of zero trust, as put forth by NIST in their publications.**

This paper will explore:

- Key zero trust definitions and concepts, as put forth by NIST
- Security implications of zero trust
- Practical implementation steps of zero trust with Privileged Access Management and Secure Remote Access solutions
- How BeyondTrust enables organizations to achieve zero trust
- Design considerations for zero trust architectures

¹<https://csrc.nist.gov/publications/detail/sp/800-207/final>

²<https://csrc.nist.gov/publications/detail/sp/1800-35/draft>



NIST 800-207 and 1800-35B: Defining Zero Trust & Zero Trust Architecture (ZTA)

Definition of Zero Trust

The National Institute of Standards and Technology (NIST) defines Zero Trust (ZT) as “an evolving set of cybersecurity paradigms that move defenses from static, network-based perimeters to focus on users, assets, and resources.” NIST further explains that the collection of concepts that comprise the zero trust principles are designed to “minimize uncertainty in enforcing accurate, least privilege per-request access decisions in information systems and services in the face of a network viewed as contested.”

In practical terms, this entails eliminating persistent trust, performing continuous authentication, granularly restricting access to the minimum needed, applying segmentation and microsegmentation strategies, and continuously auditing access.

Privileged Access Management (PAM) is a foundational technology stack for implementing each of these essential zero trust security controls.

➤ *“The interactions between the products in the [BeyondTrust] suite have been brilliantly and carefully orchestrated in a way that allows us to maximize our chance of getting as far down the Zero Trust road as we possibly can given the state of the products in the security market.”*

Brandon Haberfeld, Global Head of Platform Security, Investec



The primary goal of zero trust is to protect enterprise resources (particularly, but not solely limited to, data). As outlined in [NIST SP 1800-35B](#), this goal has become increasingly important in response to today's network challenges:

- Networks have become decentralized, perimeterless environments with resources distributed across both on-premises environments and multiple clouds.
- Many users need access from anywhere, at any time, from any device to support the organization's mission.
- Data is programmatically stored, transmitted, and processed across different boundaries under the control of different organizations to meet ever-evolving business use cases.
- It is no longer feasible to simply enforce access controls at the perimeter of the enterprise environment and assume that all subjects within it can be trusted.

Thus, network location can no longer be treated as the prime component to the security posture of the resource. Instead, the zero trust model operates under the assumption that no asset or user account can be implicitly trusted based solely on their physical or network location or asset ownership. Subject and device authentication and authorization must be required before establishing a session to an enterprise resource.

Definition of Zero Trust Architecture

[NIST Special Publication \(SP\) 800-207](#) defines a Zero Trust Architecture (ZTA) as "an enterprise's cybersecurity plan that utilizes zero trust concepts and encompasses component relationships, workflow planning, and access policies. Therefore, a zero trust enterprise is the network infrastructure (physical and virtual) and operational policies that are in place for an enterprise as a product of a zero trust architecture plan."



NIST further articulates that the primary focus of a ZTA is “protecting data and resources. It enables secure, authorized access to enterprise resources that are distributed across on-premises and multiple cloud environments, while enabling a hybrid workforce and partners to access resources from anywhere, at any time, from any device in support of the organization’s mission.”

The concepts and principles in a NIST SP 800-207 zero trust architecture are applied to “enterprise networks that are composed of pre-established devices and components and that store critical corporate resources both on-premises and in the cloud.” According to SP 1800-35B, ZTA performs real-time, continuous behavioral analysis and risk-based assessment of the access transaction or session. “For each access request, ZTA verifies the requester’s identity and role, the requesting device’s health and credentials, and possibly other information. If defined policy is met, ZTA dynamically creates a secure connection to protect all information transferred to and from the accessed resource.”

Each access request is evaluated by verifying:

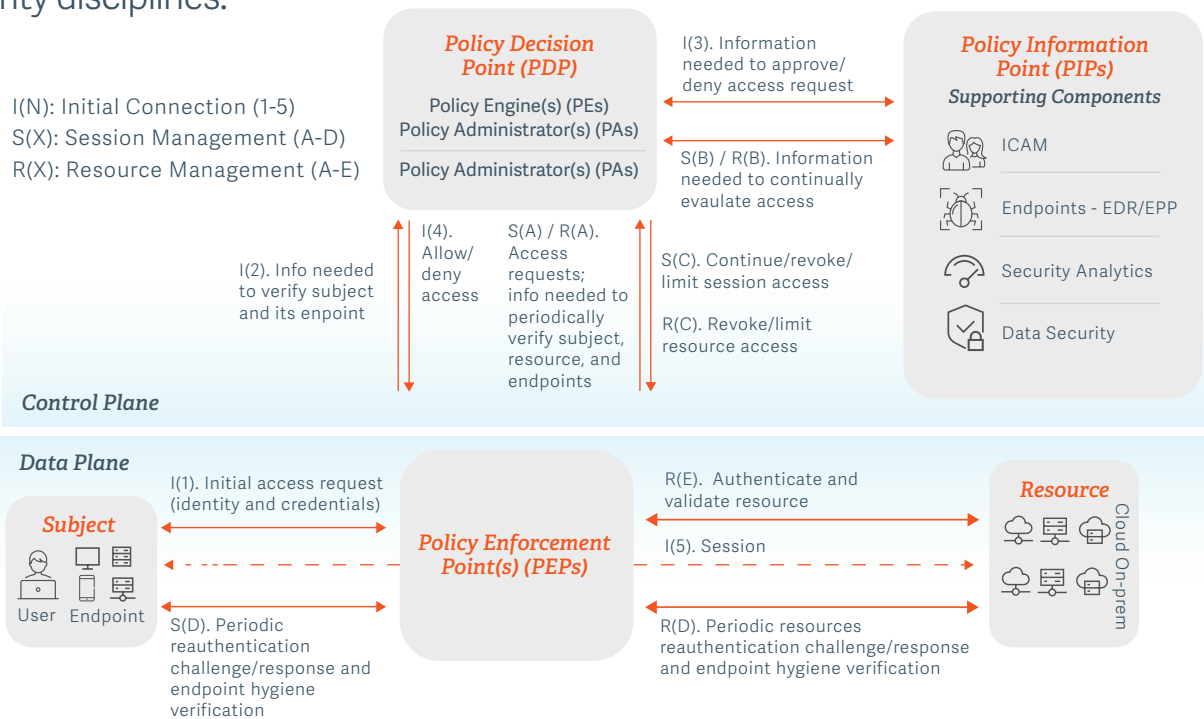
- The context available at access time, including the requester’s identity and role
- The requesting device’s health and credentials
- The sensitivity of the resource





Zero Trust Reference Architecture

Here is an example of a ZTA, as outlined in NIST SP 1800-35B, applied to multiple security disciplines:



The core ZTA components depicted in the architecture diagram above include:

- Policy Engine (PE):** The PE handles the ultimate decision to grant, deny, or revoke access to a resource for a given subject. The PE calculates the trust scores/confidence levels and ultimate access decisions based on enterprise policy and information from supporting components. The PE executes its trust algorithm to evaluate each resource request it receives.
- Policy Administrator (PA):** The PA executes the PE’s policy decision by sending commands to the PEP to establish and terminate the communications path between the subject and the resource. It generates any session-specific authentication and authorization token or credential used by the subject to access the enterprise resource.
- Policy Enforcement Point (PEP):** The PEP guards the trust zone that hosts one or more enterprise resources. It handles enabling, monitoring, and, eventually, terminating connections between subjects and enterprise resources. It operates based on commands received from the PA.



Other important definitions include:

- **Policy Decision Point (PDP):** The PE and PA combine to comprise the PDP, which executes the decision on whether a subject is permitted access to a resource.
- **Policy Information Point (PIP):** The PIP provides telemetry and other data to enable the PDP to make informed access decisions. This includes PAM, EDR, and identity-based threat and response solutions.
- **Subject:** An end user, application, and other non-human entity that requests information from resources.
- **The Gateway:** Responsible for enabling, monitoring, and terminating the connection between the subject (user or application) and resource via the agent so all activity can be assessed and documented.

Zero Trust Enclaves

As a security best practice, network and access controls should be configured, tested, and monitored for traditional on-premises resources, to prevent inappropriate access. Unfortunately, in many modern environments, these security controls have been loosened or removed to support myriad use cases—from hybrid cloud environments to remote workers, to solutions that span multiple geolocations and networks. This does not mean that the security controls have been removed, but rather that sufficient exceptions have been implemented to enable additional business use cases, and the added risks have been accepted. This “unhardening” of the environment runs contrary to the zero trust model and, in many cases, uses static credentials or secrets for solutions to interoperate. Thus, a compromise in one location jeopardizes the entire environment.



To achieve zero trust while supporting these newer use cases, a new perimeter must be established. This entails taking a zone-based approach to creating a secure enclave. Access into the enclave itself, as well as between resources that reside within it, should be tightly restricted.

Therefore, an enclave is like a network perimeter within an unsecure network, and the containment, with all the security controls, is built within.

Gateway Enclave

A gateway enclave model establishes a granular segmented perimeter with assets accessible only through a gated and monitored network path. While resources within the enclave can have loosened security controls to meet the operational business requirements, they are still monitored for inappropriate behavior when activity originates from outside the enclave. Think of the enclave as a mini trusted network within another network.

Gaining access to the enclave requires security controls beyond what is enabled via a typical bastion host or restricted network tunneling and routing controls. This is key for establishing segmentation.





The following indicate other necessary requirements to enable zero trust functionality within an enclave:

- An external resource holds the policy and administration for all access.
- Access is brokered by a logical engine that processes the policy and attributes to determine or limit access.
- Access is fully documented—from authorization through complete session management and session recording.
- Sessions are monitored for behavioral activity, from inception through any required lateral movement, to complete a task. Inappropriate activity can be responded to in real-time.
- Access is not absolute and should follow workflows, like just-in-time, including ephemeral connectivity, and integration into ITSM and other change management solutions.
- No sessions or traffic can access the enclave unless processed through the control plane. In other words, all lateral movement into the enclave should be denied using traditional network and access security controls. The enclave should operate as a secure segment of the network that is fully gated and lacks backdoors.

Success with this model demands a balance between all these requirements. Authorized users and applications must initiate a session that can be verified and connected into the gateway enclave. Session monitoring, credential injection, and least privilege must be applied to overcome the security and compliance concerns governing an organization and be consistent with zero trust objectives.



Resource Enclaves

A resource enclave is a variation of a network zone or VLAN. It is a collection of resources (assets) (applications, operating systems, network devices, databases, etc.) with a hardened perimeter around all the assets. In lieu of a perimeter being a broad network zone, it is isolated to only critical assets within the resource enclave for a given purpose. Essentially, a resource enclave is a secure network zone with limited external access and is fully segmented. Access of any type must come through a gateway, as described by a gateway enclave. Essentially, a gateway enclave defines the path for secure access and the resource enclave defines the resources contained within.

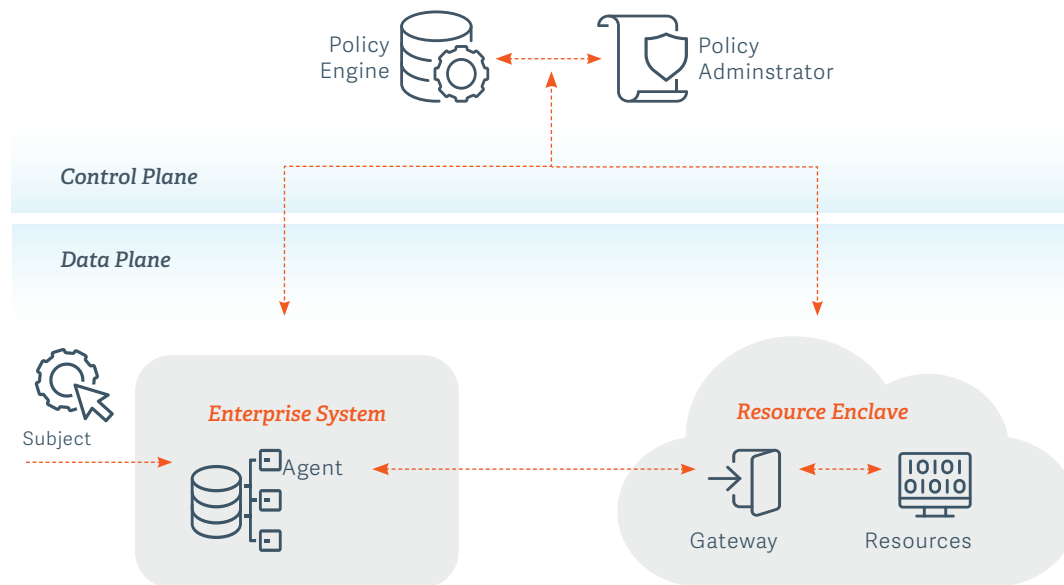
In any organization, there will potentially be dozens, or hundreds, of resource enclaves secured by traditional access control lists and network access hardening. This represents an expansion of the microsegmentation concept to unique implementations to meet business or technology requirements. A resource enclave is adaptable to legacy systems, and the architecture must be compatible with zero trust, either natively or using a gateway enclave approach.

Therefore, resources within the enclave rely on traditional security control models and do not necessarily mitigate threats (i.e. inappropriate lateral movement) within the enclave itself unless connectivity is fully monitored. For example, threats from lateral movement are only mitigated between subjects and from resource enclave to resource enclave.

As a matter of principle, all sessions and connectivity must be enforceable, whether the user or application is operating from the corporate office or from a remote location. This paradigm can help a zero trust architecture (ZTA) play a pivotal role to overcoming legacy network challenges, when a broad network perimeter is no longer used for security enforcement.



Consider this diagram of a simplified NIST-based zero trust architecture for a secure enclave gateway:



In a pure zero trust environment, the gateway and resource enclave can exist on a single asset and, therefore, the application is properly restricted. While this is the goal, it is not realistic for many organizations adopting zero trust for the first time. The concept of an enclave can scale up or down depending on your resources and assets that are being implemented during a zero trust initiative.

Privileged Access Management (PAM) is Foundational to Zero Trust

Privileged Access Management (PAM) consists of the cybersecurity strategies and technologies for exerting control over the elevated (“privileged”) access and permissions for users, accounts, processes, and systems across an IT environment. Use cases range from enforcing least privilege, to managing privileged accounts/credentials, to securing all privileged remote access, and more. PAM is a fundamental security control that falls within the broader Identity and Access Management (IAM) umbrella.



> *“The majority of the systems within the buildings being accessed are not traditional IT systems. They are building control systems, like smart elevators, surveillance systems and HVAC units, where it is not possible to install antivirus software. We recognize that privileged access management is one of the most of important tenets of a modern cyber security program and a must have for a zero-trust architecture and robust BYOD security framework.”*

Curtis Jack, Manager of Technical Engineering, Oxford Properties Group

 [View the case study](#)

Today, enterprises and public agencies are challenged with securing significantly more privileged accounts and supporting vastly larger remote, hybrid, and decentralized networks than in years past—many of them are not even considered traditional privileged accounts operated by administrators or system owners.

We have more tools, users, and machines requiring privileged access than ever before. Resources and privileged accounts can be provisioned at great scale across cloud environments. However, even ephemeral resources can create a risk surface and must be appropriately accounted for with security controls.

Increasingly, resources that require authentication, privileges, and access may reside outside of corporate governance. This can include other untrusted resources or identities, accounts, and processes. These realities have given rise to the concept of the Data Plane, which is important to manage.

A zero trust architecture (ZTA) addresses modern privileged access security challenges. A ZTA enforces granular, secure, authorized access near the resources—whether located on-premises or in the cloud—for a remote workforce and partners based on an organization’s defined access policy.



When applying the granularity of privileged access management (PAM), zero trust can ensure all access is managed and documented for appropriate behavior, regardless of the size and design of an enclave. Today, this is a particularly important challenge to solve in order for organizations to meet regulatory compliance requirements, as well as cyber insurance eligibility requirements.

How BeyondTrust Privileged Access Management (PAM) Enables a ZTA

As envisioned by NIST (SP 800-207), a zero trust security model eliminates persistent trust and enforces continuous authentication, least privilege, and adaptive access control. This strategy also applies segmentation and microsegmentation for secure access. A zero trust approach is about constant visibility and control over who is doing what on your network. These are all core capabilities modern privileged access management (PAM) solutions, like BeyondTrust, are designed to elegantly address.

➤ *"The interactions between the products in the [BeyondTrust] suite have been brilliantly and carefully orchestrated in a way that we are maximizing our chance of getting as far down the Zero Trust road as we possibly can given the state of the products in the security market."*

Brandon Haberfeld, Global Head of Platform Security, Investec.

▶ [View the case study](#)



BeyondTrust Privileged Access Management solutions support the smart, practical implementation of NIST's zero trust security model—without disrupting business processes. Our solutions eliminate persistent trust by ensuring that all privileged access and permissions are being continuously audited, and least privilege access is provisioned just-in-time and revoked immediately upon completion of a task, change in context, or if a certain amount of time has expired.



With BeyondTrust, you can start with the PAM use cases most urgent to your organization, and then seamlessly address remaining use cases over time. Each use case, once addressed, will provide enhanced control and accountability over the accounts, assets, users, systems, and activities that comprise your privilege environment, while eliminating and mitigating multiple threat vectors. Moreover, BeyondTrust provides centralized privileged access security controls across your heterogeneous environment—on-premises, cloud/multicloud (AWS, Azure, etc.), Windows, macOS, Unix, Linux, etc. No other vendor provides such deep and expansive control over privileged access.

The more use cases you address, the more PAM synergies emerge, and the more impact you'll realize in reducing enterprise risk, improving operations, and achieving zero trust objectives.

At a high level, BeyondTrust PAM provides the following capabilities across on-premise and cloud environments:

- Discovers, onboards, and catalogs all privileged identities, accounts, and assets.
- Enforces adaptive access and continuous authentication to ensure all devices, users, accounts, and identities have a high confidence in their actual identity. In other words, they are who they say they are above just positive authentication.
- Right-sizes privileged access and entitlements by applying least privilege, including just-in-time access, to all sessions, endpoints, and applications.
- Enables secure, least privilege remote access for vendors, employees, and service desks for sessions and trusted applications.
- Implements segmentation and microsegmentation to isolate assets and users, and to prevent lateral movement.



- Monitors and manages every privileged session, providing continuous visibility and control over who is doing what, and why, so any suspicious behavior can trigger immediate revocation of permissions and access.
- Extends Microsoft® Active Directory authentication, single sign-on capabilities, and Group Policy configuration management to Unix and Linux systems, simplifying the secure management of identities and the path to implementing zero trust enterprise-wide, regardless of operating system or application.

Table: How BeyondTrust PAM Fits within the NIST Zero Trust Architecture

		NIST ZTA Core Concepts			
		Policy Engine (PE)	Policy Administrator (PA)	Policy Enforcement Point (PEP)	Gateway
BeyondTrust Solutions	Password Safe	Found in the management capabilities governing secrets management, and the role and attribute-based access models defined by the Policy Administrator. BeyondTrust Password Safe's policy engine decides on access availability based on a wide variety of criteria and tracks privileged access—even if it is ephemeral.	Creates, updates, and manages the policy for end users and machine identities to grant access and automate a privileged session. Access to the resource enclave is granted to the Policy Administrator and can be managed through the BeyondTrust BeyondInsight Console.		<p>The Gateway is implemented using:</p> <ol style="list-style-type: none"> 1. A BeyondTrust appliance as a worker node when installing Password Safe completely on-premises, regardless of whether Internet access is present. 2. A Resource Broker when using Password Safe to manage assets on-premises (preferred). <p>The Gateway responds to a privileged access request for a session or application from a trusted Policy Engine. This is fundamental to zero trust since a user is never granted direct access to a resource, as when using a native operating system protocol. All connectivity is managed and monitored by the agent, and only session activity (screen, terminal, or web page) is rendered to the requesting subject.</p>



NIST ZTA Core Concepts				
	Policy Engine (PE)	Policy Administrator (PA)	Policy Enforcement Point (PEP)	Gateway
Privileged Remote Access	<p>Found in the management capabilities governing remote access, and the role and attribute-based access models defined by the Policy Administrator. BeyondTrust Secure Remote Access can manage assets and users in any network zone, regardless of perimeter, as long as Internet access is available.</p>	<p>Creates, updates, and manages the policy for end users; grants access; and automates application remote access. This is the basis for zero trust. Access to the resource or application is granted to the Policy Administrator and can be managed through the BeyondTrust Secure Remote Access Console.</p>	<p>Implemented using a Secure Remote Access Jump Client. It responds to a remote access request for a session or application from a trusted Policy Engine or intermediate proxy, called a Jump Point. A user is never granted direct access to a resource, as when using a native operating system protocol. All connectivity is managed and monitored, and only session activity (screen, terminal, or web page) is rendered to the requesting user.</p>	
Privilege Management for Windows/Mac & Unix/Linux	<p>The Policy Engine can be found in management capabilities of the rules, policies, and log engine governing least-privilege endpoint access, and in the role and attribute-based access models defined by the Policy Administrator.</p>	<p>The Web Policy Administrator creates, updates, and manages the policy for end users, grants access, and automates application access. This is the basis for zero trust. Access to the resource or application is granted to the Policy Administrator and can be managed through the BeyondTrust Privilege Management Windows & Mac Interface (when deployed via the Cloud).</p>	<p>The Policy Enforcement Point is the least-privilege client installed on Windows, macOS, Unix, and Linux endpoints. For Windows and macOS, it initiates privileged applications and performs application control on the endpoints on behalf of the user or application. The Policy Enforcement Point capability compares the application execution request to the defined policy and launches (or denies) the application with the appropriate privileges, without actually using privileged credentials. This capability is fundamental to zero trust since the application or user is never granted administrative credentials, but can execute an application with privileges.</p> <p>For Unix/Linux, it initiates commands on the host on behalf of the user or application—without the end user actually logging in—and renders the results transparently to the end user.</p>	



Privileged Password Management

Digital Transformation Ushers in Privileged Account and Access Sprawl

Organizations have experienced an explosion of privileged accounts (human, application, machine, etc.) that must be managed. These privileged accounts are supporting on-premises technology, the cloud, hybrid environments, and many of the SaaS, IaaS, and PaaS solutions that help power a modern business.

Privileged credentials and other secrets must be managed using a model that can support your modern environment. Access to passwords and secrets must address human and machine identities, whether on-premises or in the cloud, or being licensed from a vendor, supplier, or solutions provider.

Enable Secure, Adaptive Privileged Access & Oversight with Privileged Password Management

Privileged password management (also known as privileged account and session management) provides an optimal way to architect access to sensitive resources, making it well-matched for current challenges. The approach to accomplish this involves using a Zero Trust Enclave Gateway, as described earlier.



BeyondTrust's privileged password management solution, Password Safe, can ensure your resources are managed and protected from potentially inappropriate connection abuse, and that all resources contained within an enclave are executed within a zero trust model. This means no end users or machine identities are ever trusted for a direct privileged session—unless their access can be brokered through a gateway. All session activity is fully monitored. This holds true for any location in which a resource enclave may reside, irrespective of the perimeter.

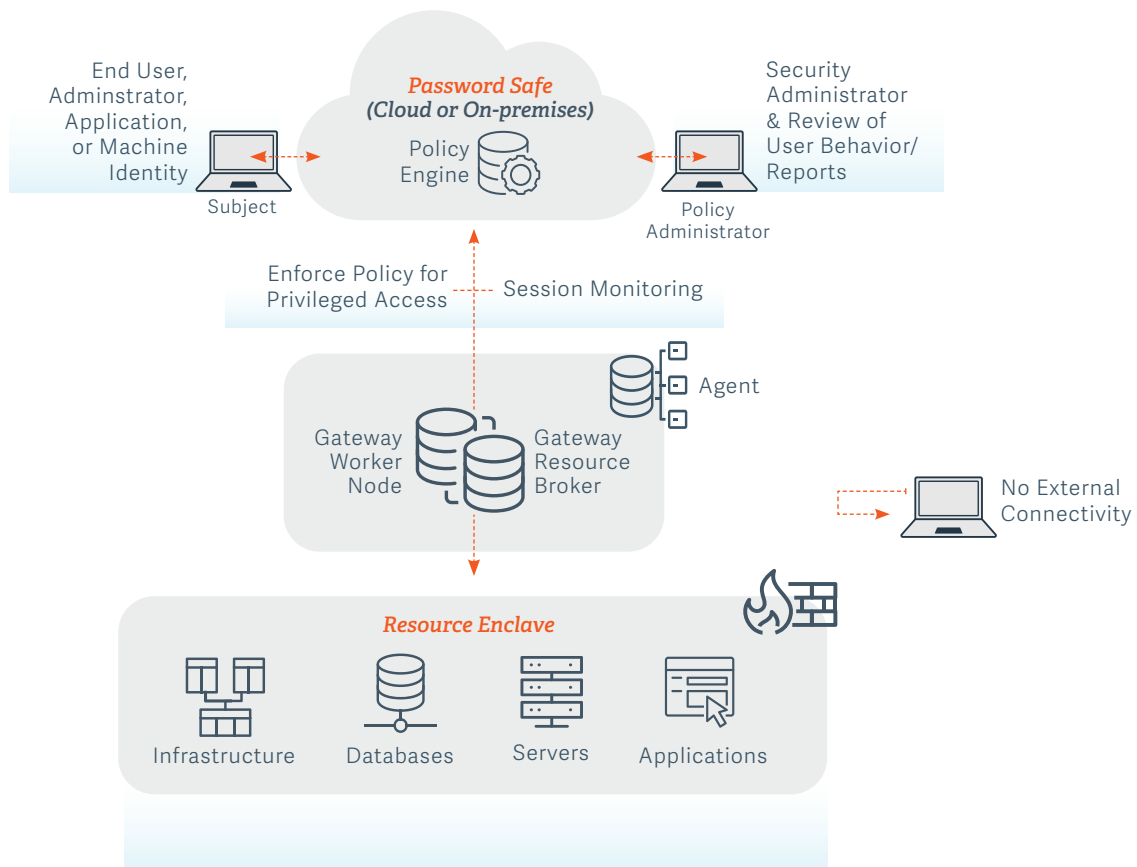
Essential Password Safe zero trust use cases include:

- Securely storing and allowing for retrieval of passwords, credentials, and other secrets to provide access to human and machine identities.
- Automatically rotating passwords on resources based on workflows, scheduling, and on-demand, and providing a feature-rich API to enable other automation.
- Providing a historical record of passwords and secrets for access to backups, virtual machine snapshots, and disaster recovery environments.
- Enabling complete session management, including monitoring, recording, playback, and analysis of sessions in both real time and from a historical perspective for compliance auditing, forensics, and training purposes.
- Allowing access only through a secure gateway technology to prevent or monitor lateral movement, broker access via discrete components, and enumerate and enforce an appropriate access policy.
- Allowing for the privileged password management architecture to be applied to existing environments, with minimal distribution.

Based on the design put forth in this paper, all BeyondTrust Password Safe use cases can follow the models of least privilege and just-in-time access, integrate with ITSM solutions, and help achieve the benefits of a zero trust enclave deployment for administration and access.



Password Safe Architecture for a Zero Trust Enclave Model



In the diagram above, the Enterprise System is the Password Safe Console or REST API to assist with automation and nonhuman accounts. Also, note that the agent's role is tightly coupled to the policy engine. Its purpose is to release the passwords or secrets from encrypted storage to the subject, and to perform checks and rotation on existing managed resources. This includes accessing historical passwords and secrets and managing decisions from the policy engine, including the duration of access.

Any partial implementation of this model can be an improvement for privileged password management when considering a software-defined perimeter. This architecture is much more secure than allowing unrestricted access to any resource from any place using static credentials, especially if the password or secrets are single-factor and reused due to limitations of legacy solutions.



Secure Remote Access

Remote Access Risks Abound in the Modern Work Environment

As a security best practice, native remote access protocols should be disabled for corporate-issued computing device(s). Unfortunately, in many environments (especially for users working from home), this security control has not been implemented, and remote devices may be accessing corporate resources using inadequately secured remote access pathways.

The rationale behind enabling protocols like RDP, SSH, and VNC has been a source of contention between information technology and information security teams. One argument is the need for low-cost remote access technology natively supported by the operating system. This is generally advocated by IT due either to a lack of security experience, or budgetary constraints. Security and compliance teams, on the other hand, are wary of the inherent vulnerabilities, wormable exploits, and the lack of auditing and secure network routing of native protocols.

There must be a balance between these approaches. Authorized users need to initiate a secure remote access session to any device, any place, regardless of protocol. In addition, session monitoring, credential injection, and least privilege must be applied to overcome the security and compliance concerns governing an organization. These capabilities must be in place, whether the employee is working from the corporate office or from a remote location. This is especially true if the remote session is only for support of an application, and not a terminal-based remote session.


How can a user operate an application without the need for VPN (Virtual Private Network) and meet the tenets of zero trust? Secure Remote Access enabled for zero trust can solve this problem for applications and overcome traditional protocol routing challenges as well.



How a Zero Trust Architecture Can Overcome Native Remote Access Protocol Challenges

A ZTA can play a pivotal role in overcoming native remote access protocol challenges. A zero trust remote access implementation can accommodate almost any environment and allow a just-in-time (JIT) workflow with zero standing privileges (ZSP), meaning there are no accounts with privileges always-on and enabled.

The diagram below illustrates risks based on remote working within a decentralized, perimeterless environment, and reflects the combination of assets, connectivity, remote access technology, and prime locations for a threat actor to infiltrate an environment. In addition, based on normal connectivity, applications could operate remotely or using a remote session. To secure all activity, connectivity itself must be secured and all access provided only via a secure remote access technology.



In the graphic above, each "threat" represents a risk:

- Three Threats: Unacceptable, critical risk
- Two Threats: Medium-level of acceptable risk
- One Threat: Low risk for remote access
- Zero Threats: Best case for acceptable remote connectivity



Note that using a personal device with a business-issued VPN client is always a critical risk, regardless of whether the connection is wired or wireless. This is because the device is unmanaged and the organization has no control over how it is used, updated, or operated, including what other software or malware may be installed.

In this decentralized environment, threats exist when accessing sensitive internal resources from:

1. Personal or Bring Your Own Device (BYOD) hardware that is unmanaged, unpatched, multi-user, end-of-life, or otherwise susceptible to phishing or malware. In addition, BYOD users are typically their own local administrators, amplifying the risk.
2. Unsecure home networks based on WiFi connectivity where the connection is potentially unsecure, has a weak password, is wide open, or may allow a man-in-the-middle attack due to a common SSID or poor encryption. Also, other devices could compromise the wireless network or monitor communications. This includes privileged accounts outside of corporate governance used by home networks accessing consumer SaaS solutions.
3. VPN technology, which typically uses split tunneling and should never be installed on personal devices, that could compromise communications and provide a conduit for lateral movement via the flaws in the home network. Because VPN technology only operates at the network layer, it is unable to monitor remote sessions or applications. And typically, remote users only need application access (application layer) for a specific program or session.



Addressing Remote Access Risk with Zero Trust

To mitigate the threats, a combination of zero trust, IT managed devices, IT governance, and privileged access management can succeed where traditional technology alone may pose an unacceptable risk.

- **IT Managed** – Managed security controls for risk assessment, including core disciplines for vulnerability and patch management.
- **Connectivity** – Minimizing network risk with a wired connection in lieu of unknown wireless connectivity.
- **Privileged Access Management** – Remote access sessions are initiated at the application layer based on role, including credential obfuscation. This eliminates the need for network layer traffic per application per user. Privilege elevation is strictly controlled locally and across the network, also eliminating local administrative credentials and admin rights for end users.
- **Governance** – Documenting all privileged activity for compliance, including remote access user behavior.
- **Zero Trust** – Implementing a cloud-based management architecture for all remote access sessions honors zero trust, and strict application control is enforced. This applies the concept of zero trust with least privilege to all sessions, regardless of their origination or destination. It ensures the risks can be fully mitigated by never exposing root or administrative privileges outside of the extended perimeter, nor to the end user.



This combination of technologies and strategies works because you are not only securing the source device, but also minimizing network risk by:

- Using a wired connection
- Strictly controlling remote access
- Not performing any protocol routing like SSH
- Recording all session activity for compliance and behavioral analysis

Finally, applying the concept of zero trust with least privilege ensures the risks can be fully mitigated by never exposing root or administrative privileges outside of the perimeter. VPNs themselves cannot achieve this without the use of privilege management solutions.

Extend PAM & Zero Trust Controls Beyond the Perimeter with BeyondTrust Privileged Remote Access

VPNs cannot effectively manage risks inside of a defined perimeter or enclave. For zero trust to succeed, the network and environment need to be secured before a ZTA can be implemented.

BeyondTrust Privileged Remote Access is designed to manage sessions, and, with a few considerations, can be implemented using a zero trust model. The solution extends Privileged Access Management best practices beyond the perimeter. Privileged Remote Access enables organizations to apply least privilege and robust audit controls to all remote access required by employees, vendors, and service desks. Users can quickly and securely access any remote system, running any platform, located anywhere, and leverage the integrated password vault to discover, onboard, and manage privileged credentials. Thus, BeyondTrust Privileged Remote Access provides the granular, least privilege controls that are impractical with VPNs and many other commonly used remote access technologies.



Where native protocols or native tools are desired, BeyondTrust also offers infrastructure access management capabilities via our Privileged Remote Access solution. This capability combines native workflows, zero trust principles (like JIT and ZSP), and a comprehensive audit trail.

For technology workers, like developers and cloud ops engineers, the need to use VPNs (and sometimes multiple are needed in disparate environments) has been a necessary evil. While VPNs are encrypted and are generally secure, they do offer broader access than is typically needed. Additionally, from a workflow perspective, the user needs to remember what VPN should be used for each specific device, and they must know credentials for that system.

With Privileged Remote Access, the user can stay with a single console, use their native tools (i.e. SQL, SSH, and RDP), and inject credentials from a secure vault, and a full (and consolidated) audit trail will be generated for every session. This methodology gives the users their desired workflow for managing across infrastructure, and gives the organization the airtight security needed during an audit.

Core Privileged Remote Access use cases for zero trust include:

- Applying least privilege and robust audit controls to all remote access for employees, vendors, contractors, and service desk personnel.
- Managing and automatically injecting credentials into remote sessions and applications so the end user never sees or has knowledge of them for appropriate usage (integrates with BeyondTrust Password Safe for even more expansive privileged credential management).
- Securing infrastructure access—implementing a secured jump server with multi-factor authentication, adaptive authorization, and session monitoring for administrator consoles. This also applies to access that crosses trusted network zones.



- Providing access to web pages, such as the Azure or Microsoft 365 portal, through a locked-down and embedded Chromium browser.
- Enforcing boundaries between development, test, and production systems for SecDevOps best practices.
- Providing application-level microsegmentation that prevents users from executing applications and other resources they are not authorized to access.


Privileged Remote Access provides the following advantages over VPN alone.

How BeyondTrust Privileged Remote Access capabilities compare versus a typical VPN:

VPN vs. BeyondTrust Privileged Remote Access		
CAPABILITY	VPN	BeyondTrust
Remote Access	●	●
Secure Connectivity	●	●
Network Layer Access (Protocol Tunneling)	●	●
Encrypted Traffic	●	●
Application Layer Virtualization		●
Remote Desktop		●
Proxied RDP Access		●
Proxied VNC Access		●
Proxied SSH Access		●
Application Session Monitoring		●
Application Session Recording		●
Just-in-Time Access		●
Zero Trust Architecture		●
Privileged Access Management (PAM) Integration		●
Secure BYOT		●
ITSM Integration		●
Password Management & Credential Storage		●
Cloud or On-Premises Deployment (Physical or Virtual Appliance)		●
Agentless Access		●
Extensive Operating System & Platform Support		●
Prevention of Lateral Movement		●
Audit Trail & Session Reporting		●




Next, consider this diagram of a simplified NIST-based zero trust architecture for remote access:



Based on this design, all remote sessions should honor the models of least privilege and just-in-time access, integrate with ITSM solutions, and follow a zero trust architecture for policy and administration to meet established security best practices.

The diagram below illustrates how this occurs as a part of daily operations with BeyondTrust's Privileged Remote Access technology and zero trust:



*No direct access by remote users to an asset.
Zero Trust policy enforcement through policy engine, network routing, and dedicated remote access clients.*



Any partial implementation of this model can be an improvement in secure computing for a software-defined perimeter. This architecture is significantly more secure than allowing a home computer with VPN access into your environment to perform administrative functions.

Endpoint Privilege Management

Windows & macOS Users Have Excessive Privileged Access

For nearly all organizations, the simplest way to take control of their Windows and macOS endpoint estates is to implement least privilege by removing standing local administrative rights from their end users. Unfortunately, many organizations have not implemented this key security control to its full effect, leaving them vulnerable to cyberattacks of many kinds.

To successfully implement least privilege, the organization must find a solution that allows the following two goals to operate in harmony:

- **Productivity:** Many tasks that end users need to complete in their day-to-day work requires admin rights, including installing applications or changing system settings. If their admin rights are removed without the right solution in place, their productivity can be completely roadblocked.
- **Security:** The organization must remove all local admin rights, exercise control over which applications are allowed to be installed or run by different end user groups, and limit which tasks end users can execute on their endpoints, all to reduce their threat surface.



These goals underscore the importance of a zero trust architecture for endpoint privilege management, especially given the flexible modern work environment. Adopting the right endpoint privilege management solution as a part of a zero trust architecture can accommodate corporate office-based and work-from-home environments, and close the gap often found between managing privileges and applications, all while achieving strict governance over authentication.

Enabling a ZTA with BeyondTrust Privilege Management for Windows & Mac

BeyondTrust [Privilege Management for Windows & Mac](#) elegantly achieves both of the goals established for any successful endpoint privilege management solution. It removes excessive admin rights and applies modern application control, while dynamically elevating rights to end users so they have just the access they need at just the right time, without blocking their productivity.


Core Privilege Management for Windows & Mac use cases for zero trust include:

- Enabling zero trust security by removing excessive admin rights and eliminating persistent privileges.
- Applying granular application control and enforcing least privilege across applications, web browsers, systems, and other resources by giving users just enough access at just the right time.
- Stopping attacks that take advantage of trusted applications, like Office, Adobe, and web browsers, by applying built-in, context-based security controls.
- Drastically reducing the cyberattack surface and protecting against malware, ransomware, and phishing attacks.
- Providing a single audit trail of all user activity with graphical dashboards and reports.



Together, Privilege Management for Windows & Mac, along with Privilege Management for Unix & Linux (covered in the following section) and **Active Directory Bridge**, comprise BeyondTrust’s Endpoint Privilege Management solution. The solution combines least privilege management and application control to minimize the endpoint attack surface and eliminate unwanted lateral movement across your entire heterogeneous environment.

Consider this diagram of Privilege Management for Windows & Mac, according to the NIST ZTA:



*No local or remote trusted administrative privileges for users.
Applications are trusted for elevation or block-listed from runtime.*

If you are considering rearchitecting, redeploying, or modernizing your endpoint security model, you can achieve zero trust for privilege elevation, least privilege, and application control using this paradigm. This model satisfies all the NIST zero trust requirements and allows endpoint privilege management to extend the implementation to additional security models, such as just-in-time access.

From an architecture perspective, the entire enclave resides on the endpoint and represents the ideal state for any zero trust architecture and approach to segmentation.



Unix and Linux Privileged Risk Jeopardizes Sensitive Data and Assets

Unix and Linux administrators rarely, if ever, physically operate the keyboard directly behind their asset – if one is even connected at all. This is nearly always true, when operating within the cloud.

Administration, and even root access, is granted to the individual. That individual uses remote access technology and protocols like SSH to perform a task. If administration for Unix and Linux is primarily performed remotely, it begs the question—where does that remote administration originate? Is it on-premises and within a trusted network, or is the user operating remotely, such as from a home office, or maybe their couch? And, if you consider every remote access session is some form of privileged remote access, since access is actually granted from someone operating remotely, then the keyboard and mouse are nowhere near the actual computing device.

Today, unsecure networks of users working from home—or anywhere, for that matter—are an extension of our IT perimeters. Consequently, we have introduced new attack vectors and potential regulatory compliance issues that need to be addressed. For Unix and Linux administration, this represents an unacceptable risk since a typical business's most sensitive data and applications tend to reside on these mission-critical platforms.

All connectivity is dependent on secure credentials that follow the model of least privilege, just-in-time access, and single-use authentication.



Enabling a ZTA for Your Server Estate with BeyondTrust Privilege Management for Unix & Linux

BeyondTrust Privilege Management for Unix & Linux enables organizations to granularly control privileged access, achieve compliance, and vastly dial down cyber risk. The product can apply factors, such as time, day, location, and application or asset vulnerability status, to enable better privilege elevation decision-making. The product extends capabilities far beyond sudo, with centralized administration, session monitoring and management, file integrity monitoring, and powerful productivity enhancements.

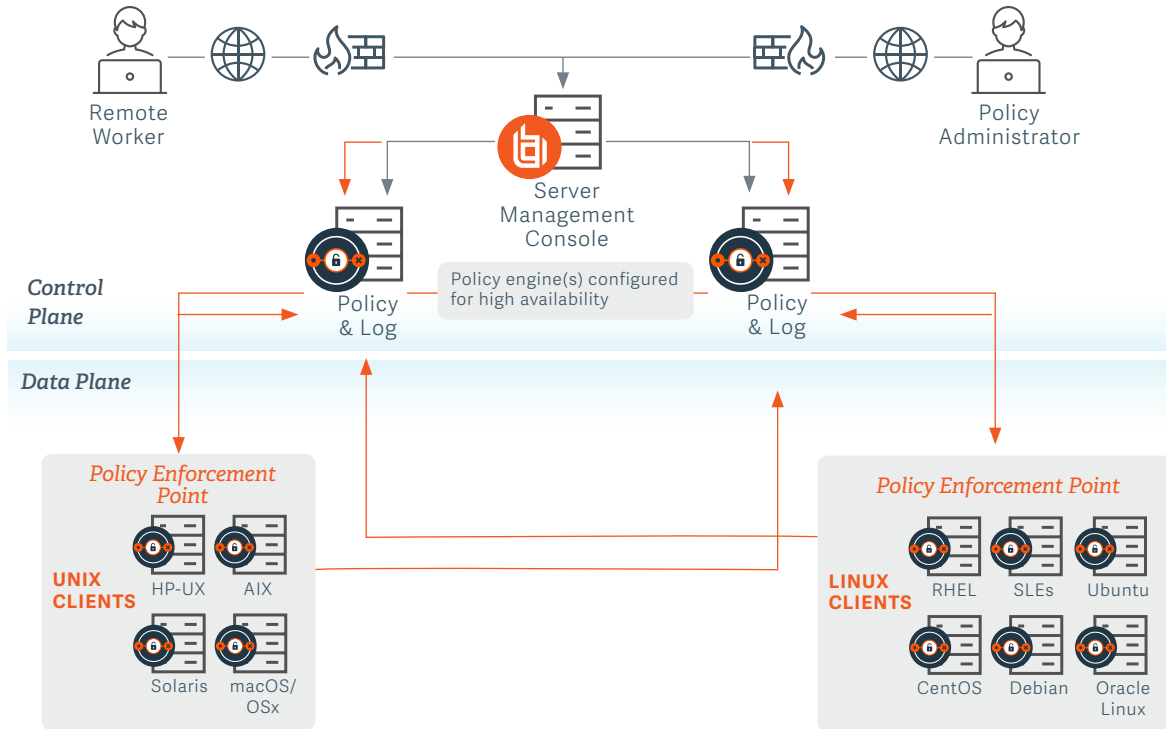
Core Privilege Management for Unix & Linux use cases for zero trust include:

- Removing admin rights and enforcing least privilege across all Unix/Linux endpoints.
- Enabling users to securely run specific commands and sessions remotely, without logging on as admin or root.
- Advancing toward a zero-standing privilege (ZSP) state by dynamically elevating privileges just-in-time for processes, application, etc.—but not for end users.
- Enforcing separation of duties and privilege separation to limit the privileges associated with any account or process.





Consider this diagram of Privilege Management for Unix & Linux, according to the NIST ZTA:



Facilitating Consistent Zero Trust Security Controls across Windows, Unix, & Linux with AD Bridging

BeyondTrust Active Directory (AD) Bridge, part of the Endpoint Privilege Management solution, plays an important role in supporting and simplifying a zero-trust strategy for Unix/Linux machines, while also helping organizations comply with other regulatory mandates.

AD Bridge centralizes authentication for Unix and Linux environments by extending Microsoft AD’s Kerberos authentication and single sign-on. Users leverage their AD credentials to access Unix and Linux systems for a seamless experience.



Organizations can attain consistent policies and controls across the enterprise by extending native group policy management tools to include settings for Unix and Linux, and transition users from desktops to remote machines or between systems, without requiring credential reentry. Leveraging Microsoft Group Policy across non-Windows platforms also enables centralized configuration management, reducing the risk and complexity of managing a heterogeneous environment.

Historically, security administrators have struggled to convert, apply, enforce, and audit NIST and other policies across the enterprise. AD Bridge makes it easy to map the NIST settings automatically and apply them to the machines—whether they are on-premises or in the cloud. NIST settings can be enforced by Active Directory GPOs (Group Policy Objects). Thus, if there is an attempt to alter those settings—even by root—they will be immediately reapplied and an alert will be generated indicating a policy violation. This alert can be forwarded to your SIEM, such as Elastic search, which is natively integrated with AD Bridge.

Core Privilege Management for AD Bridge use cases for zero trust include:

- Providing a single, familiar tool set to manage both Windows and Unix/Linux systems
- Centralizing management of group policies
- Providing thorough audit details to Audit and Compliance teams
- Expanding single sign-on (SSO) and file sharing, and controlling access to non-Windows systems.



Zero Trust Design Considerations

Technical Debt & Legacy Systems

Legacy applications, infrastructure, and operating systems are most certainly not zero trust-aware. They have no concept of least privilege, balk at application control, and lack privilege escalation models or authentication models that dynamically allow for modifications based on contextual usage. They have no concept of a remote worker, nor modern network architectures, let alone the cloud. They rely on direct network connectivity to operate them, and they lack session monitoring, or lateral movement monitoring capabilities. In fact, their security is probably highly network-dependent.

Redesigning, re-coding, and redeploying internal applications can be costly and potentially disruptive. Only a serious business need can justify undertaking these types of initiatives. Adding security controls to existing applications to make them zero trust-aware is not always feasible. It is likely that your existing applications have no facilities to accommodate the connection models in the specification and are not coded to operate in a resource enclave model as specified by NIST.

Therefore, depending on the architecture of your custom or legacy third-party application, consider using:

- Zero trust with privileged password management as the mechanism for any managed session, when connectivity should be monitored and managed.
- Zero trust with endpoint privilege management as the mechanism for remote worker privilege elevation, least privilege, and application control.



- Zero trust with privileged access management as the mechanism for remote worker authentication, least privilege, and session monitoring.
- Zero trust with secure remote access as the mechanism for remote worker connectivity.

This approach will allow you to implement zero trust controls and enhance your security posture, but without having to re-engineer established systems.

Any zero trust implementation requires a layered or wrapper approach to enable legacy systems. However, a pure zero trust approach entails enveloping all resources – regardless of their location – with these concepts. You can, however:

- Log remote session activity, record interactive screen sessions, and monitor events to look for potentially malicious behavior. This is a partial implementation of zero trust with Secure Remote Access and may be sufficient for some environments to mitigate risks. This is an important consideration when a single remote session may interact with multiple systems behind the scenes that are not being managed, and thus represents a high value to a threat actor.
- Log privileged activity, capture process launches, and monitor events to look for potentially malicious behavior. This is a partial implementation of zero trust with Privileged Access Management and may be sufficient for some environments to mitigate remote application access risks.
- Log screen activity, capture process launches, tally keystrokes, and monitor logs to look for potentially malicious behavior. This is a partial implementation of zero trust with Privileged Access Management. In conjunction with well-defined access control lists, the implementation is good enough to manage, monitor, and mitigate the risks of remote access to your Unix and Linux assets from virtually any source location.



Peer-to-Peer Technologies

If you think your organization does not use peer-to-peer (P2P) networking technology, you may be unaware of the default settings in Windows 10. Starting in 2015, Windows 10 enabled a peer-to-peer technology to share Windows Updates among peer systems to save Internet bandwidth. While some organizations turn this off, others may not know it exists.

This represents a risk of privileged lateral movement between systems. While no critical vulnerabilities and exploits have materialized for this feature, it does present communications that violate the zero trust model in much the same way a mail server needs to communicate with every end user's mail client. There should be no unauthorized lateral movement—even within a specified microperimeter.

In addition, if remote workers have protocols like ZigBee or other mesh network technology for IoT, you will find that they operate counter to zero trust. They require peer-to-peer communications to operate, and the trust model is based strictly on keys or passwords, with no dynamic models for authentication modifications.

With the above in mind, consider the following:

1. It is impractical to build a resource enclave for the entire installation. Resource enclaves should be well-defined and as small as possible. They require peer-to-peer communications to operate, and the trust model is based strictly on keys or passwords, with no dynamic models for authentication to support access via a gateway. They can be managed using a privileged password manager, but not in a zero trust architecture.



Therefore, if you decide to embrace zero trust and Privileged Password Management, consider hardening your resources to allow an enclave model, and disallow any inappropriate network communications to anything within that defined enclave.

While there will be exceptions for communications within the enclave itself, conceptually, the perimeter stops at the collection of resources. Privileged Password Management using a gateway should be the vehicle for allowing any external access.

2. If you decide to embrace zero trust and Privileged Access Management, consider hardening your endpoint security model to disallow any inappropriate network communications on the same subnet as the source or destination. While there may be exceptions for devices like local printers, conceptually, the perimeter stops at the device itself, is defined through applications and privileges, and remote access only controls the connection point-to-point.
3. Investigate whether your Unix and Linux environment is using similar technologies, has P2P, or has mesh network technologies enabled—even for policy or network management. These present a huge stumbling block to embracing any trusted remote access paradigm because lateral movement will always inherently be present in some form.



Partnering with BeyondTrust for Your Zero Trust Journey

Public and private sector organizations across the world are partnering with BeyondTrust to implement zero trust security tenets and architectures.

> Additional Resources

CASE STUDY [Investec's Journey to Zero Trust: From Theory to Practice](#)

CASE STUDY [How Oxford Properties Group Enabled a Zero Trust Architecture with BeyondTrust](#)

WHITEPAPER [Mapping BeyondTrust Capabilities to NIST Zero Trust \(SP 800-207\)](#)

This paper explores the 7 tenets of zero trust in detail and shows how BeyondTrust can help you align to them. The resource also maps BeyondTrust PAM use case to zero trust tenets in additional detail.

If you're ready to discuss your zero trust project with BeyondTrust, we welcome you to [contact us](#).



About BeyondTrust

The BeyondTrust platform



ABOUT BEYONDTRUST

BeyondTrust is the worldwide leader in intelligent identity and access security, empowering organizations to protect identities, stop threats, and deliver dynamic access to empower and secure a work-from-anywhere world. Our integrated products and platform offer the industry's most advanced privileged access management (PAM) solution, enabling organizations to quickly shrink their attack surface across traditional, cloud and hybrid environments.

With a heritage of innovation and a staunch commitment to customers, BeyondTrust solutions are easy to deploy, manage, and scale as businesses evolve. We are trusted by 20,000 customers, including 75 of the Fortune 100, and a global partner network.

Learn more at beyondtrust.com