WHITE PAPER

THE NIST ZERO TRUST APPROACH

How To Achieve It with Unix & Linux Remote Access

BeyondTrust

TABLE OF CONTENTS

1

Introduction - What Is Zero Trust?

While zero trust has become a trendy catchword in IT, this model is very specific about how things should be designed and operate, and it may not work for everyone. In practice, zero trust is best suited for new deployments, or to strictly control access of a user or automation to sensitive resources, especially when they are connecting remotely.

When applying the granularity of privileged access management (PAM), zero trust can ensure all access is managed and documented for appropriate behavior. Today, this is a particularly important challenge to solve as so many IT administrators—like many other employees—are working from home networks and connecting via Wi-Fi. This environment presents heightened risks to highly sensitive Unix & Linux resources being connected to and remotely administered.

This white paper will:

- Define zero trust
- Dissect the security implications of Unix/Linux administration from home networks
- Present the zero trust model developed by NIST
- Discuss the practical implementation steps of zero trust on Unix & Linux systems
- Briefly highlight how BeyondTrust enables organizations to achieve zero trust on Unix & Linux systems

Zero Trust Defined

By definition, a zero trust security model advocates for the creation of zones and segmentation to control sensitive IT resources. This also entails the deployment of technology to monitor and manage data between zones, and more importantly, authentication within a zone(s), whether by users, applications, context, attribution, or other resources.

Zero trust redefines, and strives to remove, the architecture of a trusted network inside a defined perimeter. To that end, zero trust implementations can be on-premise or in the cloud. Zero trust establishes that only trusted resources should authenticate and interact with each other, regardless of network architecture. This vision of zero trust is relevant today because technologies and processes like the cloud, virtualization, SecDevOps, edge computing, edge security, software defined perimeters, passwordless administration, and IoT have either blurred, or dissolved altogether, the idea of a traditional firewalled and segmented network. The lack of a perimeter has given rise to concepts like control planes and data planes. The control plane offers a vehicle to manage any resources that require authentication, while the data plane does likewise for the resources themselves that require authentication to operate.

By definition, a zero trust security model advocates for the creation of zones and segmentation to control sensitive IT resources.

Within each of the planes, zones can be created to manage resources and can be delegated using micro-segmentation down to the host or data layer to enforce a zero trust model. This implies that a resource, like a server, or even a database, can have multiple zones to support the data collection and monitoring needed to achieve zero trust. Zero trust essentially establishes a dynamic model of trust, verification, and a continuous re-evaluation of trust for further access, to prevent any unauthorized lateral movement. The control for authentication is established in one place and the operation executed in another.

2 Remote Administrative Workforces

For Unix and Linux administrators, it is rare, if ever, that they 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 and 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-premise 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.

In recent times, the COVID-19 crisis has upended the way we live and work. This period will long be studied by historians, economists, workforce productivity experts, physicians, and many others. However, will people also pinpoint 2020 as the year during which enabling secure remote access became absolutely essential for the sake of business continuity? How did organizations effectively make the shift? How did we certify the access was appropriate when the source of the changes had occurred outside of our secure network perimeter? Unfortunately, many organizations are still unable to answer that basic question, which has huge security ramifications.

Insecure home networks have introduced new attack vectors and potential regulatory compliance issues that need to be addressed.

Today, our home networks are now serving entertainment, school, work, and providing an active conduit into our business to provide secure Unix and Linux administration. As a result, we are allowing our insecure home networks to be an extension of our IT 'perimeters' to perform tasks in our business environments. 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 tends to reside on these mission-critical platforms.

Figure 1: The risks based on the combination of assets, connectivity, remote access technology, and prime locations for a threat actor to infiltrate an environment based on a privileged remote worker.

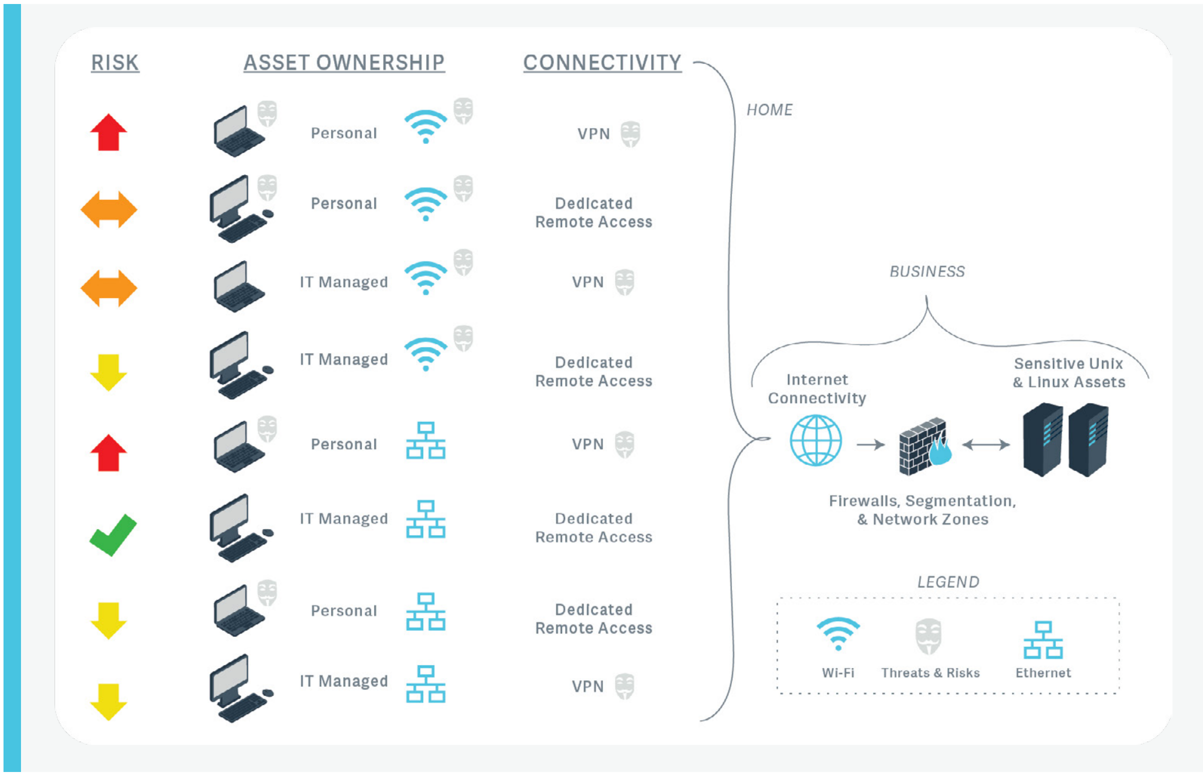

In Figure 1 above, each "mask" represents a risk:

- **Three Masks:** Unacceptable critical risk
- **Two Masks:** Medium level of acceptable risk
- **One Mask:** Low risk for remote access
- **Zero Masks:** 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.

A combination of zero trust, IT managed devices, and privilege management for Unix and Linux can succeed where VPN technology alone may pose an unacceptable risk.

To briefly summarize, 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 may be susceptible to phishing or malware.
2. Insecure home networks based on Wi-Fi connectivity where the connection is potentially insecure, has a weak password, is wide open, or may allow a man-in-the-middle attack due to a common SSID or poor encryption. In addition, other devices could compromise the wireless network or monitor communications.
3. VPN technology using split tunnelling, unmitigated security flaws, end of life software, or installed on personal devices could compromise communications and provide a conduit for lateral movement into the environment based on flaws in the home network.

To mitigate the three threat areas described previously, a combination of zero trust, IT managed devices, and privilege management for Unix and Linux (including dedicated remote access technology) can succeed where VPN technology alone may pose an unacceptable risk. This combination of technologies and strategies works because you are not only securing the source device, but you are also minimizing network risk with a wired connection, strictly controlling remote access, not performing any protocol routing like SSH, and recording all session activity for compliance and behavioral analysis. Finally, applying the concept of zero trust with least privilege on Unix and Linux 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.

3

Success with the NIST 800-207 Zero Trust Model

So, what does a successful deployment of a zero trust model look like with regards to Unix and Linux privileged remote access?

The NIST 800-207 Zero Trust Model clearly states that the goal of zero trust is to focus security on a small group of resources (zones) in lieu of wide network perimeters or environments with large quantities of resources interacting "freely". It is a strategy where there is no implicit trust granted to systems based on their physical or network location (local area network, wide area networks, and the cloud), but rather access is granted by a trusted source for either a user or application.

Consider this simplified NIST core zero trust architecture presented below.

Figure 2: Representation of the NIST 800-207 core zero trust architecture.

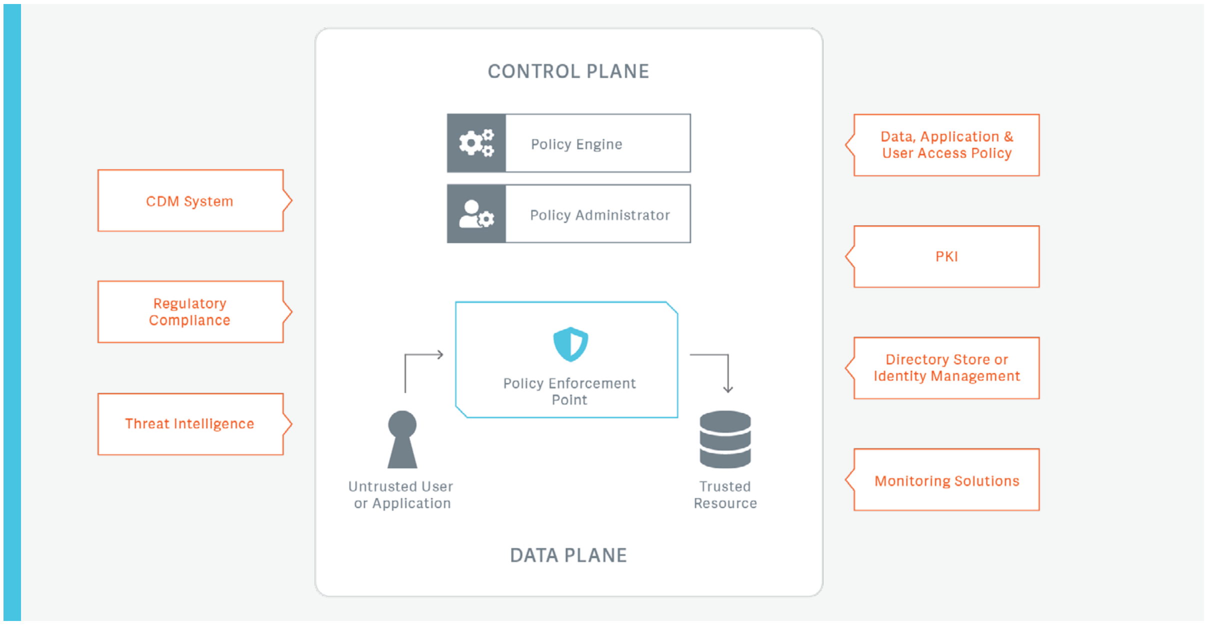

The key components of the control plane and data plane are typically found in privileged access management solutions:

- The *Policy Engine* is responsible for the decision to grant access to a resource. It uses as much data as it can based on roles, attributes, and threat intelligence to determine if access should be granted.
- The *Policy Administrator* is responsible for establishing the connection between a client and a resource. It provides the negotiation between the resources to “state” that the connection is allowed.
- The *Policy Enforcement Point* is responsible for enabling, monitoring, and terminating the connection between the Untrusted resource (user or application) and Trusted Enterprise Resource.

If we map this to BeyondTrust’s Privilege Management for Unix & Linux solution, we find:

- The *Policy Engine* can be found in management capabilities of the rules, policies, and log engine governing least-privilege endpoint access, and the role and attribute-based access models defined by the Policy Administrator.
- The *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 BeyondInsight Unix & Linux Management Console.
- The *Policy Enforcement Point* is the least-privilege client installed on Unix and Linux Hosts. 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.

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

***Figure 3:** Illustration of the BeyondTrust Privilege Management for Unix & Linux zero trust architecture in a basic high availability implementation.*

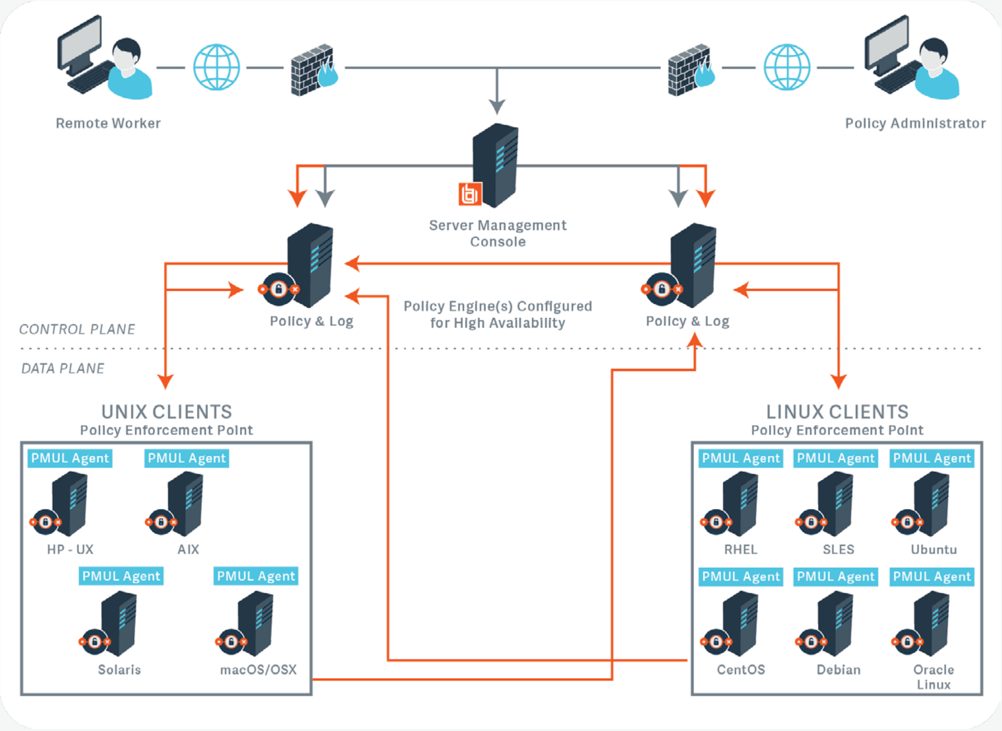

Privilege Management for Unix & Linux uses this model for user and application access regardless of the network topology. And, if access is secured using a dedicated privileged remote access technology, you can achieve zero trust access for Unix and Linux without the need for protocol tunneling or a VPN. While any partial implementation of this can be a great step toward more secure computing, it represents only a hybrid approach. An example of this would be to use a VPN technology in lieu of a dedicated privileged remote access solution. In fairness, this scenario is much more secure than allowing a home computer with VPN access into your environment to perform administrative tasks on your sensitive Unix and Linux resources – especially as root.

4 Design Considerations for Zero Trust on Unix & Linux

Zero Trust has been developed in response to industry trends that include remote users, dissolving network perimeters, and dynamic cloud-based assets. It focuses on protecting resources, not logical network segments, as network segmentation is no longer seen as the prime component to the security posture of the resource. This, in itself, begins the discussion of why zero trust combined with Privileged Access Management can solve privileged remote worker challenges given the current worldwide coronavirus crisis. The acceleration of remote working caused by the pandemic will result in a long term, sustained expansion of remote working.

With the above in mind, consider the following three topics as you embrace this zero trust model for Unix & Linux:

Technical Debt

If your organization develops its own software for consumption, and the applications are more than a decade old, you have technical debt. Redesigning, recoding, and redeploying internal applications can be costly and potentially disruptive. There needs to be a serious business need to undertake these types of initiatives.

Adding security parameters to existing applications to make them zero trust-aware is not always feasible. Odds are, your existing applications have no facilities today to accommodate the connection models in the specification, nor are they coded to operate in small groups as specified by NIST. Therefore, depending on the architecture of your custom application, consider using zero trust and privileged access management as the mechanism for remote worker authentication, least privilege, and session monitoring. This will allow it to be a successful add-on to your existing solution without re-engineering established systems.

The acceleration of remote working caused by the pandemic will result in a long term, sustained expansion of remote working.

Legacy Systems

Legacy applications, infrastructure, and operating systems are most certainly not zero trust-aware. They have no concept of least privilege, session monitoring, or lateral movement, and they do not possess authentication models that dynamically allow for modifications based on contextual usage. In fact, their security is probably highly network-dependent.

Any zero trust implementation requires a layered or wrapper approach to enable these systems. A pure zero trust approach entails enveloping all resources – regardless of their location – with these concepts. You can, however, 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 secure 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 are probably 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 are still not even aware it exists. This represents a risk of privileged lateral movement between systems that is fundamentally uncontrolled.

Starting in 2015, Windows 10 enabled a peer-to-peer technology to share Windows Updates among peer systems to save Internet bandwidth.

While no vulnerabilities and exploits have materialized yet for this Windows 10 feature, it does present communications that violate the zero trust model. There should be no unauthorized lateral movement—even within a specified microperimeter. In addition, if you use protocols like ZigBee or other mesh network technology for IoT, you will find that they operate completely 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.

Therefore, if you decide to embrace zero trust and privileged access management, please investigate whether your Unix and Linux environment is using similar technologies, has P2P, or 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.

5 Conclusion

Today, enterprises and public agencies are supporting vastly larger remote workforces than most of them had ever anticipated in years past. This remote workforce includes highly privileged IT administrators operating from home, remotely performing sensitive tasks on critical Unix and Linux resources.

Fortunately, organizations can largely mitigate these dangerous risks via a simple deployment of BeyondTrust Privilege Management for Unix & Linux using the NIST 800-207 zero trust model. Thousands of customers worldwide rely on BeyondTrust's solution to secure these critical resources from potential abuse when being remotely administered.

ABOUT PRIVILEGE MANAGEMENT FOR UNIX & LINUX

BeyondTrust Privilege Management for Unix & Linux is an enterprise-class, gold-standard privilege management solution, enabling organizations to granularly control privileged access, achieve compliance, and vastly dial down cyber risk. Apply factors such as time, day, location, and application, or asset vulnerability status, to make better privilege elevation decisions. Enable users to securely run specific commands and sessions remotely, without logging on as admin or root. Extend capabilities far beyond sudo with centralized administration, session monitoring and management, file integrity monitoring, and powerful productivity enhancement.

ABOUT BEYONDTRUST

BeyondTrust is the worldwide leader in Privileged Access Management (PAM), empowering organizations to secure and manage their entire universe of privileges. Our integrated products and platform offer the industry's most advanced PAM solution, enabling organizations to quickly shrink their attack surface across traditional, cloud and hybrid environments.

The BeyondTrust Universal Privilege Management approach secures and protects privileges across passwords, endpoints, and access, giving organizations the visibility and control they need to reduce risk, achieve compliance, and boost operational performance. Our products enable the right level of privileges for just the time needed, creating a frictionless experience for users that enhances productivity.

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 78 of the Fortune 100, and a global partner network.

beyondtrust.com

V2020_07_ENG