The Threat Within:

The Case for Zero Trust Access Control in the Era of Hybrid-Cloud Computing

The Threat Within: The Case for Zero Trust Access Control in the Era of Hybrid-Cloud Computing

Real Incidents Translate to Significant Losses

- 2011: A former employee at the U.S. subsidiary of Japanese pharma Shionogi plead guilty to deleting 15 business-critical VMware host systems, costing the company $800,000.
- 2010: An IT employee at Bank of America admitted that he hacked the bank's ATMs to dispense cash without recording the activity.
- 2010: A contract programmer fired by Fannie Mae was convicted of planting malicious code intended to destroy all data on nearly 5,000 internal servers.
- 2010: A Goldman Sachs programmer was found guilty of stealing computer code for high frequency trading from the investment bank when he left to join a startup.
- 2010: A Utah computer contractor pleaded guilty to stealing about $2 million from four credit unions for which he worked.

Executive Summary:

In protecting your data, you've invested considerable time and resources into firewalls, IDS/IPS, malware detection and numerous other security measures. That's how you keep the bad guys out, right?

Well, not exactly. Because all of the perimeter defenses in the universe can't provide a shred of protection against what could amount to your enterprise's most formidable threat: the very people you've entrusted to run and manage these systems.

That's right. The weak link may be working in your office right now, or may be accessing your critical infrastructure from a coffee shop.

They're the administrators paid to support applications, database servers, and network devices that drive strategic business objectives. But these individuals are granted what amounts to "keys to the kingdom" level of access, bringing significant levels of risk. For the purposes of this white paper, we'll call them privileged insiders.

We know you think of yourself as circumspect. You bring good people on board, and you give them what they need to do their jobs. But good people can lapse into bad routines.

They may use the same administrative passwords for years, store them in unprotected spreadsheets and other documents, or share them too freely. Or perhaps a busy senior administrator will hurriedly grant a new hire "keys to the store" when limited access would have sufficed. In fact, some servers and network devices have just a single administrative account. So granting all-or-nothing access may be the only option. And with the rise in the 24/7 mobile mentality, a system admin working remotely at a coffee shop can unwittingly expose ID and password information by using an insecure wireless connection.

These and other scenarios make it easy for modern adversaries who are well funded and equipped to plot, scheme and otherwise invent clever ways to exploit your most trusted users. Things can go bad quickly. Especially when these adversaries can take advantage of the treasure trove of information shared in social media to launch highly focused spear-phishing attacks, targeting privileged insiders and snatching up precious administrator credentials.

Insert a rogue administrator into the equation and then you really have a potential mess. An admin already has privileged access. He knows the network. He can avoid existing security controls, or access network and security systems to cover his tracks.

When the inevitable breach occurs, you realize the extent of exposure created by easy access, excessive privileges,

and poorly protected credentials. And because you don't have firm control over who went where and what exactly each person did, it's nearly impossible to find who was responsible and how the damage was done.

This is not mere theory. Both regulators and auditors are taking notice. WikiLeaks served as a clarion call, exposing the immense dangers of "insider threats." As a result, many regulators are updating requirements and auditors are scrutinizing security programs to identify this important risk. They certainly don't want the equivalent of a WikiLeaks scandal to unfold under their watch, and are willing to hand out stiff penalties for compliance violations. HIPAA violations can cost up to $50,000 each. Fines for PCI non-compliance can total up to $100,000 per month.

"Hackers target these individuals and their credentials to gain access to critical systems."

This white paper carefully examines privileged insider risk and how organizations are addressing it with what's called the least-privilege approach — one with real merit but significant limitations for modern enterprises. To adequately address the privileged insider risk in today's environment, companies must go several steps beyond least privilege to a level we call zero trust.

Four Key Vulnerabilities to Consider:

- Over-privileged users.
- Unprotected privileged account credentials.
- Shared administrative accounts.
- Risks introduced by hybrid-cloud computing.

Risks and Vulnerabilities

Simply stated, privileged insider risk exists because the network infrastructure supporting your business and housing your critical data demands constant upkeep. The individuals who perform this work are granted elevated rights and privileges to do so. They can go rogue, or inadvertently lapse into patterns that create vulnerabilities. Hackers target these individuals and their credentials to gain access to critical systems.

There are four key vulnerabilities to consider:

Over-privileged users. Some individuals have more access and authority than what's necessary to perform their duties. As administrators move on from project to project, they typically collect new permissions as they go, often without relinquishing credentials from past projects. Sometimes their accounts remain active even after they've left the organization or finished a contract. In many cases, these users are not even employees of your organization – outsourced support teams, contractors, and third-party vendors who need to access systems for support and maintenance purposes also collect credentials – while working with dozens of different organizations.

Unprotected privileged account credentials. Many organizations have not taken adequate steps to protect privileged account credentials, such as user IDs and passwords. Individuals and teams store them in unencrypted files and exchange them in email and chat sessions. In addition to unprotected privileged user credentials, many large enterprises have hundreds or even thousands of unprotected passwords in applications. Such credentials enable one system to communicate and exchange data with other applications or databases. And these passwords are often placed directly into source code

What's at risk?

42%
OF RESPONDENTS SAY PRIVILEGE USER SECURITY RISKS WILL INCREASE

"Forty-two percent of respondents say the risk to organizations caused by the insecurity of privilege access users will increase over the next 12 to 24 months and 42 percent say it will stay the same. Cloud-based applications, virtualization and regulations or industry mandates are the primary reasons for this belief."

Source: December 2011 Ponemon Institute Study: Insecurity of Privileged Users: Global Survey of IT Practitioners.

"…many large enterprises have hundreds or even thousands of unprotected passwords in applications."

"…organizations can no longer rely on simply trusting individuals or devices that are inside the perimeter. There is no 'inside' anymore."

and scripts or stored in clear text in configuration files. Anyone with access to these credentials gains all the privileges associated with the accounts they unlock – what if this is a hacker using malware and keystroke logging program to grab them?

Shared administrative accounts. When multiple individuals use the same admin account, you really have two problems rolled into one. Since these credentials are shared, they're typically changed infrequently. Further, if multiple admins use the same account, everyone using the account is effectively anonymous, since there is no reliable way to know who was logged on when or what they did. This can be a really sticky issue. Some older servers and network and security devices support only a single admin account. In other cases, where the technology enables multiple administrative accounts, the issue can be the cost and complexity of setting up and managing policies for multiple accounts, especially across a range of devices and operating systems.

Risks introduced by hybrid-cloud computing. The wholesale movement to virtualized and public/private cloud computing expands attack surfaces in ways that were, until recently, unimaginable. A high-priority concern is the introduction of new control points – virtual hypervisor and cloud management consoles – bringing traditional privileged user management challenges to new technology platforms. Application and security architects must grapple with the task of migrating existing controls into these new environments. With the ability to create – or destroy – thousands of servers and computing resources in an instant, these technology platforms also introduce a number of new risks that must be addressed.

The addition of these new administrative environments exacerbates total cost of ownership issues for IT administrators. Organizations already have to deal with the problems introduced by so-called "islands of identity." Privileged users, and their roles and responsibilities, have already been defined across mainframe, server, workstation, network, and application/database technology domains. Without a way to leverage that existing identity store, administrators are faced with the daunting challenge of replicating – and maintaining those identities in additional environments. The potential for wasted effort and inefficiency is high.

All of these factors have fundamentally eroded the concept of perimeter defenses and have changed the trust model forever. Since critical infrastructure, and the employees and third parties who manage it, can now be inside or outside of the perimeter, organizations can no longer rely on simply trusting individuals or devices that are "inside" the perimeter. There is no 'inside' anymore.

When "Least Privilege" Isn't Enough

When organizations understand a serious problem exists and respond to it, they often turn to the concept of "least privilege." Least privilege (a.k.a. minimal privilege) access control is an excellent start. It dictates any specific user/module/process/program is

allowed to access solely the systems/ data required for its legitimate purposes. You may recall the DAPE principle: Deny All Permit by Exception. Best practices for privileged access control from Gartner and other IT analysts include the minimization of accounts, limitation of scope, restriction of privileges and use of risk-appropriate authentication. The least-privilege principle also accounts for what authorized parties can do on these systems (for example, limiting the commands they can run).

But least privilege is only a first step. It doesn't cover the full set of controls an enterprise must implement to tackle today's expanding threat landscape. We call this advanced set of access controls zero trust.

Zero trust access control includes all of the components of least privilege, but extends far beyond. To be clear, least privilege as a model serves as a great starting point. But it alone is not enough. It was not designed with the modern nature of threat in mind. It cannot account for the complexities stemming from mobility, virtualization, the cloud and the "new perimeter." Zero trust does.

Zero trust expands on the least-privilege concept to include a full range of controls, like monitoring, alerting and session recording. The extra Zero trust controls minimize risk and ensure organizations are well protected and able to detect insider threats. Zero trust also ensures they can respond rapidly when issues arise, armed with information that's essential to fully investigate breaches.

Pillars of Zero Trust

Xceedium's zero trust model includes eight essential controls:

1. Control Access. Limiting access is a fundamental tenet of the least-privilege principle and the cornerstone of zero trust. Zero trust implements DAPE through policies that explicitly grant access to network nodes — both one-to-one mappings of individuals to nodes, or group-level provisioning. To meet our zero trust standard, the access-control mechanism must conceal nodes that are not explicitly included in an individual or group access policy. New cloud and virtualized environments – which enable the creation and elimination of resources at a breakneck pace – pose a particular challenge. It's essential access controls be as dynamic as the environment they're meant to protect

2. Protect and manage credentials. Protecting passwords and other credentials is a lynchpin of access control. First, administrative-level passwords need to be stored in an encrypted vault and remain encrypted when traversing the network. This applies to credentials for privileged users and those used in application-to-application or application-to-database scenarios. Further, policies for creating, updating, and decommissioning of passwords must be enforced and the system must support multi-factor authentication prior to releasing passwords. To meet our zero trust benchmark, the system must also initiate privileged sessions without users seeing or knowing an actual password. This is an important concept. If privileged users don't have any access to the keys to the

"Least privilege (a.k.a. minimal privilege) access control is an excellent start … But least privilege is only a first step."

"Areas where organizations most need to improve include: monitoring privileged users' access when entering administrative root level access, understanding privileged users entitlements that violate policy and enforcing access policies in a consistent fashion across all information resources in an organization."

Source: December 2011 Ponemon Institute Study: Insecurity of Privileged Users: Global Survey of IT Practitioners

What's in a Name?

The severe risks privileged insiders pose has inspired a budding industry. But a single name for the segment has not yet arrived.

CURRENT INDUSTRY TERMS:

Gartner – Privileged Account Activity Management (PAAM), including:

- Super User Privilege Management (SUPM)
- Shared Account Password Management (SAPM)
- Application-to-Application Password Management (AAPM)

IDC – Privileged Identity Management (PIM)

Forrester – Privileged User Management (PUM)

kingdom, they can't give them away – on purpose or accidently. Of course there are emergency 'break-glass' situations in which passwords need to be known, so the system must also be able to issue single-use or time-limited credentials in such situations.

3. Prevent anonymous activity on shared accounts. For zero trust, the system will clearly — without ambiguity — link a shared, privileged admin account (like root, su, or admin) to the specific person using it. As previously indicated, older systems will often have only one admin account. But even with newer technology, many large enterprises choose to minimize the number of active administrative accounts because of the overhead costs of managing them. Regardless of the setup, you must know exactly who logged on to a shared account and what exactly he or she did while on it. Otherwise, any investigation of an incident is hobbled from the beginning.

4. Contain users. Containing users remains an important aspect of comprehensive access control. If users are able to 'leapfrog' from a device for which they have been granted access to another device for which they have not, then access control is really only first-system access control. Windows environments, for example, make leapfrogging incredibly easy with RDP hopping. Controlling leapfrogging is one half of zero trust containment; vaulting and never letting users view passwords is the other. So even if a user gains logical or physical access to a device by going around the technical-access controls, he or she won't have the password for entry.

5. Control commands. Limiting access and containing users isn't enough. You also must control what can be done once an individual logs into a system. Zero trust mandates the limiting of commands that can be executed to only those necessary.

6. Record sessions. Recording sessions will establish documentation of systems control. As Gartner notes, "The ability to monitor privileged account activity is essential to a privileged account activity management (PAAM) solution. Monitoring is what proves the effectiveness of the controls used to manage these accounts."

7. Log everything. Granular logging of everything includes the logging of all interactions privileged users have with a system (as they go about their network and database management tasks, for example). However, it also includes logging any interactions a super user has with the system providing zero trust (e.g., creating, reading or updating policies; viewing recorded sessions; etc.). This degree of awareness ensures the organization has everything it needs to prove compliance and produce the forensic data required to fully investigate issues.

8. Alert for policy violations. Considering the sheer quantity of events that modern security operations centers must deal with, the zero trust system must raise a red flag when something bad may be occurring. Despite enormous improvements in log management and SIEM systems, it remains very difficult for systems to detect anomalous behavior by those with privileged accounts. So proactive alerting that is tied to policy violations serves as a critically important detection mechanism.

Zero Trust in the New Enterprise

Large organizations have unique requirements that must be met in order for a zero trust access control solution to be viable in their complex environments. Some of the most critical requirements are:

Support for the Hybrid Cloud

Achieving zero trust controls in hybrid-cloud environments requires four core capabilities. First, the ability to fully implement essential Zero Trust controls in technologically diverse environments (such as traditional data centers, virtual infrastructure, and public/private clouds) to ensure a robust and effective level of security. Second, the ability to federate identity across technology domains to minimize administrative overhead and associated expenses. Third, provide support for unified policy definition, enforcement, and reporting across the whole of the hybrid cloud, which is critical for ensuring and demonstrating compliance with audit and regulatory standards, as well as reducing administrative overhead. Lastly, the ability to automatically enforce policies. Virtualization and cloud computing have created a significantly more dynamic environment – where new servers can appear and disappear by the dozens or hundreds. In this elastic atmosphere, policy enforcement has to be automatic. It is simply not practical to have a program that requires a human to go to the keyboard to apply policies to new infrastructure components at cloud speed.

Reliability, Availability and Scalability

Organizations of all sizes need technologies to work reliably, but for large enterprises this is an even higher priority. Technologies also must accommodate clustering and other high-availability configurations for performance, failover, and disaster recovery scenarios — especially for a system that controls access to your most important infrastructure.

Low Total Cost of Ownership (TCO)

In large enterprises, which may have thousands or hundreds of thousands of network nodes, the administrative cost of managing solutions can bring an unbearable tax. Enterprise-ready solutions must minimize the number of administrators required to install, configure, and deliver ongoing policy management. Asking your team to jump from one piece of software to another — or to learn multiple, different user interfaces — just to define and maintain coherent policies places a huge administrative burden on the organization. Systems also should be architected to leverage network services and minimize the installation and management of 'agents' on end nodes.

Furthermore, a viable access control solution must reduce the friction that privileged users experience when accessing the target nodes they have to manage — especially since they might access and manage many network nodes every day. Security and control is great, but adding any extra time and effort to already busy administrators is not only costly — it may lead to an outright revolt.

Peace of Mind

"With Xceedium we have an all-in-one solution for these higher risk users which gives us the peace of mind that we are meeting our objectives to safeguard our network and the sensitive information it contains."

Security Expert, Department of Homeland Security

7 Seven Signs that You May Need Xsuite

- Your organization must comply with security regulations (HIPAA, FISMA, PCI DSS, NERC/CIP, SOX or others)
- You've suffered an insider breach, but you can't pinpoint who was responsible
- Your database/network administrators are allowed to anonymously access privileged accounts or systems
- You use third-party outsourcers or contractors to manage critical IT infrastructure
- You have hard-coded passwords to key applications and databases in scripts and application code
- Your organization is a global enterprise
- You have important business, mission or operational data to protect

Wide Platform Support

An enterprise-ready solution should include support for a wide number of devices ranging from Windows, Unix and Linux to mainframes, network devices, and security systems.

Integration

No enterprise security system can be an island. A privileged access control system is no exception. It must fit in well with the rest of the security and IT operations infrastructure and integrate with IAM, authorization, SIEM and other security solutions, as well as with network management and workflow tools.

Auto Discovery and Policy Provisioning

Traditional approaches to defining and provisioning policy will fail in today's highly dynamic hybrid cloud environments. When new systems can be created by the hundreds – or thousands – these manual approaches can't keep pace with the rate of change experienced in hybrid data centers. Instead, security systems must be able to automatically discover resources as they're created, and apply policy and grant access automatically.

Industry Shortfall; Home-Grown Solutions

Multiple software vendors have created products to address the escalating risks privileged insiders pose, but many are point solutions constructed of multiple, but not well-integrated, software modules.

Many organizations also have tried to leverage existing security tools, cobbling together solutions such as firewalls, router ACLs, jump-boxes, and VDI technologies to mitigate these risks.

Unfortunately, piecemeal solutions often don't deliver the necessary protections, and they certainly don't measure up to the zero trust benchmark. These home-grown solutions often result in new and convoluted policy and configuration rules — adversely impacting the performance and effectiveness of the tools' primary objective. All the while making the solution more brittle and costly to manage.

The Bottom Line:

The imposing nature of privileged insider threats, combined with rapidly changing business and technology dynamics, has placed enterprise networks and data at significant risk. Twisting traditional security defenses to help mitigate the risk can lead to greater exposure.

Many organizations are revisiting the concept of least privilege. And while least privilege is an excellent starting point, it simply falls short. Zero trust access control builds on the least-privilege concept by identifying eight essential capabilities required to adequately protect organizations from the elevated insider risk large enterprises with modern IT systems face.

Xceedium Xsuite — A Proven Zero Trust Access Control Solution

Xceedium's Xsuite is a comprehensive, integrated and highly scalable solution; its largest single enterprise user now supervises more than 200,000 nodes. Xsuite implements each of the essential zero trust controls, while simplifying manageability through a unified policy management system.

Xsuite is unique in its ability to natively support the whole hybrid-cloud. Core functions, such as access control policy definition and enforcement, credential vaulting and protection, and more, are supported for all environments in a consistent manner. It's possible to define consistent uniform policies across different platforms – and enforce those controls across the entire hybrid-cloud environment.

Xsuite extensions enhance its core functions with customizations and specialized integrations for IBM mainframes (3270- and 5250-based connections), virtualized infrastructure, and Amazon AWS-based public and private clouds.

Xsuite has received notable recognition from respected industry analyst firms:

"We are seeing rapidly rising demand for solutions addressing insider threat," Sally Hudson, a research director at IDC, has said. "Products that tie together privileged access control with privileged user and application password management are in high demand due to increased need for meeting regulatory compliance ... Xsuite should help customers address these issues."

About Xceedium

Xceedium, Inc. provides zero trust privileged identity management solutions that enhance security, improve productivity, and reduce the costs of meeting compliance regulations. Xceedium's Xsuite protects customers from the numerous threats that privileged users and trusted insiders pose to networks and data. Privileged users include employees and third parties — such as vendors, consultants and contractors — who manage critical IT assets from inside the network or from remote locations.

Global government organizations and leading enterprises — ranging from healthcare and financial services organizations to retail and telecommunications companies — rely on Xceedium for proactive insider threat protection, active monitoring, and detailed forensic data for high-risk engagements and administrative sessions.

For more information, please visit www.xceedium.com.