illumio

EBOOK

In Defense Of Active Directory

How to defend one of the most critical and targeted systems in modern environments

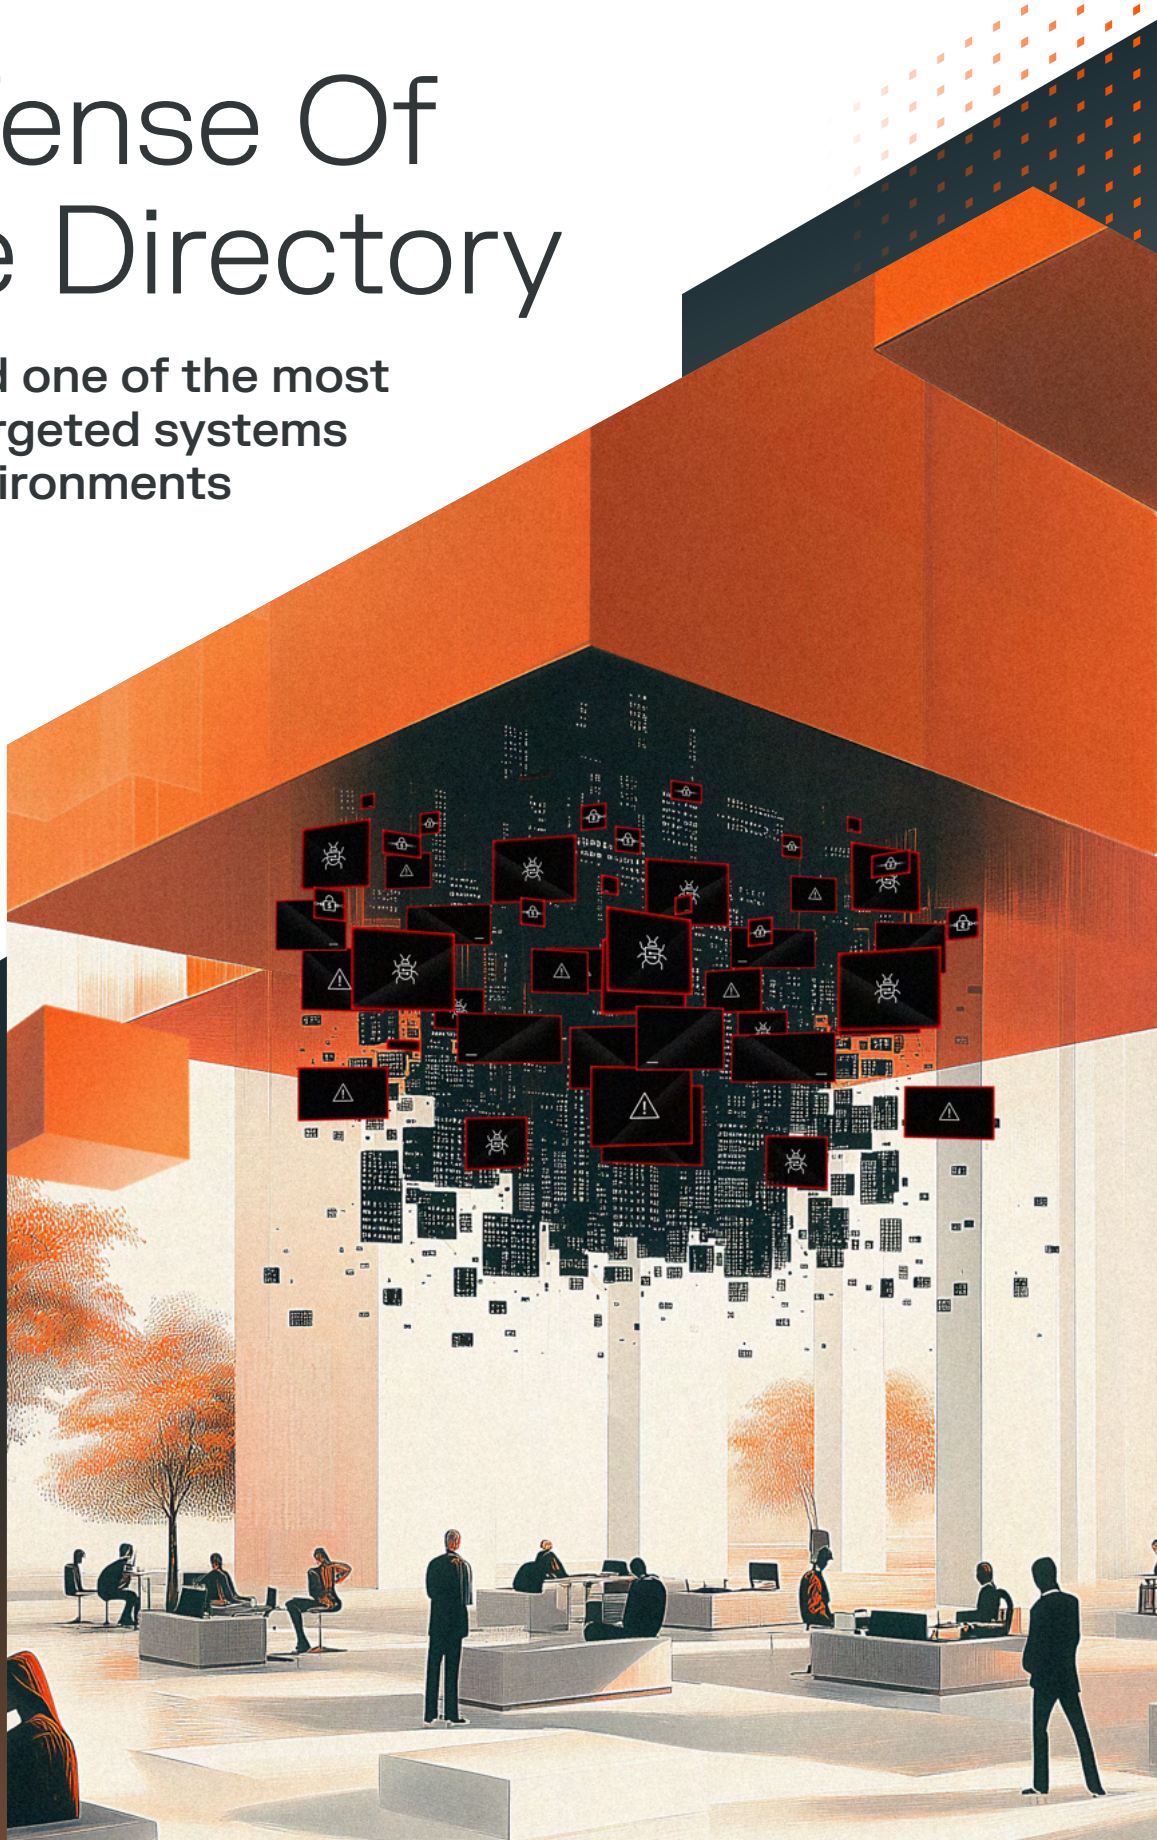

Contents

The security challenge at the center of modern environments

Active Directory still plays a central role in how identity works across most environments. It helps determine who can log in, what systems they can reach, and how trust flows across users, devices, and applications. When it works, nobody notices. But when it's compromised, critical systems and data can be exposed.

Attackers understand this better than many defenders do. They don't treat Active Directory as infrastructure. They treat it as the fastest path to total control. Compromise it, and they can move, escalate privileges, establish persistence, and operate undetected for months. A foothold can become full domain control.

This is why Active Directory remains one of the most targeted systems in modern environments — and why defending it requires a different approach. Active Directory may be legacy, but it's still foundational.

Legacy, but still foundational

Despite years of cloud adoption and identity modernization, Active Directory is still here. Many organizations still rely on Active Directory for on-premises identity and syncing users and groups to Microsoft Entra ID to extend access to the cloud.

That persistence is often described as "legacy debt." In practice, it's just how the business runs.

Active Directory continues to support large on-premises environments and critical systems that rely on authentication protocols such as Kerberos and NTLM. Those dependencies exist because enterprises have accumulated applications, services, and infrastructure over years — sometimes decades. And they all rely on a shared identity layer.

Unwinding that dependency takes years and carries real operational risk. As a result, even organizations deep into cloud adoption still run Active Directory as part of a hybrid identity model.

Attackers see the opportunity. A system that defenders would like to modernize away is still deeply embedded, broadly trusted, and tightly connected. That makes it an ideal target.

This is especially true in environments where stability matters more than speed.

Hybrid identity models

Many enterprises rely on legacy applications, strict data rules, and core systems that can't go down. Meanwhile, changes to identity spread fast. Even small mistakes can break the workflows the business relies on.

A familiar pattern emerges. Cloud identity services like Microsoft Entra ID handle SaaS, cloud apps, and remote users, while Active Directory continues to support many internal systems, applications, and authentication flows. For most companies, that's the long-term setup, not a stepping stone.

Hybrid identity extends reach and flexibility, but often carries forward existing trust relationships and access paths. The cloud expands. The perimeter shifts. Yet inside many environments, Active Directory continues to anchor how systems connect and how access flows.

This hybrid model can introduce risk. As organizations modernize, identity paths — relationships, permissions, and trust models — often persist unless actively redesigned. In many cases, Active Directory remains a primary authority for identity and trust. And anything that plays a central role in access control becomes a high-value target.[1]

Modern access layered onto longstanding trust: that's what attackers exploit.

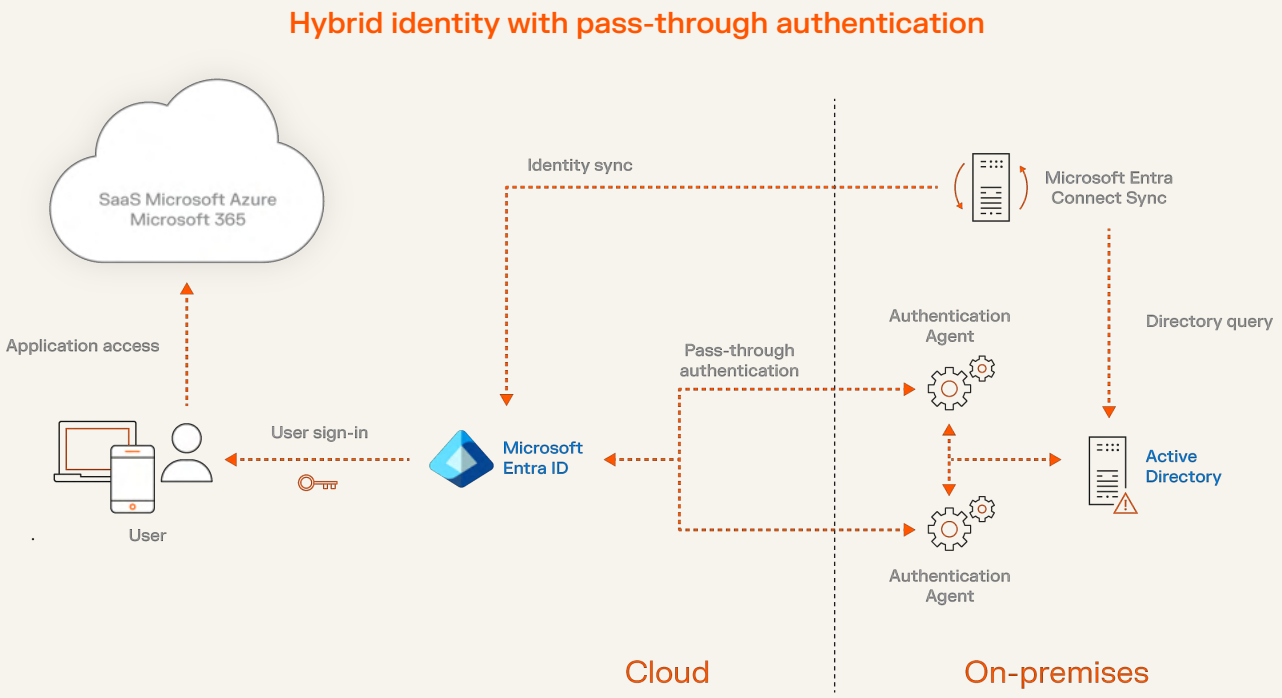

Figure 1: Hybrid identity architectures extend access into the cloud while continuing to rely on on-premises Active Directory for core authentication and system relationships.

[1] https://techcommunity.microsoft.com/t5/microsoft-security-blog/your-identity-infrastructure-is-your-security-boundary/ba-p/256344

Domain controllers: the path to Active Directory control

Active Directory and domain controllers are distinct, but inseparable in practice. Active Directory defines identity, permissions, and trust across the environment. Domain controllers host that directory and enforce it — processing authentication, validating access, and replicating identity data.

Attackers are after control of identity, not just a single system. Active Directory is the objective; domain controllers are how that control is achieved. Compromise a domain controller, and you control the directory.

Figure 2 shows why. Domain controllers sit at the center of authentication and policy enforcement. Users, endpoints, and applications depend on them to function. That central role is essential — and exactly what attackers exploit.

As attackers move through the environment, they follow existing trust paths. Those paths lead back to domain controllers, making them a natural focal point during lateral movement.[2]

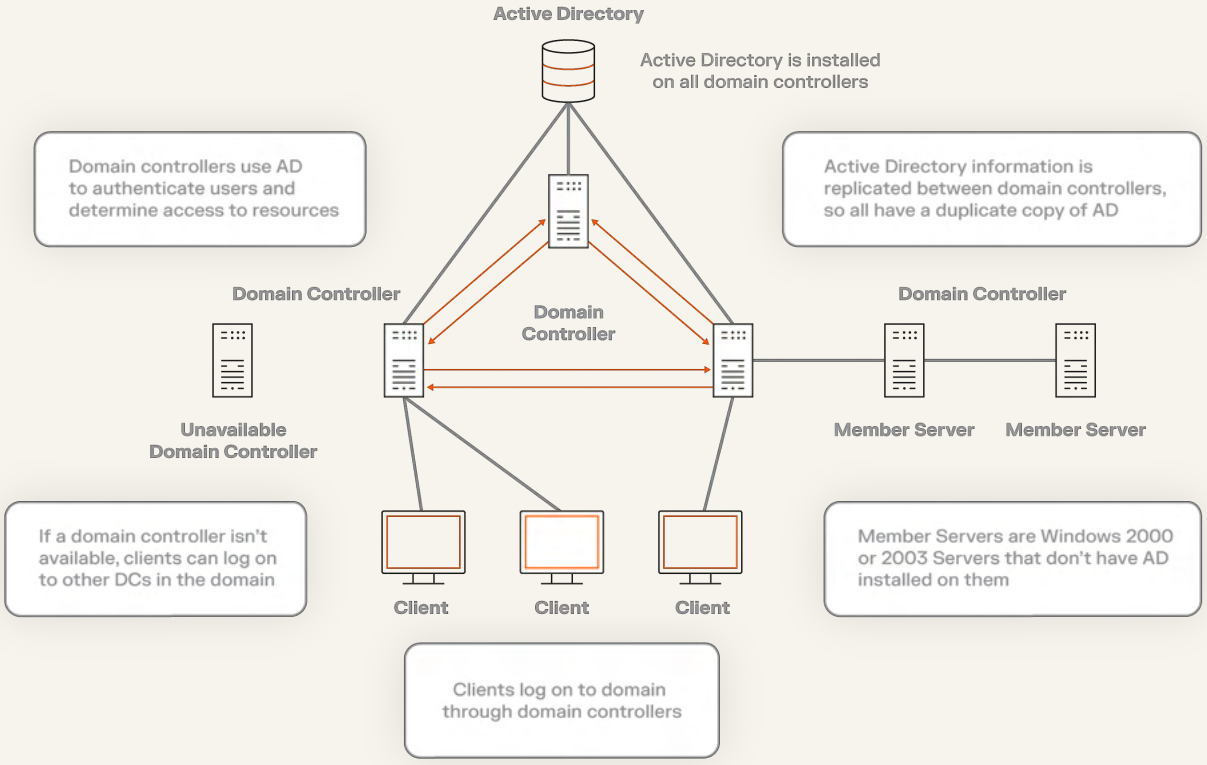

Figure 2: Active Directory domain controller connectivity model

[2] https://www.sciencedirect.com/topics/computer-science/active-directory-user?__cf_chl_tk=I4eXbVwnEaTgiZ8x2cEkNln0KiP7fBWM3ImpHQlryJo-1777299782-1.0.1.1-L.ktuh6leFYCR7zZjQw9N58hT3ZUfUb10tDGFLY1aIQ

Real-world breaches involving Active Directory

Year	Incident	What Happened	AD / Identity Impact	How It Happened
2017	WannaCry Ransomware	Wormable ransomware exploiting SMB.	Spread across domain-joined systems.	AD trust relationships accelerated the spread of ransomware.
2017	NotPetya Attack	Destructive malware disguised as ransomware.	Credential harvesting and lateral movement via domain tools.	Unprotected AD enabled the rapid, destructive spread of malware.
2020	SolarWinds Supply Chain Attack	Compromised Orion updates enabled espionage.	AD FS and SAML token abuse for persistence.	Identity compromise enabled stealthy access.
2020	Ryuk / Conti Ransomware	Human-operated ransomware campaigns.	Credential dumping and domain takeover before deployment.	Domain control preceded the ransomware impact.
2021	Microsoft Exchange ProxyLogon	Exploited Exchange vulnerabilities.	SYSTEM access and credential theft leading to AD compromise.	Entry points led to identity takeover.
2022–2023	BlackByte Ransomware	Targeted critical infrastructure.	Used AD for lateral movement and persistence.	Reinforces identity as attack backbone.
2023	LockBit Ransomware	Global ransomware operations.	Domain enumeration and domain controller compromise.	Identity was targeted before the ransomware deployment.
2023 (disclosed)	Volt Typhoon Activity	Nation-state persistence in infrastructure.	LOTL techniques leveraging AD trust.	Stealthy identity abuse led to widespread breach.

From breach to control: how attacks unfold

Most attacks start small: a phishing email, a reused password, a single compromised endpoint. At the initial breach, the attacker has limited access and no domain-level privileges.

From that foothold, attackers begin reconnaissance. They search the compromised system for credential material — password hashes, Kerberos tickets, and authentication tokens left in memory from active user sessions. Native tools or common frameworks can extract these artifacts from processes such as Local Security Authority Subsystem Service (LSASS). Once attackers have them, they can authenticate to other systems as that user without ever needing to crack a password. The environment already trusts the credentials. (See Figure 3.)

Techniques such as pass-the-hash and pass-the-ticket let attackers move laterally as legitimate users. Each step yields new systems and more credentials, bringing attackers closer to domain controllers and full control of the environment.

Techniques attackers used after initial compromise

Technique	Description
Pass-The-Hash / Ticket Attacks	Reuses stolen authentication data to gain access
Exploiting Weak Configurations	Abuses misconfigurations to gain elevated access
Living Off The Land (Lolbins)	Uses built-in tools to evade detection
Credential Dumping	Extracts credentials from memory or systems
Session Hijacking / Token Abuse	Steals or forges tokens to access systems
Network Sniffing	Captures network traffic to collect credentials

Figure 3: Attackers escalate from a single compromised user to domain admin by reusing stolen hashes and credentials — moving laterally without needing to crack passwords.

Living off the land: attacks that hide in plain sight

Once attackers have credentials, they don't need to bring too many of their own tools. The ones already installed — PowerShell, PsExec, WMI, SMB, RDP, WinRM, SSH — are enough to move between systems, query the environment, and work toward domain control.[3] This approach is known as living off the land (LOTL), and it's now the dominant pattern: a 2025 Bitdefender analysis of more than 700,000 incidents found 84% of major breaches involved LOTL techniques.[4]

What makes it effective is that it's invisible to most defenses. LOTL activity often looks like normal activity. It runs in memory, blends into normal operations, and leaves fewer obvious signs than malware-based attacks.

This dwell time often exceeds 200 days while attackers expand access one compromised system at a time — endpoint to server, server to infrastructure, and eventually to domain controllers.[5] The cycle is the same at every step: compromise, escalate, move, repeat.

In Active Directory environments, that cycle moves quickly. The environment is highly connected by design.

Compromise, escalate, move, repeat

Lateral movement happens in two stages — and an attack stalls if either stage fails. What makes it effective is that it's invisible to most defenses. LOTL activity looks like normal activity.

Inside the host. Before moving anywhere, attackers strengthen their position on the system they've already compromised. That means extracting credentials, escalating privileges, and running tools in memory to avoid detection. If they can't get administrative access or find reusable credentials, movement stops here.

Across hosts. Once attackers have credentials, they move between systems using protocols the environment already trusts: SMB, PsExec, RDP, WinRM, SSH. One compromised system becomes many. Attackers move from endpoint to server, from server to infrastructure, and eventually toward domain controllers.

Both stages repeat: compromise, escalate, move, repeat. Each cycle expands access and brings attackers closer to full control — using the environment itself, not any single exploit.

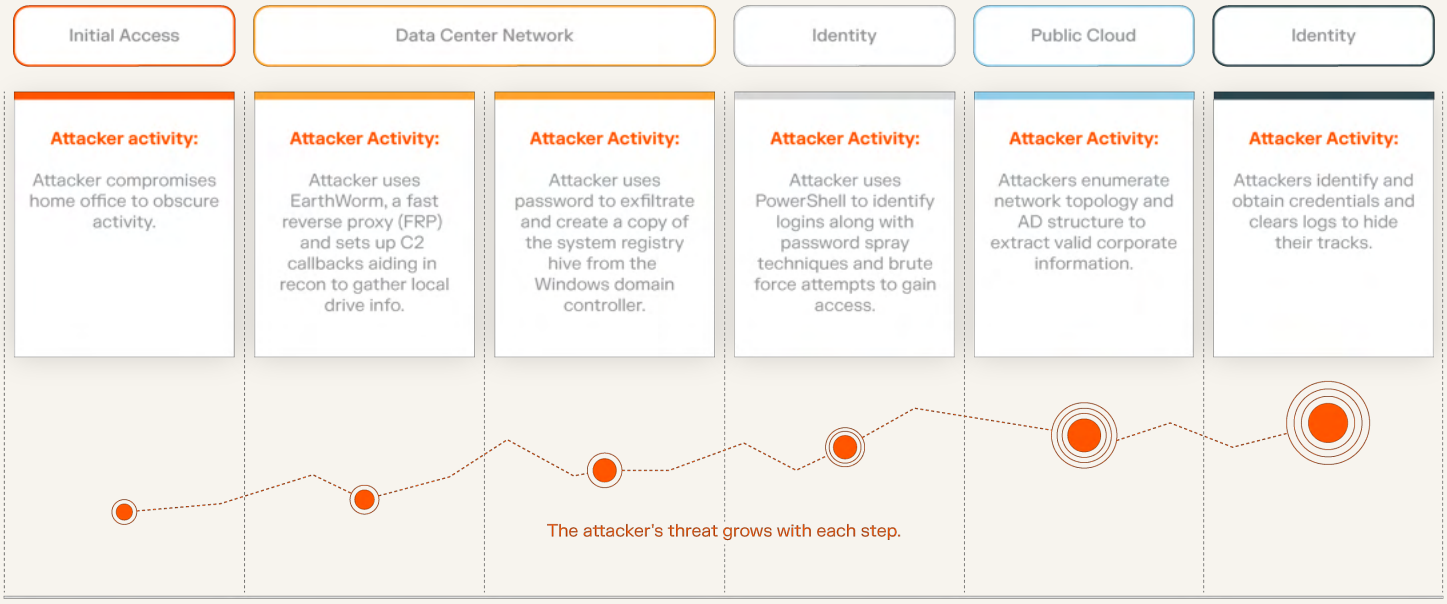

Figure 4: A simulated living off the land attack using EarthWorm, a powerful, open-source command-line tunneling tool.

[3] https://www.illumio.com/blog/modern-trojan-horse-how-attackers-live-off-the-land-and-how-to-stop-them
[4] Bitdefender. "700,000 Security Incidents Analyzed: Living-off-the-Land Tactics." June 2025.
[5] IBM. "Cost of a Data Breach Report." 2025.

Why threats matter for Active Directory

Compromising a single endpoint is a foothold. Owning Active Directory is the endgame.

Once inside, attackers aren't guessing where to go next. They use the environment itself to find the fastest path forward, and in most enterprises, that path leads directly to Active Directory. Using built-in tools, attackers can:

- Query AD to map the environment
- Identify privileged accounts and access paths
- Move closer to domain controllers
- Escalate access without deploying malware
- Reach network shares and file systems

Active Directory is what makes this possible. It can reveal domain structure, privileged relationships, and information that helps attackers identify and move toward domain controllers. Each step builds on the last, and all of them converge on the same target: the domain controller, where identity is enforced and control of the environment can be taken.

The attack isn't complex. It works because the environment was built for connectivity, not containment.

Attackers know your environment, even if you don't

Attackers don't stumble through Active Directory — they map it out.

With the right tools, they can analyze the entire environment and surface the relationships that matter: who has access to what, where privileges exist, and how to move between systems.

What emerges is something defenders rarely see: the shortest path to control. For attackers, this is clarity. For defenders, it's a blind spot. They see events. Attackers see relationships.

Many organizations don't have this level of visibility. They rely on logs, alerts, and point-in-time data. These are useful, but incomplete.

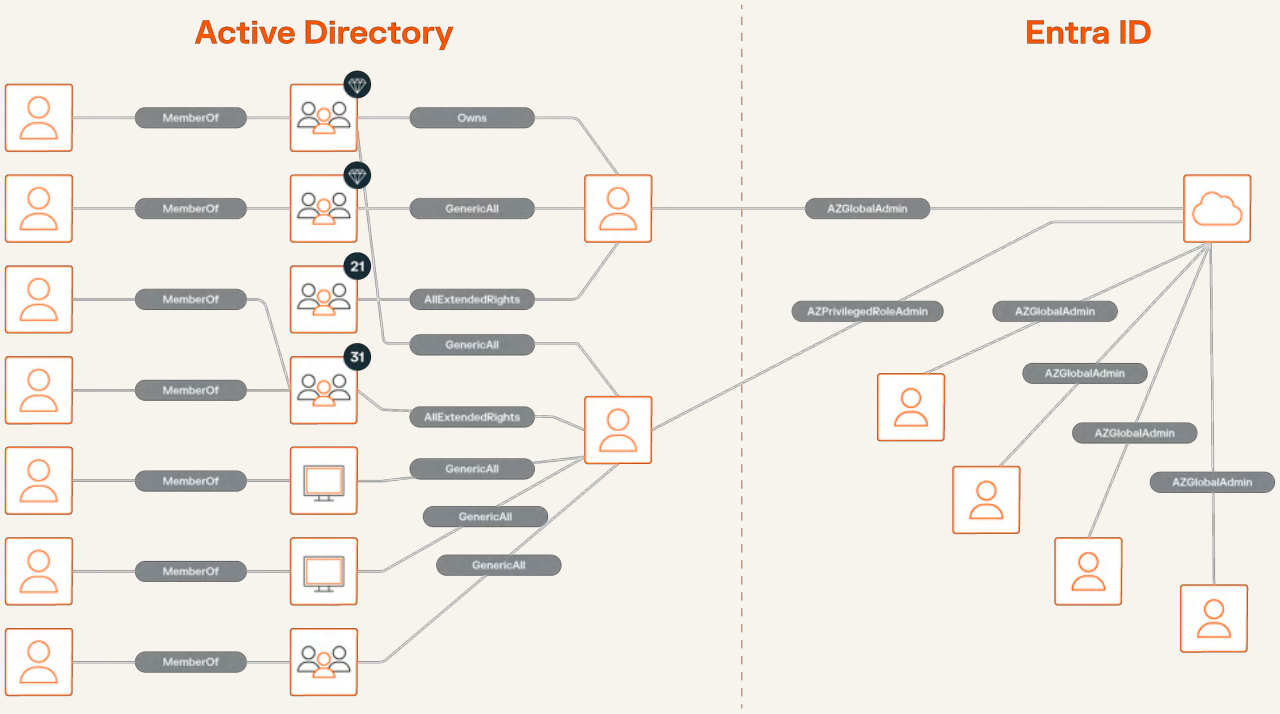

Figure 5: Attack graph tools like BloodHound map identity relationships across Active Directory, revealing hidden paths to privilege escalation and domain control.

Tools attackers use to map your environment

Attackers don't guess. They use purpose-built tools to understand identity, access, and relationships:

- **BloodHound.** Maps AD relationships and attack paths.
- **SharpHound.** Data collector used with BloodHound.
- **PingCastle.** Audits AD security posture and risk exposure.
- **PowerView.** Enumerates domains, users, and permissions via PowerShell.
- **ADFind.** Queries AD for objects and relationships.
- **CrackMapExec.** Automates credential validation, remote execution, and lateral movement discovery.
- **Mimikatz.** Extracts credentials (passwords, hashes, tickets) and enables privilege escalation.
- **Rubeus.** Abuses Kerberos tickets (pass-the-ticket, ticket requests).

Most security tools watch the wrong things

Once attackers are inside, the defender's problem changes. It's no longer *how does an attack happen?* — it's *can I see the one that's happening right now clearly enough to stop it?*

Most enterprise environments are dense: thousands of systems, identities, and permissions interacting in real time. Security teams monitor that activity with logs and alerts that capture events — a login, a file access, a process launch — but not the relationships between them.

That's the blind spot. Without a view of how systems communicate and how access flows, teams see every tree clearly but completely miss the forest.

Closing that gap means going beyond isolated signals. Four capabilities matter most:

- **Lateral movement visibility.** See how systems connect and which paths are actually in use.
- **Behavioral context.** Tell normal activity from activity that's creating risk.
- **Risk prioritization.** Know which connections and access paths would matter most if compromised.
- **Rapid containment.** Cut communication or isolate systems in minutes, not hours.

With this view, activity in the environment becomes legible. Without it, defenders are reading the event log of an attack they can't quite see.

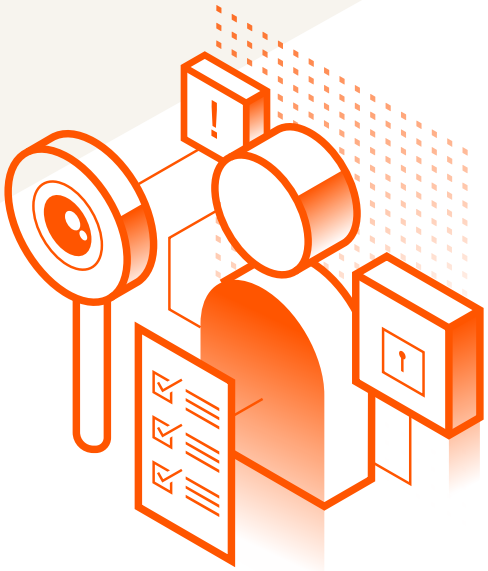

BloodHound: what attackers can see

Attackers often know your Active Directory better than you do. BloodHound is one of the reasons why.

It's an open-source attack-path mapping tool that ingests Active Directory data — users, groups, permissions, and systems — and reveals how an attacker could chain those relationships together to reach domain control.

The output is concrete: a graph that shows exactly how a low-privileged user can move, step by step, from their starting point to full domain control.

This is how attackers understand an environment — not as isolated systems, but as a connected graph of access and privilege.

Without that same view, it becomes difficult to identify:

- Stale or overprivileged accounts
- Hidden trust relationships
- Indirect paths to domain controllers

What looks secure on the surface often isn't. Because when one side knows the environment better than the other, lateral movement can become inevitable.

Firewalls weren't built for this

Even when organizations recognize the risk, they hit a second problem: they can't easily reduce it.

Active Directory runs on constant communication. Domain controllers use both static and dynamically assigned ports. In a large enterprise, dynamically assigned ports and constant system-to-system communication can scale to millions of possible communication paths.

These aren't front and back doors. They're the doors between every room in the house: every workload, every service, every identity relationship. And every one of them is a path an attacker can take once they're inside.

Traditional controls weren't built for this scale. Firewalls need static rules; Active Directory needs dynamic communication. Faced with that mismatch, most organizations choose to keep systems working, which means allowing broad internal access. The result is a functional environment with an enormous attack surface.

There's a second layer to the problem. IT teams don't design environments from scratch — they inherit them. Over years, sometimes decades, layers build up: new systems on top of old ones, modern services wired to legacy infrastructure. Attackers don't need to target the newest system. They can enter anywhere and move toward something older, less monitored, and still deeply connected to Active Directory.

The problem persists because the environment defenders are protecting was built for connectivity, not containment.

Rethinking defense: microsegmentation and controlling lateral movement

If preventing entry is no longer enough, the focus must shift — from access to movement. (See Figure 6.)

Modern security approaches reflect this shift:

- Assume breach
- Use map-based or graph-based visibility
- Monitor relationships between systems
- Contain threats quickly
- Limit privileges

This requires a Zero Trust-based approach that controls how systems communicate — not just visibility into what has already happened. (See Figure 7.)

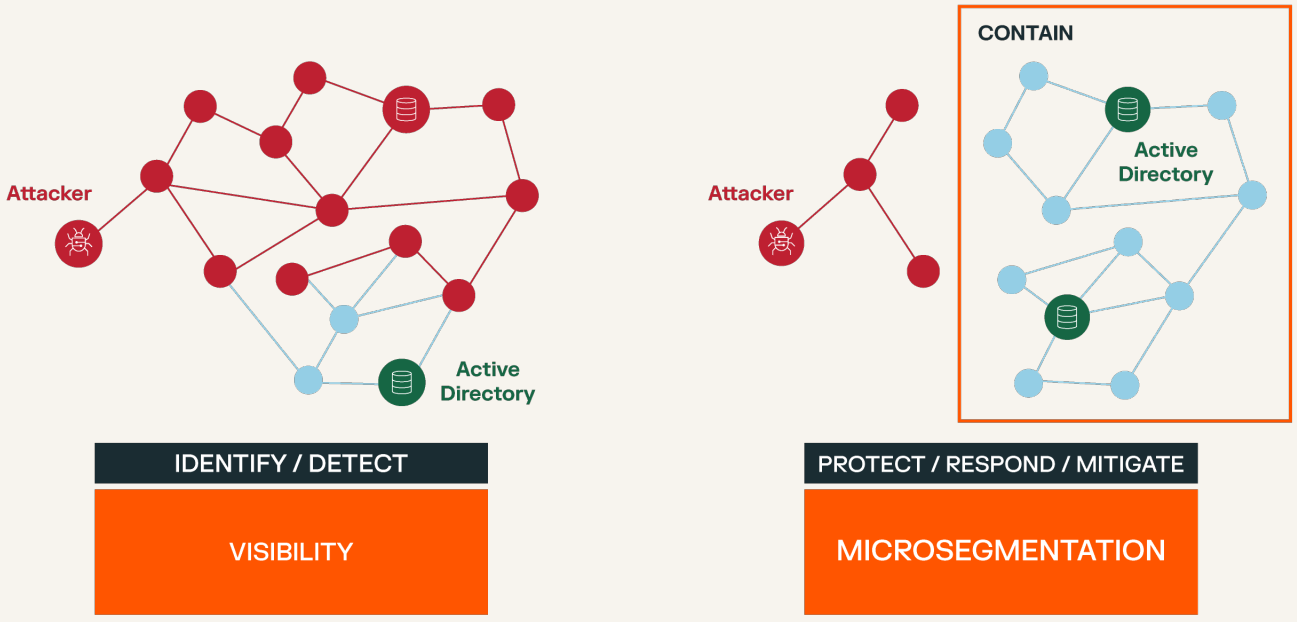

Figure 6: Visibility reveals attacker movement, but without microsegmentation, lateral movement remains unrestricted.

Figure 7: Designing for *containment* using Zero Trust-based microsegmentation

The goal is not to rip and replace the environment. It is to control how systems connect. When connections are controlled, movement is constrained. And when movement is constrained, attacks stop spreading.

When a breach is already underway, reactive tools have limited value — the damage is often spreading faster than alerts can be investigated. Continuous visibility allows security teams to detect lateral movement as it happens. The goal is to identify and contain threats in real time. This reduces the time an attacker has to move deeper into your network.

Microsegmentation: containing attacks before they reach Active Directory

The start of a breach rarely decides its outcome. What happens next does.

A stolen credential doesn't cause the damage. The unrestricted movement it enables does. When attackers can move freely, the environment becomes the weapon; when they can't move, the attack stalls. That's the difference between a minor incident and full domain control.

Limit the blast radius from the start

Not every system needs to talk to Active Directory. But in many environments, Active Directory is reachable through more paths than security teams can easily see or control. Microsegmentation changes that. It restricts which workloads can communicate with domain controllers — and from where.

Even with valid credentials, an attacker can't reach Active Directory from an unexpected system or path. Fewer pathways means less exposure and a smaller blast radius.

Watch how Active Directory is actually used

Most attacks don't announce themselves. They use real credentials and trusted protocols, which means signature-based tools miss them. What gives them away is behavior — and behavior is only meaningful if you know the baseline.

Understanding how Active Directory is *normally* used means mapping which systems query AD, which accounts authenticate where, and which protocols carry legitimate traffic between which workloads. That baseline is what makes everything else possible. Without it, anomalies don't look anomalous.

But with it, attackers' movements stand out — even when the credentials they're using are real.

Detect movement before it becomes control

There's a window between initial access and domain takeover. That window is where attacks can be stopped — if you can see them forming.

Against an established baseline, specific signatures mark lateral movement in progress:

- A system querying AD it's never queried before
- A sudden spike in authentication requests
- RDP or SMB traffic between workloads that have never communicated
- Access patterns that escalate toward domain controllers

These signatures show an attack still forming, not one already finished. Catching them early is the difference between containing a breach and cleaning up after one.

Contain immediately

Once a threat is confirmed, time matters more than investigation. The priority isn't shutting systems down — it's stopping the attacker from moving. Isolating a compromised workload at the network level cuts the attacker off from the rest of the environment instantly, without disrupting operations. The breach is contained at whatever point in the kill chain you catch it — before credentials spread, before domain controllers are queried, before Active Directory becomes the endgame.

The bottom line

Blocking every entry point isn't realistic. Attackers will get in — through a phishing email, a stolen password, an unpatched server. That much is no longer in question.

What is in question is what happens next. A breach that reaches Active Directory can end in full domain control. A breach that stalls at the first workload stays a breach. The difference is whether lateral movement is controlled or unrestricted.

Containing breaches before they reach Active Directory is an operational discipline, not a theoretical one. And it's what turns a potential disaster back into a manageable incident.

You can't stop every breach. But you can stop disasters.

Discover how Illumio stops Active Directory threats before they spread.

Try Illumio Segmentation today.