ALSID
ACADEMY

Episode 5

THE INSIDER THREAT
THE ENEMY WITHIN

“Another buzzword, hooray!” We get it. We often feel the same when confronted with a *new* form of threat. Unfortunately, this one is very real and sarcasm won’t help fight it. So let’s get a little more analytical.

It’s a fact: malicious insiders are a growing problem. According to Crowd Research Partners’ 2018 Insider Threat Report (worth checking out), 53% of respondents confirmed they suffered an insider attack in the previous 12 months and a whopping 90% of them feel vulnerable to this threat.

Root causes, you ask? Respondents ranked ‘excessive access privileges’ as the main enabler of this threat. Naturally, their most common defense tactics ought to be access right hygiene and lateral movement prevention, right?

Wrong. Perimeter and data-centric protections still reign supreme among their preferred solutions, with the welcome addition of user behavior analytics. But are these approaches really addressing the root cause? Or are they merely treating the symptoms? In this guide, we’ll characterize the nature of the insider threat and the personae which embody it. From there, we’ll assess the relevancy of various defense tactics, so you can fine-tune yours.

SPECIFICITIES OF THE INSIDER

Let's state the obvious. The danger of "insider threats" is the intimacy the threat actor shares with its victim. The first axiom of insider defense tactics is that our opponent has gained, unchallenged and undetected, the typical rights of your average employee on (some) data, systems, and networks.

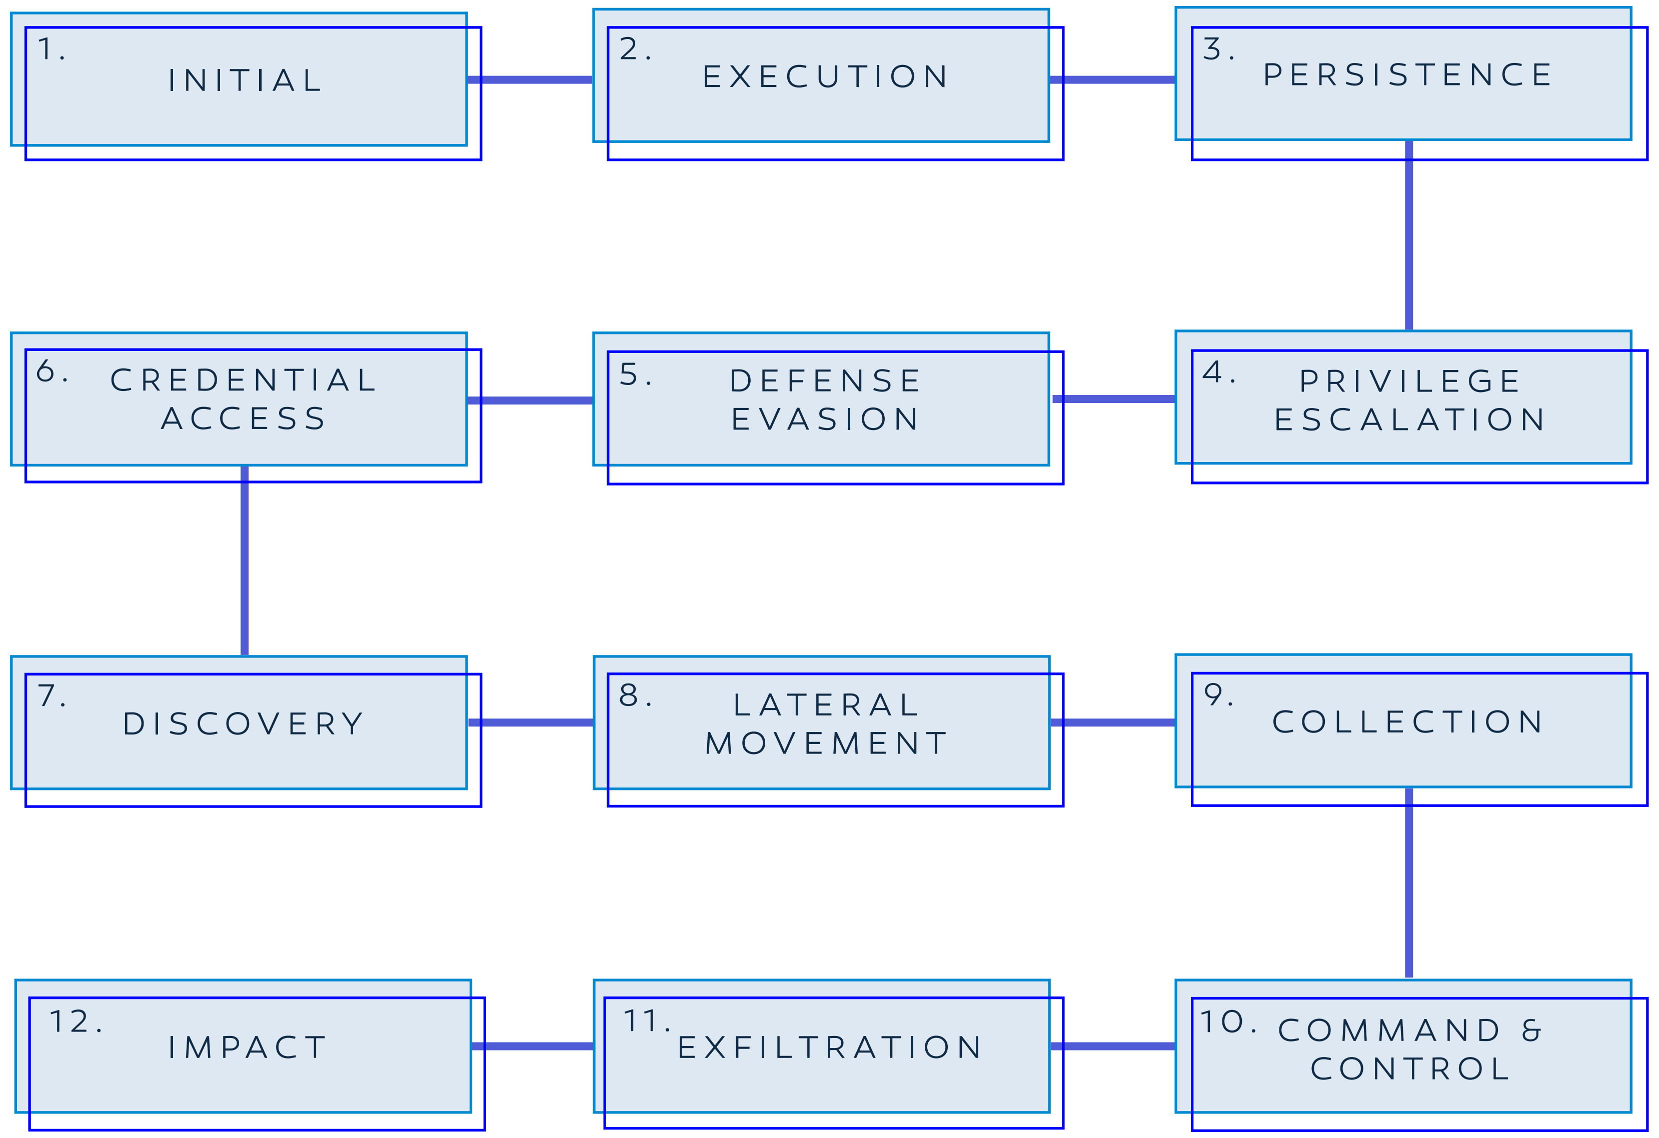

1. INITIAL

2. EXECUTION

3. PERSISTENCE

4. PRIVILEGE ESCALATION

5. DEFENSE EVASION

6. CREDENTIAL ACCESS

7. DISCOVERY

8. LATERAL MOVEMENT

9. COLLECTION

10. COMMAND & CONTROL

11. EXFILTRATION

12. IMPACT

Under the MITRE ATT&CK framework, our insider threat actor starts its attack at the equivalent of steps 6 or 7.

From there on, the journey of an insider advances according to its wickedness (from unintentional, to full-on sadistic). Everything depends on the attacker.

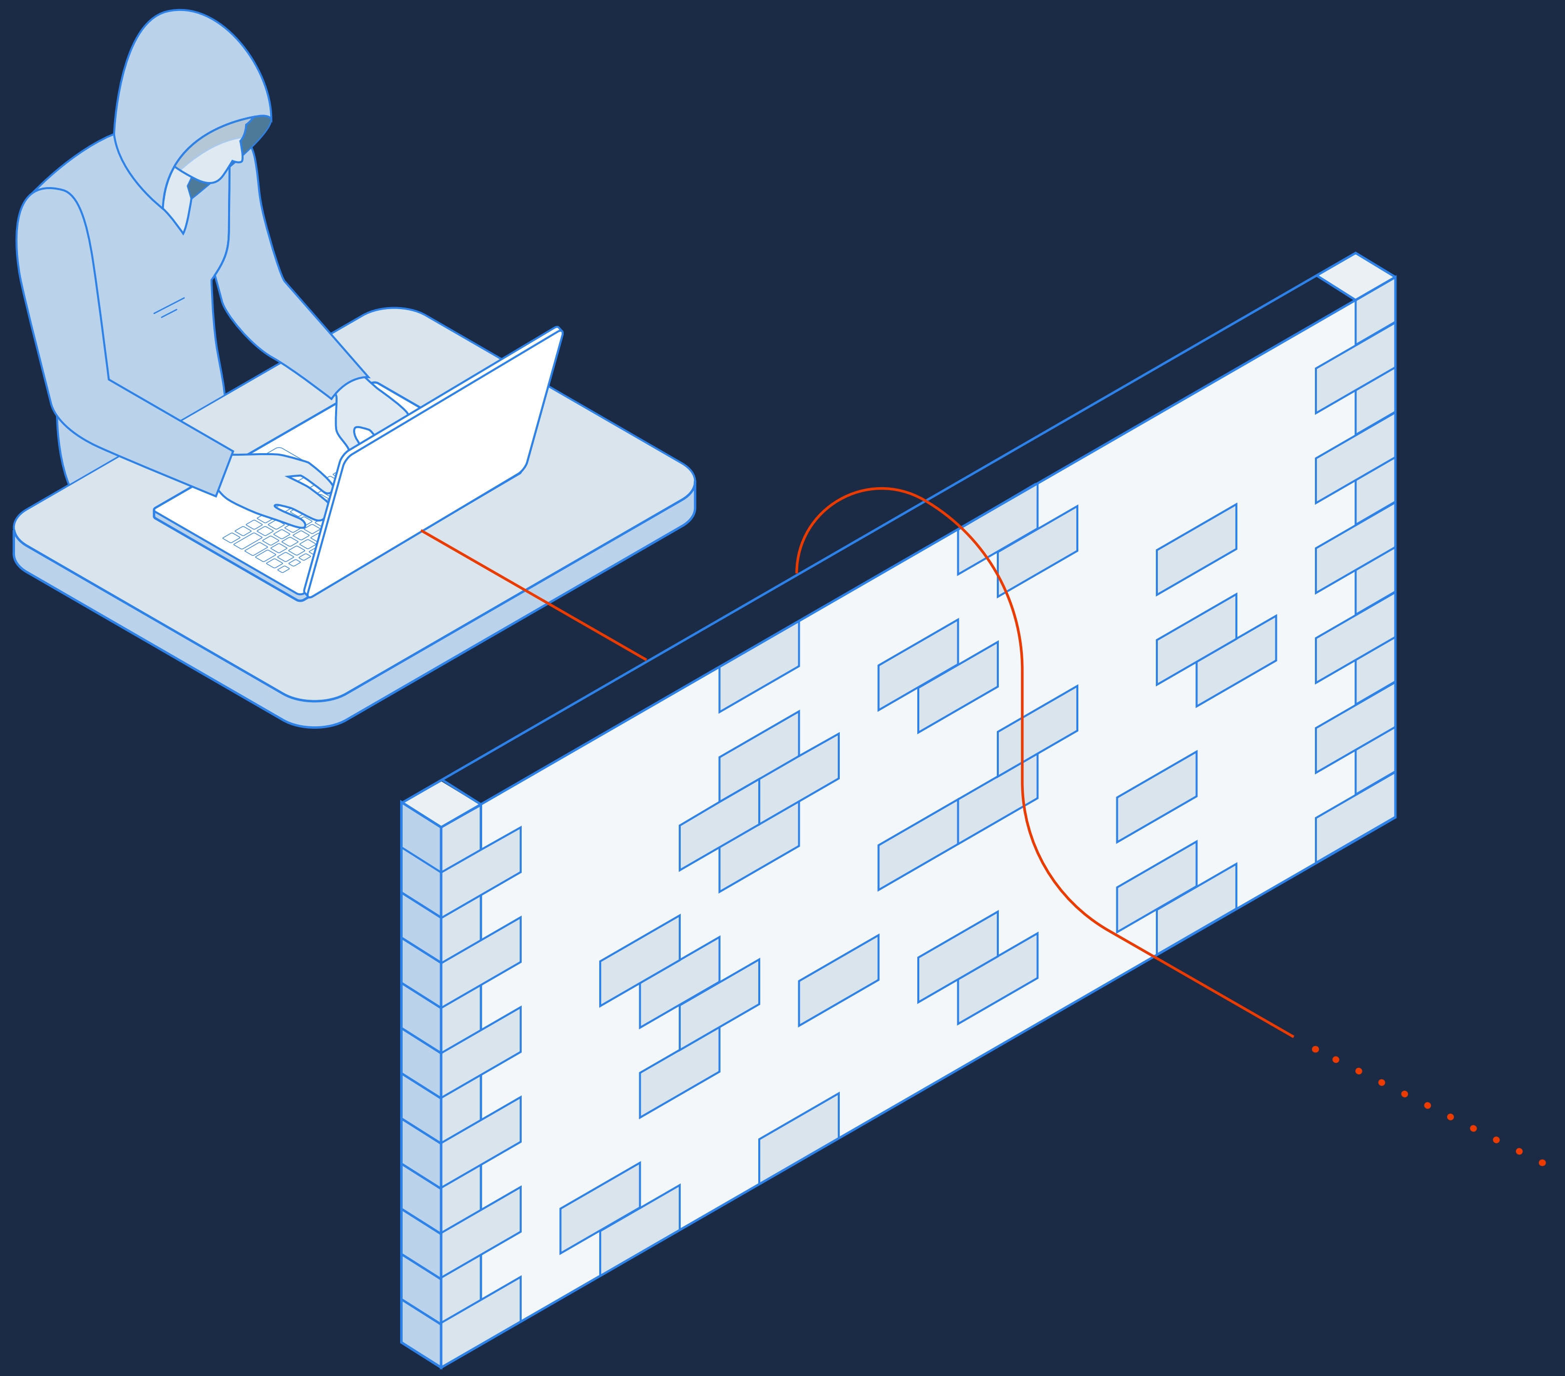

Conclusion A

There are two direct, very logical takeaways from applying this model to the insider threat:

1. Old-fashioned perimeter protections are hardly effective against insiders. They mostly make a difference at steps 1 - 2 and, to a much (much) lesser extent, at steps 3 - 5 and 11 - 12.
2. The first things any decent insider threat program should try to prevent are credential access, discovery, and lateral movement. Those are the first steps of the journey that goes from a “benign threat” to a “catastrophic security incident.” A defense tactic that tackles steps 6 - 8 will break the dynamics of the insider threat before any real harm can happen (starting at step 9).

THE INSIDER'S ECOSYSTEM

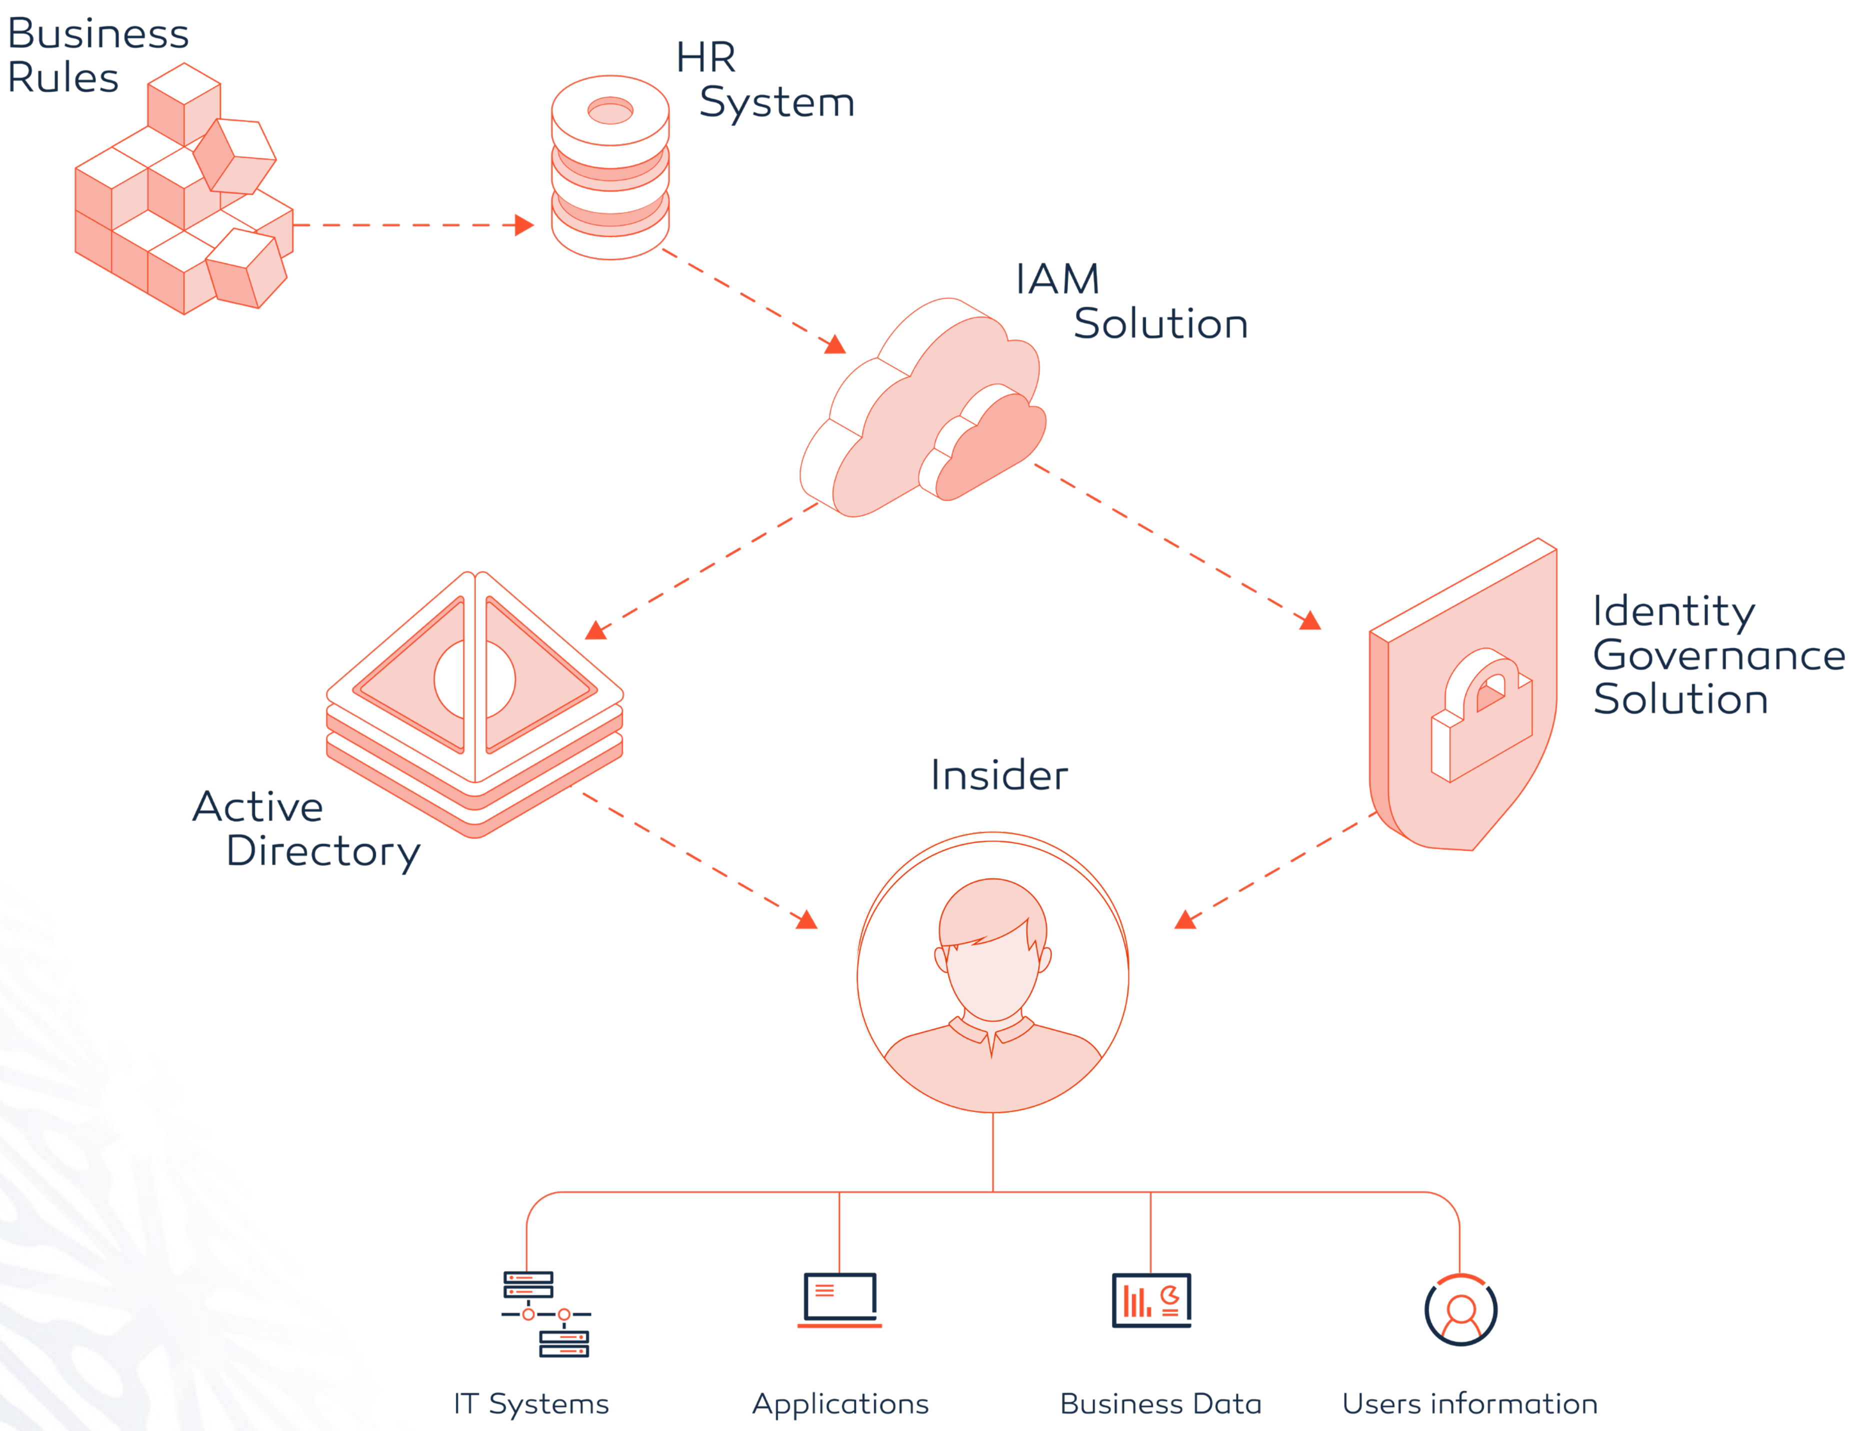

To quantify the actual rights an insider inherits from his/her particular position, we need to understand the ecosystem that governs those rights. In this respect, the typical (simplified) framework looks like:

Typically, an HR System registers the business rules that govern user types and access profiles. It's the foundation IT infrastructures use to implement ACLs and technical rights management.

An IAM Solution technically translates an organization's business rules to various directories and identity governance solutions.

Identity Governance Solutions implement identity-related processes, such as provisioning and privilege account governance.

A Directory Service *"is a shared information infrastructure for locating, managing, administering and organizing everyday items and network resources."* **(Wikipedia)** This is where the rules defined elsewhere are eventually enforced.

To sum it up, the rights an insider inherits are:

1. Defined in the HR System.
2. Translated by the IAM Solution into technical rules.
3. Eventually enforced by a Directory Service and an Identity Governance Solution (which might actually be a unified software).

That said, the chain of events is not an always-on process that's executed at machine-speed. They are human decisions, with validations and actions that make this a rather slow, somewhat asynchronous workflow.

A by-product of this is that it's inefficient to try to tamper with the HR System's business rules (step 1) to eventually corrupt the users' rights that are enforced at step 3. Such an attack's probability of success, considering there are several human approvals in this chain, would be significantly below average, and the fruits of success would prove slow to grow.

Conclusion B

The corollaries of this user management ecosystem are that:

1. The hazard level of an insider who does not intend to change his/her access rights is defined by the combination of those 4 different systems. Therefore, all insider threat programs must start beyond IT with the HR System's business rules' hygiene, and only then ensure those rules are accurately translated and enforced at a technical level.

2. For an insider who intends to change his/her rights, the most direct option consists of corrupting the Directory Service. Going upward in the process (HR System, IAM Solution) to achieve the same result is not unseen but proves way more complicated. Any decent insider threat program must include a serious defense tactic for Directory

INSIDERS & DEFENSE TACTICS

And what should we do about it.

Johnny Clumsy – The unknowing, accidental attacker

Behavior and Motives

Johnny doesn't want to do anything harmful, but he is prone to blunders. A typical consequence of Johnny's clumsiness would be a data leak either through bad internet usage (personal cloud storage, personal emails, unmanaged online tools, social media, etc.) or through lost devices (USB sticks, laptop, etc.).

Defense Tactics

Johnny is dangerous because his employer was lenient in defining business rules and translating them into technical enforcements. Therefore, the keys to limiting risks are:

- By virtue of our 'Conclusion B-1,' the most important security tactic here consists of maintaining a good hygiene throughout the whole chain, starting with the HR System. Notably, a least-privilege approach must be defined, and then enforced with compliance tools.
- And because Johnny's small privileges might still be sufficient to do harm, organizations minimize the risk of involuntary data leakage through perimeter protections (DLP, EPP, firewalls, etc.) and encryption.

Ansel Evil – The dishonorable employee

Behavior and Motives

Ansel is not used by a third-party: he, himself, wishes to harm his company. Ansel will intentionally exploit his rights over resources to steal secrets and/or destroy company assets.

In some rarer occasions, an IT-competent Ansel might go on a hacking spree to access data that are normally out of his reach.

Defense Tactics

In general, the hygiene and compliance tactics that applied to Johnny are still valid for Ansel. It should be noted, however, that against an intentional attack, perimeter protections are mostly ineffective and will only, at best, prove useful for post-incident investigations.

In the case Ansel happens to be a wannabe-hacker, our 'Conclusion A-2' shows that organizations need first to prevent lateral movement, credential access, and privilege escalation. And because of 'Conclusion B-2', this must happen at the Directory-level. Concretely, it's about the hardening and real time monitoring of this vital in frastructure. It's also worth mentioning that compliance approaches are mostly ineffective against technically capable opponents: designing an attack that chains together compliant events is not rocket science.

Janet Puppet

Behavior and Motives

Janet is a legit user with no bad intentions, but whose account has been hacked by a third-party. In that context, ‘her motives’ refers to the wide range of motives that usually animate hackers: state-sponsored espionage or sabotage, financial gains, ideological sabotage, etc. Organizations should expect Janet to be technically capable of exploiting IT weaknesses and vulnerabilities to gain access to resources that don’t belong to her.

Defense Tactics

Janet is equivalent to the worst Ansels of the world. For that reason, defense tactics are equivalent to those defined in the previous section, albeit possibly less compromising given the aftermath of Janet’s operations are likely to be way worse than Ansel’s. To summarize: hygiene is of the utmost importance, compliance and perimeter security tools are mostly ineffective, and Directory level hardening and monitoring is critical.

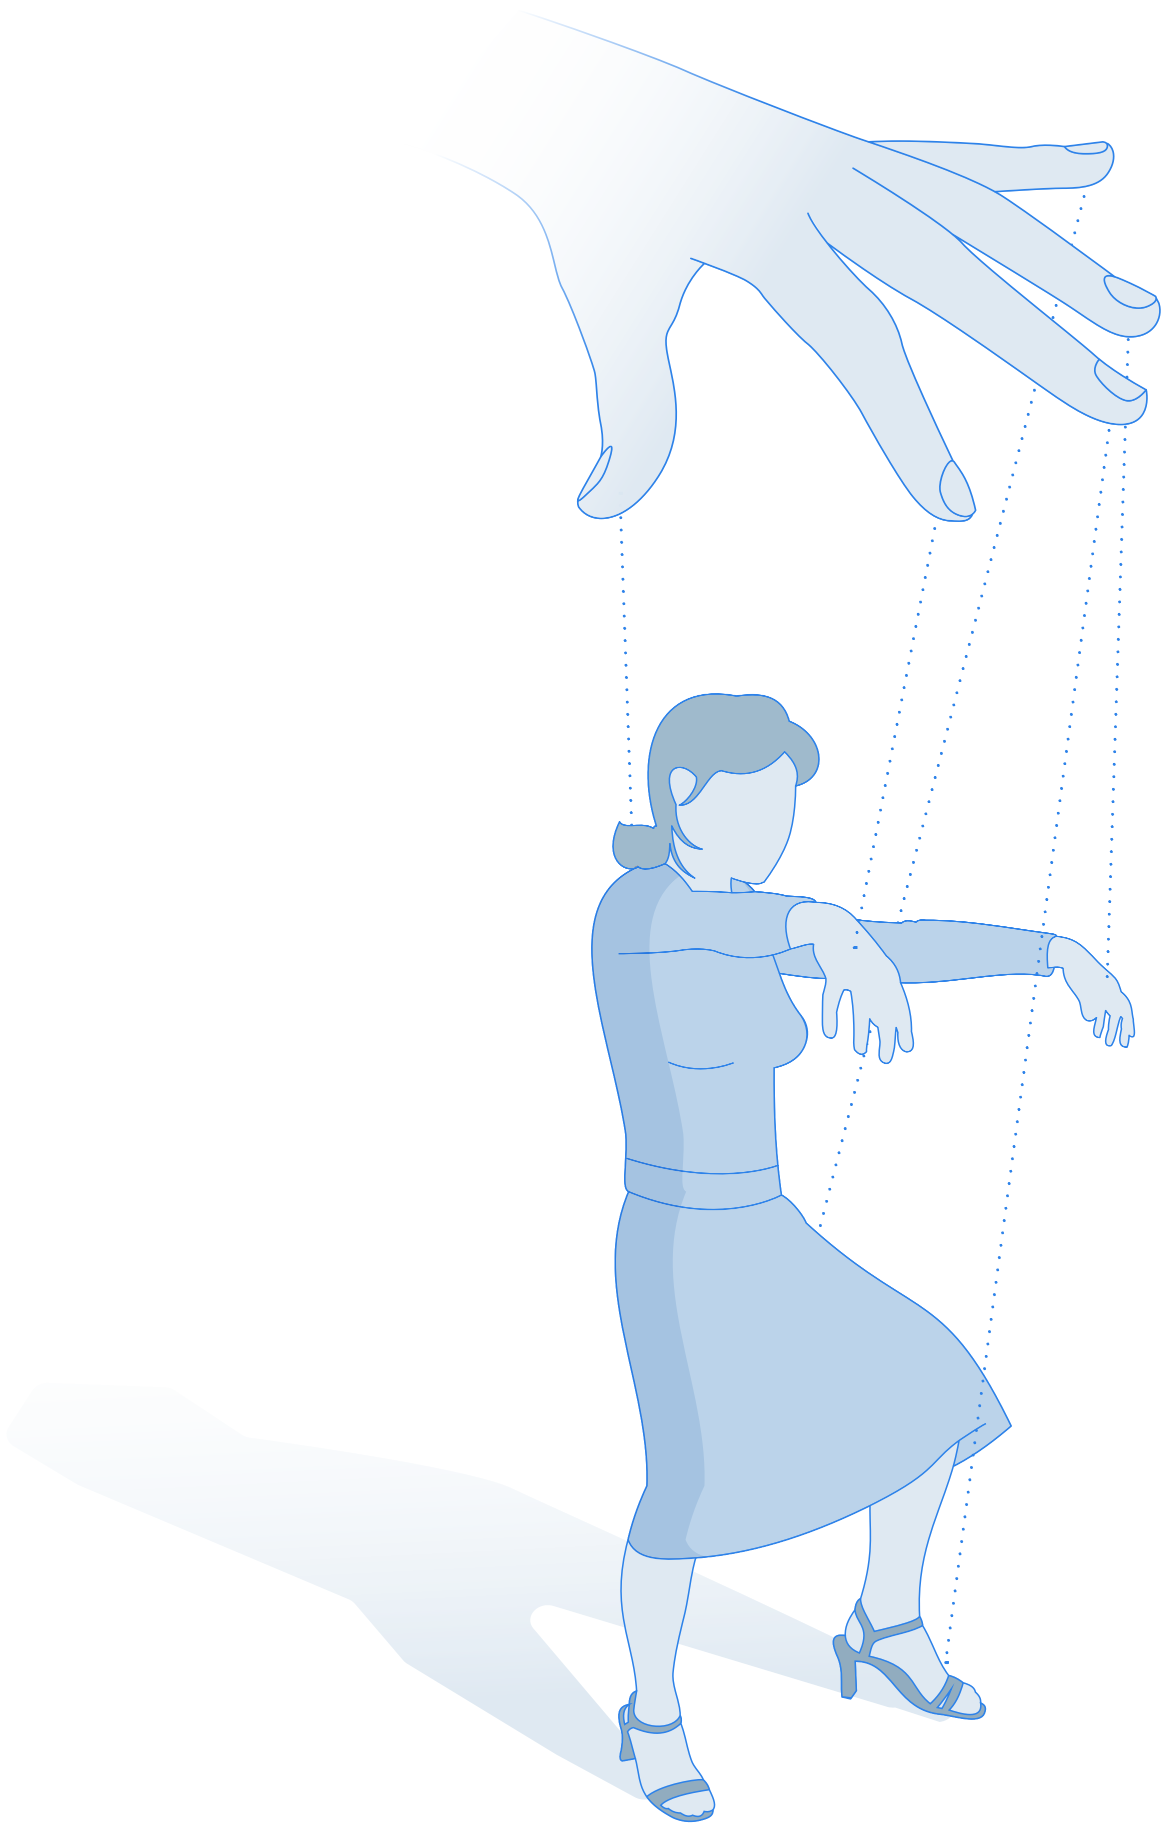

Esther Partner

Behavior and Motives

Esther can be an equivalent to Johnny, Ansel, or Janet. The only difference is that she's not an employee but a subcontractor. This, theoretically, gives her limited rights over her client's technical infrastructure.

That being said, it is hard to be generally conclusive on typical real-life subcontractors' rights. More often than not, we've seen suppliers enjoy greater privileges than their salaried peers. And in many cases, the management of these peculiar resources is... exotic: accounts are often provisioned at the IAM or Directory levels directly, ergo bypassing HR's business rules and basic hygiene. All in all, the threat level associated with Esther can be equal to that of Johnny, Ansel, or Esther.

Defense Tactics

The defense tactics effective against Esther are the fusion of all those previously mentioned. In her particular case, organizations may also want to consider the following:

- Contractually enforce regular security assessments of their suppliers' security posture.
- Contractually specify each party's responsibilities and liabilities in case of a breach.
- Reinforce Directory Services defenses if contractors' accounts are out of the HR System's scope.

SUMMARY

Take a look at the persona breakdown of each security practice's relevance.

		JOHNNY CLUMSY	ANSEL EVIL	JANET PUPPET	ESTHER PARTNER
HYGIENE & GOVERNANCE	STRICT RULE DEFINITION & TRANSLATION BETWEEN HR <-> IAM <-> DIRECTORY	+++	+++	+++	+++
	LEAST PRIVILEGE POLICY	+++	+++	+++	+++
	COMPLIANCE CHECKS	+++	+	+++	+++
	DIRECTORY PROTECTION	-	+++	-	+++
TECHNOLOGY DRIVEN	PERIMETER PROTECTION	++	+	++	++
	DATA ENCRYPTION	++	-	++	++
LEGAL & COMPLIANCE	SUPPLY CHAIN AUDIT	-	-	-	+++

Overall, it seems obvious that hygiene and governance best practices are the two most impactful initiatives when it comes to mitigating the insider threat.

From a technology perspective, hardening and monitoring Directory Services are the most effective tactics across the whole landscape.

DEFENDING DIRECTORY SERVICES

As we saw earlier, malicious insiders will try to move laterally (after discovery and privilege escalation, if necessary) through their Directory Services. In 99% of cases, this will mean exploiting Active Directory's vulnerabilities or misconfigurations. Unfortunately, Active Directory (and Directories in general) tends to become quite ugly quite quickly, so that there is no shortage of vulnerabilities and misconfigurations with which to contend.

On the defense side, defending Active Directory involves a combination of anticipating threats, detecting attacks, and responding to breaches.

Anticipating threats

Anticipating essentially means:

1. Enforcing Microsoft's and trusted advisors' (Gov-CERT, etc.) best practices.
2. Monitoring Active Directory to uncover regressions.
3. Going back to step 1 to fix regressions.

...Continuously.

Most large organizations go through step 1 once in a while, usually after a penetration test or a breach. Unfortunately, Active Directory implementations evolve quickly, and regressions are (re-)created on a daily basis, therefore creating the need for a continuous process of assessment/ fixing.

This admission rules out manual audits and pen testing, which are too slow and expensive to be run recursively at a high cadence. Organizations must evaluate dedicated technology that is AD-native.

Detecting attacks

Unfortunately, there's often confusion, when it comes to Active Directory security, between detecting attacks and detecting compliance regressions. The latter is necessary as a mean to anticipate threats but does little to address ongoing attacks: most attack patterns are chains of compliant events.

Conceptually, detecting Active Directory attacks is not exotic: it's about gaining intelligence on the tactics, techniques, and procedures used in the wild by adversaries and checking Active Directory events and objects that match those indicators.

Technically, though, this is very hard to implement with traditional monitoring technologies such as SIEM for several reasons:

- Active Directory logs are hard to collect and analyze in real-time, therefore creating a window for hackers to move unchallenged.
- Correlation rules are way harder to define and maintain for AD than for the rest of the threat landscape. They're not static IOCs: every new event triggers a graph calculation to uncover whether a new 'path to higher privileges' was created.
- Some of the most effective attack techniques, like DCShadow, don't leave a single log in the system, rendering log-based correlation useless.

As for anticipation, Directory-centric detection requires dedicated, Directory-native technologies that can tackle those unique characteristics.

Responding to threat

There's no doubt that some attacks will succeed and that Active Directory-centric security solutions must integrate seamlessly with the incident response practice.

That means providing analysts with the ability to drill down into the Directory's events (not logs) with a complete visibility on the events' details (which are truncated in logs).

For the most trivial incidents, the ability to trigger automatic remediations through SOAR integrations is a vital addition to alleviate the burden SOC and IR analysts bear on a daily basis.

Conclusion

We’ve seen that the insider threat is multiform. Fortunately, there are a few common denominators to the various actors behind these terms. Most notably, organizations must ensure that their HR governance is enforced technically throughout the entire user lifecycle, and that Directory Services, being the most obvious targets in actual attacks, get dedicated protections for anticipating threats, detecting attacks, and responding to breaches.

ALSID