451 Research

S&P Global
Market Intelligence

Identity and access security: Strengthening the resilience of cybersecurity's front lines

Analysts - Scott Crawford, Garrett Bekker, Fernando Montenegro, Daniel Kennedy

Publication date: Monday, October 18 2021

Introduction

As security industry analysts, we typically don't focus on particular attack types or threat actors as much as we study the offerings to confront these challenges brought to market by technology product and services providers. But cybersecurity is different from many other technology markets in that its directions aren't set by innovators working toward goals defined by themselves. It is a field much more like gamesmanship or military strategy, where intelligent adversaries and defenders contend against each other. In this arena, incidents and events can often change the very nature of the field. Yesterday's successful tactic can become commodity tomorrow, while newly exploited or revealed gaps may appear to change priorities for investment. And the market of products and services must respond.

In the last couple of years, we have seen attacks such as ransomware have such an impact. Indeed, the prevalence of such incidents – and the underground economy that has aligned to support them – seem to speak directly to the extent of the attacker's opportunity. Are there consistent themes in defensive gaps that make such a flourishing possible?

In some cases, the answer must be yes – and one of the most apparent such themes is the opportunity presented to attackers by gaps in identity and access management. Once an initial foothold into a target is gained, the ability of attackers to discover access privileges and relationships has enabled them to identify where and how those privileges can be acquired and exploited to do significant damage – including the ability to encrypt or otherwise compromise business-critical resources and effectively hold them hostage.

The Take

Laxity in defining and enforcing better constraints on such exposure has direct parallels with an earlier generation of attacks that also spread rapidly with serious impact. In this respect, poorly disciplined access control that allows attackers to find and exploit privileges is to today's ransomware what flat, unsegmented networks were to the 'worm' attacks of a prior era. Both gaps have allowed attackers to move with relative freedom across a landscape of targets, potentially compromising large numbers of other resources with rapid – and sometimes crippling – efficiency. Network segmentation has since become a well-established best practice in IT security.

Today, identity and access security – what we're calling IAS – is only just now becoming recognized as a similar priority, for effectively the same reasons. While they may be considered part of the broader spectrum of 'zero trust' initiatives intended to define and provision legitimate but well-constrained access, IAS techniques seek to identify the opportunities for compromise as an attacker would see them and mitigate those exposures. They are the complement of both identity and access management and zero trust, but with a primary focus on the adversary. As such, IAS tools may also be considered adjacent to the family of technologies associated with what we have described as 'outside in' security visibility (previously covered here).

Context

Here we'll explore how the phenomenon of systematic identity and access exploit became a hallmark of attacks such as ransomware. We'll summarize some of the approaches of IAS emerging to address common and all-too-frequently exploited gaps – including a direct parallel to network segmentation for access privileges – with a view toward where the field of IAS may lead.

Identity, access, and the paradox of 'security duality'

Identity and access management (IAM) is a fundamental of information security. It is the first line of defense for identifying and authenticating entities recognized in an environment (people, IT systems and resources, and groupings of each) from those that are not, and authorizing specific entitlements to the access and use of resources associated with identities.

In its simplest form, IAM may be implemented as a list of identifiers. The association of each with a set of privileges and authorizations facilitates the creation of an 'account' for each such identity. Once the number of entities reaches any kind of scale, however, managing things like access permissions and keeping authentication safe from abuse becomes unwieldy. Making this complexity manageable has, of course, long been a primary objective of IAM. The adaptation to specific environments has more recently given rise to technology segments such as cloud identity and entitlement management (CIEM) and SaaS security posture management (SSPM).

One way to better organize identities is through the use of a hierarchical naming system, such as the 'directory' concept. Analogous to namespaces in the Domain Name System (DNS), a directory groups a set of entities and associated attributes (such as access privileges) within an organization. The organization can be subdivided to smaller groups to facilitate more detailed management at more local levels. Privileges can be limited to specific entities and groups. They may apply to subgroups within larger groups. Groups can also establish so-called 'trust relationships' with other groups not always directly related within a hierarchy, enabling different organizations to function together as needed. One of the best-known examples of a directory system is Microsoft Active Directory (AD), long used by virtually any organization that must manage the Microsoft estate at any significant scale. Microsoft Azure Active Directory extends this functionality to cloud and on-premises resources as an Identity as a Service (IDaaS) offering, with some distinct differences from legacy AD.

One of the advantages of a directory system is that it also makes resources easier to find. This facilitates IT management when, for example, a directory group in a specific location can readily find resources local to that group such as printers, file shares, or even other computers or people within the group. When the need extends beyond a group but within an organization, a directory system facilitates broader access across subgroups or other so-called 'organizational units,' as well as the movement of personnel and resources within the larger organization as needed.

But with this convenience comes a risk: If privileges are not specific enough or are too broad, they may enable unintended access and capabilities. An individual or group, for example, may not have been meant to have the ability to modify sensitive resources limited to resource owners. But if an entity or group is also a member of other groups having those privileges, unintended consequences can result. In addition, long-standing accounts can accumulate privileges over time, which may never be curtailed or even kept current, if no apparent reason exists to do so (a phenomenon sometimes called 'privilege creep'). Nor is administrative privilege over a given asset always limited to administrative accounts. It can be assigned to individual user accounts as well as to groups. This is common with personal devices, if the individual also administers the device and must be able to add or modify software or otherwise control configuration, network access and so on as needed.

Highly sensitive privileges can also be assigned to accounts that serve enabling functions within the environment, such as so-called 'service accounts.' These are identities assigned to IT resources (as opposed to people) that enable them to perform certain sensitive tasks such as the automated interactions between systems. Regardless of whether the entity is human or not, if sensitive privileges are extended within an organization beyond those intended, control over sensitive IT assets may not be as limited as administrators assume. (CIEM initiatives extend this concept to the scale and profusion of components in 'cloud native' environments.)

Administrators are aware of these potential exposures and seek to contain them. This is, after all, a fundamental objective not only of IAM as a whole, but of specific segments such as privileged access management (PAM) and identity governance and administration (IGA) as well as CIEM, all of which seek to apply more fine-grained control over access. It is also an objective of zero trust initiatives that seek to contain access and privileges to the minimum required for specific functions. Administrators must be equipped with the tools for defining and administering IT and authorized access.

But this introduces a duality that characterizes cybersecurity more than any other technology field: attackers are aware of these gaps as well, and often seek to exploit not only those gaps, but also the same tools administrators use to manage and secure the environment. Evidence of the latter is abundant in so-called 'living off the land' (LOTL) attacks that harness the advantages of administrative tools such as Microsoft Windows' PowerShell. On its own, PowerShell is recognized as a highly useful management enabler. In the hands of an attacker that has secured administrative privilege, however, it can be a potent weapon – and won't be identified as malware. Behavior in its abuse must therefore be recognized by defenders.

The exploit of graph analytics is another example of how attackers leverage legitimate tools against a target. There are few more powerful ways that organizations can visualize identity, access and privilege relationships in an environment that expose risk – but the same is true for the adversary as well. In 2015, Microsoft's John Lambert famously summarized the value of graph analysis for communicating this insight in a way that has come to describe a definitive tactic exploited by adversaries: "Defenders think in lists. Attackers think in graphs. As long as this is true, attackers win."

Exploiting IAM against the victim

Once an initial penetration of the target environment has been gained – whether through a phishing attack; technical exploit of functionality enabling the use of an associated identity, account or

privileges; compromised or stolen credentials; or other means – attackers often move next to a reconnaissance of the environment to see what further targets may be available for exploit. Again, an unfortunate consequence of 'security duality' is that the same directory functionality that makes it easy for legitimate users and administrators to map out users and resources can also make it similarly easy for the attacker. If an attacker can explore a directory system to see what identities and groups have access to specific privileges, leveraging those connections to gain access to a valuable target with minimal risk of discovery is the next step.

This reality has played a key role in the prevalence and severity of attacks such as ransomware. Although open source tools such as BloodHound, which leverages the Neo4j graph database, can be used by defenders to visualize entities, relationships and gaps in IAS visible within an AD environment, they can be exploited by attackers as well. The combination of Microsoft Active Directory's wide adoption with the adversarial visualization of gaps in IAS offered by graph analytics has been a potent combination contributing to the impact of attacks. The approach has become such a common practice with attackers that it has been codified in the popular MITRE ATT&CK framework that evolved as a lingua franca for security operations.

Because enabling the business is a priority for IT, a high degree of latitude in access privileges benefits users who can freely access and use resources widely. This latitude may not be constrained if there's no compelling business reason to do so. But if an attacker discovers it, they can exploit it as well, taking over accounts that, either on their own or as a member of a group – or of a group that is itself related to another group – has access and control privileges that can be exploited in an attack. If such activity is not assumed to be abnormal, it may not be detected until too late.

Identity and access security on the rise

It is just this prevalence and severity of such attacks that is spurring a trend within the information security market: the increased visibility of tools for managing IAS. IAS tools and techniques seek to identify these all-too-often-exploited gaps in access and privilege control, highlight priorities for security teams and IT administrators, and call out tactics for remediation.

These techniques capitalize on the view of a potential target the attacker has in order to improve resilience against attack. They align with enterprise priorities for threat detection and mitigation, which, along with user behavior and phishing – among the most-exploited aspects of attacker targeting – were among the top five most-cited information security pain points among respondents to 451 Research's Voice of the Enterprise: Information Security, Budgets & Outlook 2021 survey.

Among the techniques seen in IAS security so far:

- **Beating attackers at their own game:** Exploiting an organization's IAS gaps requires the ability to make connections between identities, groups, access rights and privileges. When these connections extend across entities and groups by association, one way to grasp them readily, as noted earlier, is through the technique of graph analytics. In this case, however, it's the defender that can play the security duality to its advantage. By leveraging the same techniques attackers use, such as graph analytics, to visualize IAS exposures, organizations can also see and prioritize attack paths, and identify techniques to mitigate exposure that make it harder for attackers to act throughout an environment.
- **Enhancing vulnerability assessment and management:** The field of IAS analysis may itself be considered an aspect of vulnerability assessment. Vulnerability management systems can incorporate IAS tools and complement them with insight into how additional mitigations to exposures can be applied beyond raising attack path barriers. In security M&A, we have seen this complementarity manifest in deals uniting IAS and its adjacencies with vulnerability management vendors.
- **Better definition of identity and entitlements:** While this is overall a part of the zero trust trend, areas of focus in this definition are arising to sharpen the focus on access targets. CIEM in particular

addresses the complexity of management in the scale and variety of emerging cloud-native resources, while segments such as PAM, IGA and others aim to refine access and privileges for users and accounts and keep management up to date. In the network, the secure access service edge (SASE) helps organizations implement controls when access must be made available across wide-area networks, and not just on the dedicated enterprise network and its virtual private extensions.

- **Attack and exploit detection:** IAS exploits may often adhere to a sequence of activity that can be recognized as malicious. From an initial penetration of a host or endpoint, to the compromise of credentials, evidence of efforts to probe a directory environment, lateral movement and potential damage, threat detection systems throughout the IT environment can pick up signals that, when recognized as abnormal or adhering to distinct indicators, can indicate a potential attack in progress and defeat efforts before they have a serious impact. This plays a role in trends such as extended detection and response (XDR, covered at length in this 451 Research Technology Business Insight report).
- **Security analytics focused on IAS:** Even without telemetry specific to threat detection, organizations amass a wealth of monitoring data that can be better leveraged to recognize both IAS gaps as well as attacks. Security information and event management (SIEM) systems can be tuned to sharpen this recognition – and when they have a focus on user and entity behavior analytics (UEBA), they may be particularly attuned to identity and access abuses. Other approaches to security analytics that capitalize on the analysis of data at scale with high performance can be purposed to teasing out activity that may otherwise be difficult to distinguish from acceptable behavior.
- **The advantage of deception:** Another technique with its own distinctive approach to detection is deception technology. Deception systems deliberately appear to attackers as potential exploit targets – but they are instrumented with telemetry that alerts defenders to an attack attempt. As such, they can add real-time awareness to threat detection. When time is of the essence in responding to attacks such as ransomware, this capability can provide valuable early warning, as well as highlight where resilience can be strengthened against real-world attempts.

Is the past prologue?

It should be no surprise that a primary focus of IAS threat mitigation has its parallels to a previous era of attacks. Roughly 20 years ago, unsegmented and largely flat networks made it possible for worms to spread rapidly with damaging impact. These attacks took maximum advantage of the ability to scan other targets visible on the network, and made the most of automation to exploit those exposures and further their spread. A similar case prevails today, with comparably broad and inadequately disciplined access controls and privileges enabling attacks such as ransomware to spread just as rapidly and with even more devastating impact. Not surprisingly, similar approaches to segmentation and constraint – applied not only to networks but to identity and access controls – can be expected to play a role in threat mitigation today.

To be sure, this likely isn't a trivial undertaking for any organization. As with 'zero trust' initiatives in general, restraining access and privilege can hinder the business if not approached with care to balance business and risk priorities. In the light of today's attacks, however, many organizations must take IAS more seriously and weigh their options. More closely defined groups and access privileges, segmentation of groups and memberships to better insulate assets from potential exploit and stronger controls on authentication such as multifactor techniques, all play a role.

Looking ahead, it seems likely that there are further lessons to be applied from threats of a previous era. With the rise of mobile devices, security models emerged that sought to define privileges with greater granularity and insulate applications more distinctly from the underlying system such that an exploit of an application would not necessarily lead to an exploit of the underlying system or other applications, in part because access and control privileges between those layers were more distinctly defined. Today, we see IT vendors moving toward similar approaches, with concepts such as security enclaves taking hold in the datacenter, while on user devices, vendors such as Microsoft are introducing 'chip to cloud' initiatives with new operating system releases for strengthening

authentication and access controls on endpoints to limit the impact of attacks and extend protections to access targets. As the realities of IAS continue to take hold among defenders, we expect to see even more activity in the field as organizations seek to make the most of their advantages.