

Managing a Different Kind of Identity

By Mark Mac Auley

The recent hacking of the account information of 19,000 AT&T customers is yet the latest example of why protecting and managing identities is so important to a security professional.

The recent hacking of the account information of 19,000 AT&T customers is yet the latest example of why protecting and managing identities is so important to a security professional. You may be thinking I am referring to customer data, and I am not. It is the management and protection of corporate identities I am talking about – beyond usernames, passwords, single sign-on, Identity Management (IdM), IPS, IDS, firewalls, ACLs, and the list of layers goes on. What I am talking about is the identity of the *devices* used to access your organization's information and the assets within the organization.

Is the device known? Is it trusted? Is it a healthy device? Is it malicious or could it be made malicious? What user identity is bound to the machine?

You may be thinking, "I have tokens and multi-factor authentication." That's great to add another authentication type for *users*, but tokens don't identify *devices*, and since users are only half the equation, device identities are as important as user identities on the network.

One of the few absolutes I've encountered over the past few years is that all hacking occurs after establishing signaling. If you can't connect physically or virtually, you can't hack. It's the equivalent of dial tone, where if you do not provide a way to establish dial tone from a device/phone (known or unknown), there is no way to connect to another network or the devices on it. I can't hack something I can't connect to, or whose destination I don't know. I can't dial any number I want to, either. If I don't have international calling service on my phone, my clients and colleagues in Europe won't hear from me – good thing?

The other piece of this metaphor is that in order to deny me service, the phone company denies me dial tone. They don't come to my house and ask me to give them my phone. I still have a phone, I just can't use it. I could have any phone, in fact, but I must be able to have or get dial tone to establish signaling with another device.

So in the IT security realm things are a bit different. We have computers (phones) with IP addresses (phone numbers) that make calls every day on our networks locally, long-distance and internationally. We've spent many years managing these phones and their numbers, and mapping them to account information (usernames and passwords). Recently, Identity Management has been analyzed and deployed globally as a way to provide access to certain phones that call certain numbers, based on whose account it is and what they can dial – but this has been at the application layer. I equate it to being on hold – I'm still inside the network, and can "zero out" to try another number/extension.

Caller ID came along, and this gave us the ability to see who was calling and decide if we wanted to pick up, essentially authenticating the caller and establishing signaling with them. Just like IP address and MAC address spoofing, I could call anyone and pose as "Steve Jobs" and spoof you into allowing the signaling and picking up. The data about the endpoint (caller and device) has been compromised, allowing bad things to happen.

So how do we get back our control of something we thought we had control over, but really don't? With identity-based access control, the endpoints (*users and devices*) are distinguished with an unalterable identity that ultimately gives the security and network groups control of the dial tone of the device. It allows IT security professionals to proactively manage access across segments of networks, based on user and/or device attributes that are checked against policy before signaling is established. This is the equivalent of saying that if you have a phone and you call me, I don't have to pick up or even route the call before I know that: It's you, and from your phone; the number is your number; and it's absolutely correct.

In this approach we have spent a great deal of time focusing on networks, infrastructure, applications and data separately. What we are seeing is the exploit of the silos, and so we must understand the need for convergence of the approaches used to protect these components – otherwise exploitation will continue to occur.

The discussion must take into account the key differences between these silos: governments, financial institutions (large and small), hospitals, and educational organizations – specifically in terms of application vs. network IdM or Access Control. One thing should keep ringing through our heads – managing identities of devices and users in the network layer. The infrastructure layer and the application layer are included in this. An application-based solution, on the other hand, focuses on the application layer, and still allows access to and visibility of your network.

If you are an unknown/untrusted user on my network – hacker, contractor, visiting doctor, pharmaceutical rep – you are inside my perimeter. That's where hackers hang out to learn what is happening and how to get around: just inside the perimeter. They rarely walk in the front door guns blazing, because they don't know what they're shooting at. What subnets do you have? What applications are running in those subnets? Are you a Cisco or a Foundry shop? Are you using Oracle HR or SAP? What ports are open on the servers that host these applications? Where are the vulnerabilities? It's the difference between the Lewis and Clark expedition and a GPS-enabled road trip.

In the application layer, you cannot control the network to block access to assets. This brings up an issue and a recent personal example I want to look at – offshore vendor management. US-based companies struggle with needing to extend their control and enforce their security policies offsite and out of country. The question has been asked, how do we limit access to networks with production systems to a few key people, while enabling access for larger groups to other systems on other subnets? In this, we also want to have some level of control over what assets these other groups or users can see and interact with. In essence, the problem is about customer databases, dev, test, staging environments, and other applications that have high levels of restrictions for export control or compliance.

Whenever I come across this question, I explain that this scenario and others, where you need or want to extend the security policy enforcement outside your walls, is exactly what identity-based access control is designed to address. To this end, an IdM solution that operates in the network layer can provide several significant benefits over an application-layer IdM solution:

1. Endpoint control and security policy enforcement outside company and geographic borders
2. Protection of network, infrastructure, applications and data with a single piece of technology
3. Greater flexibility of policy enforcement by component: users, groups and devices

Here is the last question I will leave you with in this three-part series on Identity Management. What is the quantifiable difference in risk between internal and external user access?

I will argue that an internal user brings significantly more associated risk than an external user – yet most of the focus seems to go to the external user, especially when bad things happen. You could argue it the other way too, given the rise of breaches, and that is exactly what our lawmakers are doing right now. I believe that from a policy perspective, they are trying to create transparency and equal weighting of internal and external breaches, and force organizations to give equal weight to protecting the assets of the organization – to really think about how they should do it. The issue I see is that insiders are far more dangerous – they are already inside, and are considered known and trusted.

Think about the real-life differences for a moment. Do we require most external or unknown users to have a background check before they try to access our organization's assets? How about a drug screen? Hiring Manager approval? No, but we require it all of internal users. The underlying issue in all of this is trust. Establishing trust is a process, not an event. It is up to each organization to look at and implement a solution, program, mindset, or process to put the establishment and monitoring of trust inside and outside the organization. It's a delicate balancing act we all know too well: Provide enough openness and access to allow the organization to operate, while insuring the highest levels of security, so as not to paralyze the operation entirely. Oh, and don't forget to insure privacy of your identities, inside or outside the organization.

The next thing on the horizon is “NAC.” I have been talking to many companies about it, and the first question I ask is, does NAC mean “Network Access Control” or “Network Admission Control”? It's the difference between machine health and device access, and I believe it's the next “Identity Management.” It means different things to different people, and the definitions will mature along with the business problems that are solved by implementing NAC.

We'll discuss NAC next time around.

About the Author

Mark Mac Auley is a member of the ISSA, IAPP, InfraGard, and the IISFA. Mark runs a blog at <http://identitystuff.blogspot.com> and has consulted on and sold identity management solutions for the past four years. He lives in New England and is currently the Northeast Manager at Trusted Network Technologies. He can be reached at identitystuff@gmail.com.