Strategy for Identity Management

By Michael Carper

As a specialty apparel retailer (Coldwater Creek®), we face many of the same challenges as other retailers. We do a significant amount of seasonal hiring requiring a great deal of on-boarding and off-boarding. Secondly, we have a good number of "silo-ed" platforms that require a separate username and password. These challenges necessitate a strategy for managing identity within the enterprise.

As rare and unique as some names may seem, there are certainly cases where multiple employees with the same name will work for the same company. I've experienced instances where a person with my name was working for a company I had worked on projects with, and another case where a realtor of the same name currently works in my market. In the case of the realtor, I've gone into various stores recently and they've asked for my name to pull up my account. When they ask to verify my address or phone number, it quite often belongs to someone else. I sometimes wonder if he starts flying a lot in the future, will he somehow inherit my frequent flyer miles?

So, how can we know that the "Jane Smith" who works in finance and has access to general ledger information is the same "Jane Smith" who has lesser access to other systems—but is not the same "Jane Smith" who is a sales associate in one of our stores? To address this, there are four essential elements that must be considered: definition of roles, enterprise directory services, single sign-on, and multi-factor authentication.

Definition of Roles

Perhaps the most important element of any identity management strategy is the definition of roles. There are many approaches to this task—but it is rarely easy. It is generally an effort best led by the human resources (HR) organization. Some companies may take the approach that every job description is a role. In other cases, it might be fair to use a common role for a good number of job titles.

For example, a hospital may have many different job titles for a nurse. Some of these nurses may actually work in the same unit and require access to the same systems. In that case, it might be appropriate to use a common role for all of those nurses. On the other hand, other nurses may work in surgery, or in a clinic, and they may have dramatically different system access requirements. To make things even more complicated, some nurses "fill-in" in other areas when staffing requirements mandate it—effectively working in one role during one shift, and working in another role on another shift.

This is where the rule of "least privilege" comes into view. That is, a given employee should not have the ability to access information or systems in excess of what's minimally required to execute the duties of the job. When new employees join a company, their identity can then be created based on a template for that role—rather than to mirror "Jane Smith's" account. Jane may have worked previously in sales and human resources and may have access to much more than what her current role requires, if accounts are not managed by role and the rule of "least privilege."

Additionally, applications should be delivered based on role. Some nurses may require access to medicine administration systems to see what medication should be provided to a given patient. Another nurse may need access to lab systems to see the results of tests. If these systems are different applications, there's no reason to present all nurses with this application. If these systems are part of the same application, then field-level security may be required to ensure that only the appropriate information fields are available.

Enterprise Directory Services

Directories have been around for quite some time. At the very least, they're about as old as telephone service. In most markets, the phone company still prints a directory of its subscribers and distributes it to them. Even in the smallest of towns, one could find duplicate—but valid—entries. Beyond that, all the data in the directory was probably not valid on the day it was printed due to people moving and changing numbers.

In most companies today, directories were created to support the e-mail system. Given that Microsoft Exchange is the most pervasive e-mail system in corporate America, almost everyone has adopted Microsoft Active Directory® (AD) to support Exchange. That makes perfect sense—for Exchange. The problem is that the e-mail directory ends up becoming the primary directory for most companies. I'm sure there are some companies where every single employee has an e-mail address—ours is not one of them.

Perhaps as few as 50 percent of our employees have an e-mail address. The remaining employees are sales associates, food service employees, janitorial, or work in similar roles—often part-time. Though a sales associate may not have an e-mail address, they'll need access to the point-of-sale system. Additionally, all employees are in the payroll system, but there's no way to know that the "Jane Smith" in the payroll system is the same person whose e-mail address is jsmith@yourdomain.com. Then there's the next layer of challenge that comes from the request that sounds like this—"We have a new employee joining finance. His name is Robert Lucas. Please mirror his account to Jane Smith's." I think you can all see the point.

The catch is—generally speaking—no one can use the payroll directory for authentication. Besides, what about contractors—normally not entered into the payroll system? Most companies authenticate their systems to AD which was built to support an e-mail system. The best way to resolve this might be to build another directory. The sole purpose of this new directory is to be a "directory of directories." It is the directory that knows about all the various directories in the environment. More importantly, it knows which directory is the "source-of-truth" for a given directory attribute. It may not hold all the directory attributes, but would hold pointers to the sources-of-truth.

For example, the human resources (HR) system may be the source-of-truth for many things—employee number, title, work address, and so on. But it would not be the source of truth for an e-mail address—that actually would be Exchange. Neither of these directories would be the source-of-truth for a phone number. So the job of this new directory, or meta-directory service, is to propagate the information from the sources-of-truth to other directories as appropriate.

With the enterprise directory service in place, and roles now defined, system access can be provisioned automatically from the HR system. When a new employee joins the company, they will be processed in HR and assigned to a particular role. Based on that role, accounts are created in various systems—point-of-sale (POS), order management, inventory, accounts payable, shared data directories, the e-mail system, and so on as appropriate for that role.

Also based on the role, software applications are delivered. Most all roles might require Microsoft Office, for example. But perhaps only project managers should have Microsoft Project Professional installed. The installation of this software should be done automatically based on role. Given that, we can provision access when new employees are hired. We can now de-provision access when employees leave the company. We can also change their access when their role changes. Better still, IT resources or requests are no longer a cog in the process. This will all happen as a routine part of HR's existing processes of managing employees in their system.

Single Sign-on

Authentication should be done, when possible, against the meta-directory. It's common for most modern software applications to have the capability of authenticating against AD or an LDAP directory. The strategy should be to reduce authentication to a single event for the user—the initial logon for the workstation.

Some legacy applications cannot support this type of authentication. To accommodate this, a single sign-on platform should be used. This platform should have a profile built for every business application that cannot authenticate against the meta-directory. It should be transparent, so that the user will not be prompted or hindered.

This provides for more stringent password standards. My recommendation is that the primary authentication should require at least eight-character passwords and should include non-alphabet letters—such as numbers and characters. However, because we're going to harden other passwords behind the single sign-on platform, we may not need to change the primary password as often as once thought. Changing every thirty days is probably over-kill. Passwords on all other systems should be set to the most stringent standards of which the given platform is capable.

Perhaps fifty-character passwords should be used on other platforms. The key here is that the users will not know these passwords—the single sign-on platform will pass the credentials to these systems. This strategy also ensures that a person would not be able to avoid the workstation's primary authentication to login to a system directly.

The single sign-on platform should provide for self-service password resets. This would enable users to provide some unique information about themselves—such as favorite food, first pet's name, and so on—and be taken from the logon screen to a form that enables them to change the password. Given that we've reduced the number of passwords users need to remember to only one, the chances are better for remembering the password in the first place.

Multi-factor Authentication

Now that we are provisioning users automatically, based on role—and the users only have one authentication event—the next step is to get rid of username and passwords all-together. Many industry analysts believe that "usernames and passwords" may be a thing of the past in the years to come. Having all these other elements in place is a pre-requisite.

There are three factors of identity:

- what you know,
- what you are,
- and what you have.

The traditional username and password model is a single-factor means of authentication, as it only addresses "what you know." Once that knowledge is obtained by someone else, they can use those same credentials for malicious purposes. Most security experts agree that two-factor authentication should be the minimum standard.

There are many products coming to market today to enable two-factor authentication. Among them are fingerprint readers, hand-print readers, and retina scanners to name a few. In most cases, a simple swipe card should suffice—the "ATM machine model." These have become quite common in restaurants. A waiter accesses the POS terminal by swiping a card and entering a four-digit personal identification number (PIN). This example uses "what you have"—the card, and "what you know"—the PIN. A fingerprint reader is an example of a "what you are" approach.

The multi-factor approach is applied to the primary authentication for the workstation. The single sign-on platform handles the passing of credentials to other systems. The result is speedy access, without hassle, for your users into all the systems they need access to.

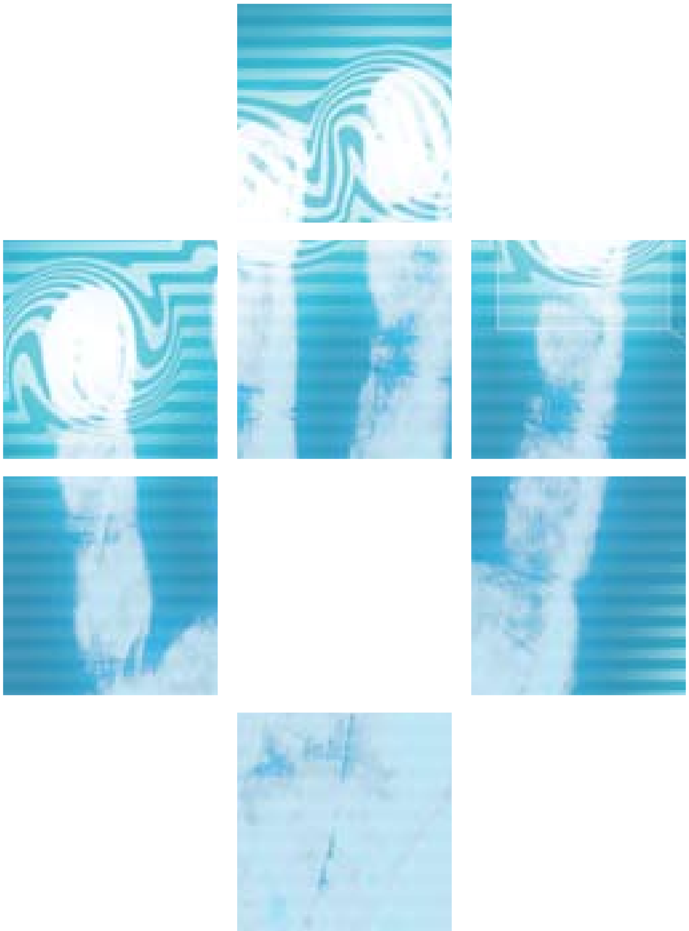

Conclusion

This approach is not one that can be pursued with haste. This is a thoughtful process that requires a great deal of strategic thought. This will change work-flow processes, in many cases. Many departments, like HR, must be actively involved. Training may be required. The technology service desk must be prepared. In the final analysis, identity management cannot hinder the business. It must be implemented as an enabler of efficiency to the business. Only there, will it provide value.

About the Author

Michael Carper is the Divisional Vice President of Technology Operations at Coldwater Creek, a retailer of women's apparel, accessories, jewelry, and gift items. He has eighteen years experience in the engineering and support of telecommunications and information technologies in the retail, healthcare, financial services, wireless, and telecommunications industries. Mr. Carper has a bachelor's degree in electrical engineering and technology from Purdue University and a master's degree in Information and Communications Science from Ball State. He was recently an adjunct professor at Ball State University, where he taught graduate courses in wireless communications. He is currently pursuing a doctorate in technology management through Indiana State University. His leadership in technology has been featured in *Computerworld*, *CIO Decisions* magazine, and in case studies by Dell, Cisco, and Intel. **mcarper@thecreek.com**