HealthManagement
Promoting Management and Leadership

Volume 1 / Issue 4 Winter 2006 - Cover Story

Managing Security Access

Author

Chris Sullivan

Vice President, Customer Solutions

Courion Corporation

csullivan@courion.com

www.courion.com

With aging populations, healthcare continues to expand into one of the largest industries across the European Union. Along with this have come vast changes in both organisational structures and supporting technologies. This, in turn, has driven tremendous complexity into the way access is granted, controlled, and revoked for both patients and those who care for them.

While patient privacy, safety and care have become the mantra for most institutions, this cannot be achieved unless access to the thousands of supporting information systems is provided quickly, accurately and securely. Against that backdrop, security professionals toil to manage risk when most of them don't even know who has access to what, much less whether that access is appropriate. To make matters worse, their raison d'être is to enable patient care and improve safety - which means that there is a very low tolerance for risk.

Put simply, in an environment of increasing complexity, staff must have access only to what they need, exactly when they need it and without hiring an army of highly paid security administrators.

The Task Seems Daunting But the Solution is not

The good news is that these problems can now be solved. Automated provisioning and compliance solutions are enterprise-ready and there's a growing body of best practices that can be applied to yield real benefits. As we explore these approaches, keep in mind that the benefits of effectively managing access come in three measurable forms – speed, sustained efficiency and transparent control.

Speed means that we can accelerate the way care is provided by giving individual clinicians what they need more quickly, or by enabling an acquiring organisation to integrate a new facility in less time.

Sustained efficiency means achieving operational savings without the solution being worse than the cure. Done correctly, dramatic savings can be pumped back into care. Done incorrectly, the provisioning system itself can become a white elephant that consumes more time and energy than it is worth.

Finally, transparent control means embedding preventative and detective controls into day-to-day processes in a way that reduces risk without imposing additional burden on the clinicians.

Best Practices for Managing Access

(1) Crawl Before You Walk

Let's get started with best practice area number one. Do not try to implement a comprehensive Identity and Access Management (IAM) programme as one massive project. I've seen no evidence that this has ever succeeded. What you are really about to automate are detailed processes for staff on-boarding, change, termination, and periodic review. These processes are dependent on security and operations policies that will vary by type of care, location and even management level. This can't be done in one monolithic effort for two simple reasons:

• Mostorganisationsdon't understandtheirown policies wellenough to spec-out a solution; and

- Even if they could define things in sufficient detail to coordinate an army of offshore .NET and Java developers, it would take years to complete – by then, the problem will have changed.

A more natural approach is to define a programme around a vision for efficiency and control and then begin with concrete projects that support specific goals. Each of these projects should be measured in business terms (to garner support), be simple and bounded (to minimise risk) and extensible towards the longer term goals.

Delivering value quickly and consistently will build support and momentum.

(2) Always Move the Ball Forward

As you consider initiatives, evaluate their impact on speed, efficiency and transparent control. You should always be advancing one of these things and it will be common to advance all three. Deploying an account request process with more formal approvals will only decrease risk if the staff actually uses it. If it's harder for end-users, they will find a way to circumvent the process and there will be less control. If you are creative, you will find a way to reduce risk and hassle.

Remember, incremental progress is better than delayed or unattain-able perfection!

(3) Know What People Have

Any business school will tell you that you can't manage what you can't see and this holds true for identities.

If you don't have a current map of who has access to what, then how do you know if people are over-credentialed? How do you disable their access when they leave? How do you even help them when they call the service desk?

Building this map can be difficult because most legacy environments are not very consistent, but there are effective tools that can help:

1. Establish a unique ID for all users;

2. Pull accounts and attribute information from core systems;

3. Map those account names to the unique ID:

- Consider policies that were in place when the accounts were created;

- Balance accuracy against the risk of making the wrong association;

4. Claim accounts that you cannot automatically map. If I am the only Sullivan at Courion then it's probably safe to assume that the AD account csullivan belongs to me, but if there's also a Clarice, you might have us identify and authenticate these against our accounts to claim them; and

5. Keep these mappings current with maintenance scripts.

Congratulations, you've just implemented some important controls and you are well positioned to automate disables completely!

(4) There is a Role for Roles

I've seen many successful role implementations and many unsuccessful ones. Roles are hard because you must work out what each person should get access to and then you must validate that with application owners to validate that access. However, since you should be figuring this out EVERY time you change someone's access anyway, why not do it once? Start small and build an approach that will both scale and accommodate change. For your first foray into roles:

1. Select a modest population, perhaps legal;

2. Work with them to define a representative set of job functions;

3. Assign appropriate access rights to them. In practice no one will know just what to assign, so ask them for representative users and consider what they have;

4. Scrutinise roles against security policies. Now you can redirect legal requestors to a simpler workflow that simply asks them to choose a pre-approved role and access can be securely granted without additional approval.

Going forward, you'll have to scale what you learned. As your approach matures, you'll want to be thinking about the following:

- Keep the number of roles manageable. Perhaps 200-300 for a 40,000 person organisation;

- Roles should be dynamic and rights assigned based on policies. In this way, when the policies change, you don't need to re-engineer the roles;

- Select tools that can automate provisioning and compliance with or without roles and be sure that they support role lifecycle management (developing, creating, changing, periodic review, governance and change control);

- Implement a governance process; and

• Avoid temptation. Under and over-credentialing are simple, but the former doesn't add much value and the latter creates risk.

Advanced Techniques

I have friends and colleagues who have implemented robust identity management programmes that are doing everything that we've discussed here. They have deployed enterprise roles for 80%+ of their users with only 200 roles. They have provisioned 5,000 new users from an acquired institution over a single weekend. They have reduced security administration staff by >70% and cut millions in operating costs. They've cut service levels from weeks to minutes, all while reducing the effort expended for internal and external audits to a fraction of what it had been.

Today, they are leveraging their infrastructures in ways that you might not have imagined. Since they've automated the employee on-boarding process, why not add in physical security and manage access badges to the floor and door level? Now that you have decided what rights a specific type of user should be granted, why not go back and track what they actually use? I have one customer who found that they provisioned 17,000 accounts in the last year for an application that was only used by a few people – that's a lot of labour and unnecessary risk.

Back to Basics

Remember, it is better to have incremental progress before delayed perfection. In this case, progress is defined in terms of speed, efficiency and transparent controls. Make sure you know the clinical and technical context that you're dealing with and execute short, successful projects that will build on each other to advance your goals. If you are just getting started:

• Build and maintain an identity map. It will help you in more ways than you can imagine;

• Scrutinise orphaned accounts. If you can't map them, they probably shouldn't be there;

• Automate the disable function. Granting new rights quickly and efficiently can be challenging because you need to understand how policies translate to system attributes. If you have an identity map, disabling is pretty easy – set the revoke attribute; and

• Get started with roles.

Finally, measure results! Institutions value patient safety and care and will support those who can show how they enable it.

Published on : Mon, 1 Jan 2007