IBM Global Technology Services

www.ibm.com/services

IBM Podcast

Stopping insider attacks: how organizations can protect their sensitive information

Ben Edwards, IBM CHQ Communications, interviews Mark Ramsey, IBM Global Business Services and Stuart McIrvine, IBM Software Group, on the growing threat of insider attacks

EDWARDS: We're here today with Mark Ramsey, Global Business Services, Global Data Analytics leader for IBM Center for Business Optimization; and, Stuart McIrvine, IBM Software Group, Tivoli Director, Corporate Security Strategy, to talk about security and privacy for businesses and in particular what they need to consider in light of the increasing phenomenon of insider attacks. Mark and Stuart, welcome to this IBM podcast.

RAMSEY: Thanks, Ben.

MCIRVINE: Hello.

EDWARDS: Hi. Stuart, if I can just start with you. Maybe you could just describe to us what insider threats are and why they differ from other sorts of security threat.

MCIRVINE: The insider threat really is a threat that comes from someone who's generally authorized to access a system, okay? So the insider, it doesn't just mean it's

somebody necessarily inside the company; it could be a business partner, it could be a customer, et cetera. But, they've been at least provided the authority to access certain applications and data, et cetera. And that differs from what you would more commonly recognize as a hacker. And that's someone from the outside who really has no authority to penetrate the perimeter, have access to systems, et cetera.

EDWARDS: It sounds like it's a much more difficult threat to tackle.

MCIRVINE: Yes, it is, because you know, many companies today have done really a pretty good job in shoring up their perimeters. They've added firewalls, intrusion detection systems, intrusion prevention anti-virus capabilities, et cetera.

So really kind of bolstered up that perimeter, and you know, prevented a lot of attacks that come in from the hackers, et cetera, on the outside. And you know, and there are plenty of very advanced technologies today to do that.

Then they also, when you start to look inside the companies, they have identity management systems, access control systems, where they have to register the users that are allowed to access the systems and only these users can access them.

The problem is, once you've actually registered these users, right, then you know, there are other more sophisticated techniques that are actually required to make sure that, let's put it this way, that these users kind of behave within the appropriate paradigm for the company.

EDWARDS: So Mark, what sort of techniques are we talking about here?

RAMSEY: I think a lot of it is around, like Stuart mentioned, it's looking at the behavioral patterns of the users that are authorized to access the system.

So one of the things that you find is that there's been a large focus over the last few years on identity validation techniques, so things like biometrics, and passwords, and passcards, and all of those things do a very good job at validating that the person is actually who they say they are and that they then pass that level of security of being authorized to access a particular environment...

The real thing that we're looking for around insider attack is like Stuart mentioned, once that person has been authorized, are they really behaving within the boundaries of the policies that are established within an organization.

EDWARDS: Could you give me some examples of that?

RAMSEY: Certainly. There's been quite a few things published here recently where authorized users within an organization have actually violated that trust.

So for example, it could be a contact center agent within an organization that has access to sensitive information, whether it's Social Security numbers or a date of birth or even account balances within a bank, for the customer there have been several situations where that information was actually taken by those authorized users and used externally for purposes such as selling Social Security numbers to the

marketplace to use for identity theft, for example.

EDWARDS: Okay. I can see how this would work well with patterns of work which are well defined repetitive work that's well defined, the workflow is well defined.

What about I think about my working day, maybe I do have well established patterns of work but I kind of feel that it's quite unpredictable. I might be doing one thing one day, one thing another. Does it work for that sort of working pattern too?

RAMSEY: Yes, I think what you need to understand is really the approach here is to look at a peer group of individuals. So for example, even though your work pattern does vary on a day-to-day basis, if we look at all of the individuals that have the same type of job that you have, in general that should tend to group together into some similar behaviors.

And then you can look at where those deviations from that occur. So for example, back to the conversation around the contact center agent, an organization may have hundreds if not thousands of contact center agents, but the ones that operate in a certain area, so for example, maybe there's a group of a few thousand contact center agents that work on billing types of questions from customers...

Their behaviors in general will tend to group together and what you would be looking for around insider attack is situations where people are beginning to deviate from that norm.

EDWARDS: Okay. Stuart one obvious question this raises is privacy, I mean, how this connects with a company's privacy policy and how a company might reassure its employees that they're not snooping on them in some way what.... And how do you start thinking about that?

MCIRVINE: So, I mean, there are a number of different ways to look at privacy and so you know, monitoring user behavior may be one aspect of it, but you know, when we tend to look at privacy, we focus on things like how much...how many of attributes that actually identify that individual are we exposing? Right?

What we're looking at here are groups of individuals, groups of individuals who all have a similar business function and therefore should have kind of similar patterns of behavior, and not necessarily trying to single out one person in a company and monitoring their every move. Right?

So it really is looking at kind of groups of people, roles of people in an organization, okay? So you know, it's somewhat limited when we look at it in that perspective in terms of what we're exposing from a privacy point of view.

And you're really not going down and actually exposing attributes about that person trying to really identify the individual.

EDWARDS: And what sort of attributes, what sort of behavior attributes might you be tracking?

RAMSEY: As far as the attributes that are being monitored, this is something that really varies a great deal between organizations, but you can think of situations where knowing when people are actually connecting to systems, how long they're connected to systems, what type of data they access when they're connected to those systems...

Sort of across the entire life cycle of an interaction with applications that contain sensitive data, those in the aggregate would be what is used to build up the profile that Stuart was mentioning.

So if we think about it, the fact that someone is connecting to an application that contains sensitive data and they're doing that on a weekend, or they're doing that at 2 A.M. on a Wednesday when the profile is that those...that data is accessed during normal business hours, those are the kinds of things that make up this multi-dimensional analysis that really is the basis for the...the peer group.

EDWARDS: That's interesting. So you're not actually looking for specific types of behavior, but you're looking for deviations from regular types of behavior.

RAMSEY: Exactly. Now, and the reason for that is that...and you mentioned this, the jobs fluctuate a great deal, so for example, in contrast to more of a rules based system where you would establish very specific rules of what people can and cannot do, what we're looking for are those slight deviations that occur within the authorized behavior.

So it's really looking for differences to the normal behavior for an individual peer group. And this is the type of technology that we've used extensively in the fraud detection area because again, being different from the norm is something that typically is an identifier for not necessarily fraud or bad behavior, but it is something that should be further reviewed. And that's what we're looking at for the basis for this deviation detection.

EDWARDS: Okay. So how good is this sort of deviation detection approach at preventing insider attacks as opposed to just sort of detecting them after the event, if you like.

RAMSEY: Right. A very good question. I mean, we tend to look at the...at this issue in a couple of ways. One is, retrospective analysis, which is, when you're looking at what folks have done historically in building up that peer group and that profile around the normal behavior of individuals, and that's something that you can identify where you may have had a problem in the past.

But certainly the area where there's a great deal of focus is actually protecting from the event actually occurring, and that's where a real time or a prospective solution needs to be put in place.

And what we've found is by coupling together the historic analysis and the real-time, it allows us to prevent or at least shut down a situation that may be suspect. So for example, if I go back to the contact center example, once we establish that normal behavioral peer group we can then monitor the actions as they're occurring.

So for example, if a normal contact center environment, if the connections are at a certain time, and the access to data is a certain amount, if we start seeing that a particular individual is deviating from that...

So to my example, if a normal contact center agent accesses the HR system between eight and five -- normal business hours -- and suddenly we have an access that comes in at two in the morning, we can actually post an alert to the access management software which would prevent that user from connecting to the system until the particular business need was validated. So it is a way to have an early warning and potentially prevent a loss of sensitive information.

EDWARDS: Right. I guess one thought here is, as crime fighting or as fraud detection becomes more sophisticated fraud itself becomes more sophisticated too, right? So you know, the...those with criminal intent evolve and become more clever.

So perhaps as the success of this system grows people will learn not to access the system at 2 A.M.. So what would you say to that? I mean, how do we head that off?

RAMSEY: The interesting thing here is that our approach to the insider attack threat, again, is not tied to specific attributes or specific things being monitored. It's tied to looking at across a collection of attributes and understanding normal behavior.

So what I mean by that is, as the attempts get more robust, and the example I shared with you as far as connecting at 2 A.M. versus business areas, I mean, that's obviously

a red flag situation.

What we're really looking for is lots of little shades of gray versus such a large red flag. And so the important component here is that the...those normal behavior patterns continue to be updated. And that is really important for two reasons: first, it's important for the reason that you mentioned, which is the types of attacks will evolve over time, so we do not want to be tied to very specific decisions on whether a transaction is a potential insider attack or not.

And second, the actual nature of the business changes over time, so you do not want to generate a lot of false positives. So for example, if we look at something like the number of transactions that occur within a specific period of time, if the business is a cyclical business where at some times of the year they have much higher transactions than others, the system needs to be self correcting in that area.

So the approach that we've taken is one that actually continues to evolve the normal behavioral patterns so that as things get more refined either on the type of attack or from the perspective of the business changes, that the models will evolve naturally over time.

EDWARDS: Right. Stuart, as chief security officer of a business, how do I know what the risks are I face from insider threats? I mean, do they differ a lot from company to company and industry to industry?

MCIRVINE: Yes, they certainly do differ a lot. And it depends, I mean, a lot of the time it depends on the employee. I mean, the risks are higher if you have, let's say kind of low-skilled hourly-paid employees that can potentially be easily influenced by outsiders to steal information, et cetera. Right?

There are high risks when you start to look at the financial services industry, because the actual information there is extremely valuable. Right? And therefore there can be a higher incentive for even highly professional individuals, right, to go after that kind of information in the wrong way.

So I mean, you could...you could take a different slice of it for every company you look at, not even just every industry. But you know, a lot of it comes down to the type of employee, how easily some of them can be influenced, right.

How lax security is, because it's not just always that it's someone that's necessarily authorized to access the system that can do this. You know, it's some unauthorized person could be perfectly legitimate person who just leaves their ID and password lying around. The wrong person obtains access to that.

So you know, if the company's policies are somewhat lax, they may not they use passwords a lot rather than biometric devices, all these different things affect how risky this can be within a specific company.

EDWARDS: Yes, that raises a good point which is that this doesn't replace IT...ID validation, right? I mean, you need...you need that basic security of making sure we

know who has access to the system and that they are who they say they are.

MCIRVINE: Oh, absolutely and you know, we push a lot at IBM that multi-factor authentication is a far better way to go. You know, there certainly is expense involved with it, but it significantly reduces the risk and significantly increases the likelihood that the person that's logging on is who they say they are, which is an extremely critical aspect of all of this.

EDWARDS: Lots of varying degrees of risk across company and industry, but you know, on the aggregate level, Stuart, what's the magnitude of this problem?

MCIRVINE: If you look at just the amount of attacks that occur in general, I don't think anyone argues with the statistic that you know, about three quarters of all attacks actually come from known users. Okay?

And that is continuing to increase. And it's increasing because companies are getting better at dealing with the outsiders, dealing with the hackers, the people who are not authorized to access the system. So you know, that's an extremely large number, 75 percent of all attacks come from...come from known users.

EDWARDS: Right.

MCIRVINE: And the fact that that actually continues to increase is important. Then the other point that we need to take account of is we have our own security intelligence services in IBM and we monitor a lot of this stuff.

And you know, we are highlighting that some of these, just the individual attacks that are becoming extremely targeted now, they're becoming much more targeted, and the actual value associated or should I put it more appropriately, the loss associated with each attack, is increasingly significantly. I mean, it's now into literally billions of dollars now companies are losing a year.

EDWARDS: Yes, what's behind that trend? Why's the loss per attack increasing?

MCIRVINE: Because they're putting.... I mean, first of all, what's behind it now, it's not your individual hacker that's submitting a virus just for fun, seeing what kind of damage it can do.

You actually have organized criminal gangs now, right, that are trying to recruit individuals either inside or outside that can get insider authentication details to go actually extort money from companies now.

So you're seeing a shift in the way, in the type of attacks that are occurring. It's no longer just for effect now; it's actually for financial gain from organized criminal gangs from all over the world.

EDWARDS: Right.

MCIRVINE: Because this can obviously be done very remotely, too.

EDWARDS: Okay. Well, Mark Ramsey, Stuart McIrvine, thanks very much for your time and your thoughts today.

MCIRVINE: Thank you.

RAMSEY: Thank you.

[END OF SEGMENT]