

NOTICEBORED

Management briefing on Authentication metrics

Summary

This paper outlines a number of potential information security metrics relating specifically to user authentication.

Introduction

User authentication controls are designed to check the identity of individuals logging-on to a system and thereby prevent unauthorized people from accessing controlled resources (information assets such as data, systems, functions or networks). How do we tell whether our authentication controls are effective? What does “effective” mean in this context? We’ll try to address such concerns through this briefing paper but you will need to think long and hard about how to interpret our suggestions in your specific situation. With security metrics, one size definitely does not fit all.

Authentication requirements (targets)

If we are going to report the status of authentication to management, it is helpful to understand first what management expects of authentication in order to identify whether we fall short, meet or exceed their expectations. What are our targets? What are we aiming to do?

Authentication has essentially one central objective with two opposing aspects:

1. Identify known individuals who should be permitted access to our assets. These are the good guys.
2. Distinguish them from other individuals who should not be permitted access. These are the bad guys.

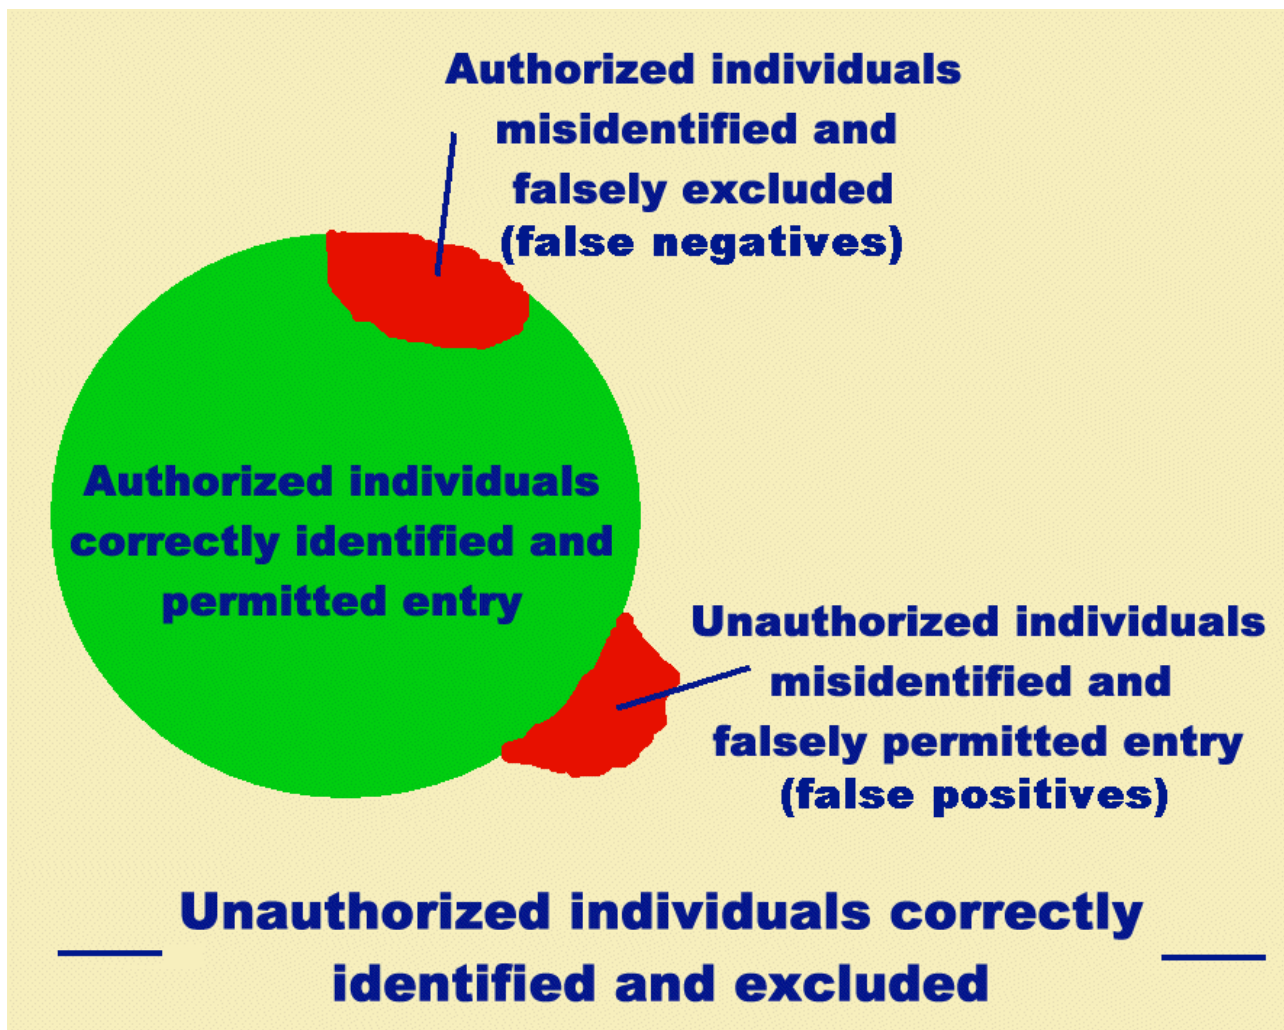
People who can be absolutely positively identified as either good guys or bad guys are relatively easy to deal with: we know immediately whether they should be permitted or denied access to our assets. The key problem with authentication relates to those people who cannot be positively identified, the ones who may be in either category. If we are uncertain, we have some difficult decisions to make.

People who are wrongly identified represent authentication failures - those bad guys who are incorrectly identified as good guys and are falsely permitted access (“false positives”) as well as the converse, those good guys who are incorrectly identified as bad guys and are falsely denied access (“false negatives”). Both types of failure can cause us problems but the risks are often different.

Successful hackers, for example, represent one type of false positives. In most circumstances, hackers would create terrible problems but not always – some systems are deliberately designed to be open to practically anyone (www.wikipedia.com for example), even including well-behaved hackers.

False negatives are authorized users who are prevented from accessing the systems they should be able to use. If the system contains secrets, this might be an acceptable situation because at least the secrets are safe. However if the system is, for example, the computerized ignition on a vehicle, it could create real problems.

These situations are illustrated on the diagram overleaf.



In reality, organizations often have a mixture of systems of each type – some where false positives are OK but false negatives are not, and some the other way around - but in most cases false positives are worse. Thinking this issue through should help clarify your authentication targets.

Potential authentication metrics

The following possible metrics follow a sequence, a development process. The idea of discussing these particular metrics is not so much to say ‘these are the metrics you should use’, rather to help you think about metrics that might be available and suitable for your specific requirements.

Metric 1. Number of authentication failures

Our first metric aims to assess the size of the red areas on the diagram above – the hard part being how to measure them. How will we identify the false negatives and false positives? If we could identify them directly, of course, we could probably eliminate them! In practice, we can only estimate the numbers using the symptoms or consequences. False positives might be estimated from the number of incidents involving unauthorized access. False negatives might be estimated from the number of calls to the Help Desk by people requesting password changes, or the number of password failures recorded in the logs that are followed in short order by successful logins. Are you collecting the data already?

Straight numbers are the simplest metrics, easiest to measure and report. However, they don’t make much sense in isolation. If one month we have 245 authentication failures, is that a good or a bad month? We really need to know how that compares to other months (see metric 3 below).

Metric 2. Proportion of failed authentications

The simple numbers in metric 1 are only part of the issue. To determine whether the number of authentication failures was OK or unacceptable, management needs to know how many successful authentication events there were in order to calculate the proportion of failures. This might be counted or at least estimated from the number of first-time successful login entries shown in the system security logs.

Proportions take a little more calculation than simple numbers but generally make more sense. "27% of authentication events failed" sounds pretty bad in any context! Combine this with some clear targets defined or at least agreed with management and we are getting somewhere!

Metric 3. Authentication trends

Period-by-period trends are even better. Management can determine whether things are going in the right direction, and (if targets are shown) can see how much further we need to go before the issues are resolved. The reporting period is quite important: too frequent and a lot of time is spent measuring and presenting the numbers, too infrequent and the variations resulting from specific activities (such as installing a new password synchronization system, or delivering some security awareness training) tend to be lost.

Metric 4. Authentication confidence level

A rather different style of metric involves surveying people regarding their confidence in authentication, for example:

How confident are you that authentication meets the business needs? Please mark the following percentage scale at the appropriate point, in your opinion.

0% 50% 100%
|-----+-----|
Not at all. Not quite enough | Just about enough Absolutely!

Comments

It is a simple matter to measure percentage values from each response and calculate the mean score. Provided enough survey forms are completed (ideally more than 30), the results should be statistically valid. The comments can provide useful feedback and quotations for use in management reports and other awareness materials.

Conclusion

This paper has hopefully stimulated your thinking. Why not discuss some ideas with your management and seek their opinions? Good luck!

For more information

Please visit the NoticeBored links collection pages on [identity theft](#) and [authentication](#) for related web sites and more reading.