SECUILIBRIUM

Balancing Information Security and Business Needs

Dissecting Risk-Based Authentication

Or, Preventing Authentication Fraud

A Secuilibrium, LLC White Paper by David Ochel

This whitepaper takes a look at *risk-based authentication*. This term has been coined in recent years to refer to authentication schemes that evaluate fraud indicators found in the context of authentication requests (typically from remote users). If these indicators (such as IP-based geolocation) show that a user is attempting to log on under circumstances that differ from an average authentication attempt for that user, authentication can automatically be stepped up, i.e., additional measures are imposed on the user in order to confirm his or her identity.

We discuss the fraud indicators being used by risk-based authentication schemes and the authentication factors available to validate user identities, discerning between *hard* and *soft* authentication factors.

A threat model is then defined as part of a pseudo risk-assessment in order to assess how and when risk-based authentication may contribute to lowering an organization's exposure, and how it compares to traditional two-factor authentication.

Effective Date: 2013-12-09

Status: Released

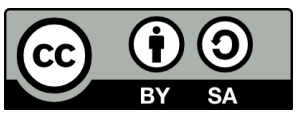

About Secuilibrium

Secuilibrium, LLC[1] provides information security consulting:

- **Strategy**

  Secuilibrium is your partner for risk and security management. We help you bootstrap an IT security governance program tailored to your organization's specific risks, or enhance existing frameworks with risk assessments, policies, and processes.

- **Architecture**

  We define or review your cyber security controls, architect security solutions that align with your primary business goals, and determine efficient and effective paths to compliance with government and industry standards.

- **Technology**

  Our broad understanding of technology, combined with our in-depth information security expertise, is available to you for expert opinion on new or existing solutions, support of audit efforts, and efforts to protect critical infrastructure.

We speak both English and German. Our services are available worldwide. Secuilibrium is based in Austin, Texas.

For feedback and inquiries, please contact us at:

Secuilibrium, LLC
3800 North Lamar Boulevard #730-314
Austin, Texas 78756
United States

Email: info@secuilibrium.com
Phone: +1-512-696-1404
Facsimile: +1-512-696-1502

[1] http://secuilibrium.com

Table of Contents

Index of Tables

Introduction

The term *risk-based authentication* can refer to (at least) two different scenarios:

Scenario 1: Asset Values Differ

Different systems within an organization hold information assets of different value. Higher-value assets attract attackers with more resources (time, tools, techniques, …) available to them, and at the same time compromise of those assets would constitute a higher loss for their owner. To counter the increased risk of an attacker stealing an authorized user's authentication credentials and gaining access to the data by impersonating that user, additional authentication steps might be introduced for those systems. (For example, two-factor authentication based not only on a user-known password, but also on a physical token held by the user.) This is a straightforward outcome of traditional risk management / threat modeling exercises.

This is **not** what is typically being referred to as risk-based authentication in today's vendors' literature.

Scenario 2: Convenience Rules!

One and the same system holds information assets of a certain value. Based on circumstantial evidence during a logon attempt to that system (for example, time of day or location of the remote user), the perceived likelihood of an attacker trying to impersonate an authorized user during that particular authentication attempt is gauged. If an irregular usage pattern is detected, a higher potential for fraud is assumed, and additional identity verification steps are enforced. If the logon attempt is within the user's usual parameters, (for example, originating from a previously used web browser,) the user is not being bothered with additional verification requests.

This is the typical functionality marketed as **risk-based authentication** these days and the subject of this whitepaper.

The reason for a surge in the marketing of risk-based authentication solutions is that mass-market, consumer-oriented organizations (banks, webmail providers, the gaming industry, etc.) have finally accepted that passwords are broken[2] and are trying to balance their need for better authentication mechanisms with the low user acceptance rates (and high cost) of complex authentication schemes. Basically, we are looking at an attempt to answer the question: "How can we maintain ease-of-use for most of our users most of the time, while decreasing the overall potential of many of our users' accounts being hacked by means of stolen passwords?"

The solution offered starts with burdening a user with requests for additional identity verification bits and pieces only if there is an increased likelihood that somebody else is trying to impersonate them.

[2] https://www.schneier.com/blog/archives/2013/06/a_really_good_a.html

That perceived likelihood is determined by looking at the circumstances (context) of the authentication attempt, and whether they are out of range of what is considered unsuspicious for that specific user or a specific user population. If the circumstances look suspicious, additional information factors are requested from the user in order to increase confidence in the user's identity.

To provide an example: If users log into an online banking solution from a PC that they have used before, and from an IP address in a region that they typically log in from, we might accept their used ID and password and let them go about their business. But if they log in from an unknown device, and/or from a country that's far away from their home address, we might start asking for additional information before granting them access. Things such as security questions we have on file, their address, social security number (hopefully not![3]), etc.

This is different from high assurance or even average corporate environments, where it must be assumed that attackers are sophisticated enough to mask the contextual parameters of authentication attempts successfully to render such a solution ineffective, and/or to obtain the additional verification factors necessary to impersonate a particular user.

Fraud Detection

In a way, the detection of fraudulent authentication attempts for IT systems is similar to credit card issuers' fraud detection systems – if an alarm bell goes off that indicates that a transaction might be fraudulent, the card issuer might send us a text message to confirm that we initiated that transaction. If more alarm bells go off, authorization of the transaction is flat-out declined.

In IT systems, a variety of parameters contribute to the context of an authentication attempt, in particular when we are looking at remote connections through public networks (i.e., the Internet).

Some of the parameters frequently used to profile authentication attempts, also referred to as fraud indicators, include:

- Originating IP/Network Address
    - Geographical region (Geo IP and related[4])
    - Blacklisted origin (e.g., known TOR relays[5])
    - Etc.
- Time (Time of Day, Day of Week, etc.)
- Logon Frequency
- Remote Device Trust
    - Pre-existing authentication token (e.g., a browser cookie)
    - System fingerprint (type of operating system, browser, etc.)
    - Etc.

[3] http://emergentchaos.com/archives/2012/01/shocking-news-of-the-day-social-security-numbers-suck.html
[4] http://en.wikipedia.org/wiki/Geotargeting
[5] https://www.torproject.org

- Number of Consecutive Failed Authentication Attempts

Based on the history of a user's logon attempts, a user- (or user group-) specific authentication profile can be established of values for these factors that represent what can be considered a *typical* or *unsuspicious* logon attempt.

Such authentication profiles for users and/or specific user groups (for example, all customer accounts located in the UK) can be pre-populated by an organization to some extent, for example with geographical parameters. Other aspects may have to be learned by the authentication system using automated mechanisms over time, such as the typical IP address ranges a user logs on from.

It is important to stress that none of these indicators constitute a reliable form of identification or authentication by themselves. Just because user David always logs on from a specific IP address, and nobody else has used this IP in the past to connect to our system, a connection attempt from that address does not validate (authenticate) David's identity. The IP address in a connection attempt could be spoofed, could have been re-assigned, or maybe a friend is visiting and using David's network to connect to the same system. As a result, this IP address can be used as a piece of circumstantial evidence in David's authentication profile to detect anomalies in David's future authentication attempts, but it cannot take the place of an authentication factor proving his identity.

Authentication

Now that we understand the factors that can be used to create authentication profiles (also referred to as risk profiles, fraud profiles, etc.), it is time to discuss the authentication factors that can be used to validate a user's identity.

By far the most popular and widely used type of authentication credential for human identities is the password, a secret kept between a particular user and a particular authentication system. It is easy to implement and does not require any additional hardware on the client side other than the computing resources already available. Unfortunately, it is also one of the weakest types of credential, as can be seen from the brief discussion of typical vulnerabilities associated with different authentication factors below.

The authentication factors used in risk-based authentication are not different from those used in other systems or schemes. They include, however, factors that do not meet the traditional definition of an authentication credential.

Hard Authentication Factors (Authentication Credentials)

Traditionally, authentication factors can be distinguished into three categories (Garfinkel & Spafford, 1992):

1. Something you know. (A password or secret.)
2. Something you have. (E.g., an authentication token.)
3. Something you are. (Biometrical properties, e.g. a fingerprint.)

Attacks to these factors are mostly restricted to vulnerabilities in the implementation of authentication systems, or to some extent vulnerabilities inherent to the properties of the underlying mechanisms:

- Passwords that are hard to guess are difficult to remember, and passwords that are hard to reverse from cryptographic hashes are even harder to define.[6] (They are also vulnerable to social engineering attacks and, unless one time passwords are employed, to keyboard logging malware.)
- Mechanisms to keep secrets confined to a physical token (such as, a smartcard or USB dongle) can be circumvented, and tokens can be lost or stolen.
- Fingerprints can be copied off users' whiskey glasses and presented to iPhone fingerprint readers for successful authentication[7], and authentication systems with reasonable false acceptance and rejection rates are hard to design and require sophisticated readers.

Regardless of vulnerabilities that may exist in implementation or underlying mechanism, in theory we can have high confidence that these traditional types of authentication factors validate a user's identity when presented as proof of an identity. We refer to them as hard authentication factors.

Note: Another, increasingly popular possession-based authentication token is a user's cell phone. Either an app on a smart phone (sometimes referred to as a soft token) can provide for an authentication factor that requires possession of the phone; or the phone's SIM card that is associated with the user's phone number in a mobile carrier's database enables the receipt of one time passwords per text message or phone call. The necessity to be in possession of the phone or SIM card makes this a "something you have" authentication token and thus arguably a hard authentication factor, although increased attack surfaces are at play here.

Soft Authentication Factors (Identity Verifiers)

On the other hand, risk-based authentication mechanisms introduce a whole new range of factors that do not fit well into these traditional categories. Many of those are actually well known to us from password reset mechanisms or telephone banking, and are also known as customer verification techniques (Federal Financial Institutions Examination Council, 2005). These factors either constitute personally identifiable information (PII), or technical solutions that assume a certain amount of trust into systems that are outside of an organization's control. They include:

- Security Questions and answers that users deposit with authentication providers.
- Personally Identifiable Information, i.e. information that can be used to identify a specific user. This includes things like home and work addresses, phone numbers, credit card numbers, which cars you have owned in the past, Social Security Numbers, etc. (Note that identification does not imply authentication!)
- Browser cookies, indicating that a user has previously used a particular browser to successfully authenticate themselves.
- Access to an external email account previously registered with the authentication provider.
- Etc.

[6] http://arstechnica.com/security/2013/05/how-crackers-make-minced-meat-out-of-your-passwords/
[7] http://www.ccc.de/en/updates/2013/ccc-breaks-apple-touchid

Vulnerabilities in mechanisms based on these factors vary in difficulty:

- Answers to security questions or knowledge of PII can often be gained by mining the Internet and other public or commercially available databases, including social networks.
- Cookies might be deposited in browsers that the user doesn't use exclusively, or might be stolen by malware.
- Email accounts may have been compromised or email delivery is intercepted.

In fact, none of these factors are authentication credentials in the traditional sense. Even security questions like "What did you name your first stuffed animal?" and "Where were you when you first heard about 9/11?" are hardly secrets that only the user itself knows the answer to. A cookie in a web browser is just evidence asserting that a user of that particular browser had been successfully authenticated using traditional means before, indicating that it may be likely that the same user is now using that browser again to access our systems. And validating that a user has access to an email account that is outside of our control just assumes that that mail server's authentication system has not been compromised – it is not a direct validation of a user's identity.

Depending on the type of authentication factor being used, we can have a somewhat increased confidence that we are actually talking to the user in question. (Although using PII to confirm a user's identity makes most security professionals' neck hair stand up.) We refer to them as soft authentication factors.

Risk-Based Authentication

Stepping Up Authentication

Dynamic, system-initiated response to elevated risk of fraud based on analysis of an authentication context is often referred to as *stepping up* authentication measures. Some examples include:

- If fraud indicators are low, the user is not authenticated at all if the requesting browser presents a cookie that indicated that the user had previously designated the system as "trusted".
- If fraud indicators are elevated, the user is asked to supply additional soft authentication factors, e.g. answers to pre-defined security questions.
- If fraud indicators are even more elevated, the user is asked to supply a second hard authentication factor, for example a one-time password sent by text message to the user's mobile device.
- If fraud indicators are unacceptable, the account is locked until the user goes through out-of-band authentication hoops. (For example, calling the helpdesk.)

There are more complex authentication scenarios in use as well. For example, asking a user for a second authentication factor only when they are trying to change their account settings, but not for normal account usage.

It is worth pointing out that many risk-based authentication solutions rely on soft authentication factors for stepping up authentication. This means that, rather than stepping up to what is commonly referred to as two-factor authentication (i.e., authenticating a user based on two of the three available hard

factors “know”, “have”, “are”), stepped up authentication takes the form of one hard factor (usually the user’s password) and one or a number of soft factors. The reasons for this have been mentioned before: it might be cost-prohibitive to distribute tokens or iris scanners to every user, and/or reduce the user experience of the system to unacceptable levels.

Obviously, the same fraud detection mechanisms could be used to trigger actions in our system that are transparent to the user. Rather than asking for additional authentication factors, a connection attempt could be elevated to operators that then closely monitor the resulting user session, etc.

System Configuration

Risk-based authentication solutions may have more or less sophisticated engines built in to automatically generate user profiles from a history of collected fraud indicator values, and to determine appropriate authentication responses from available options during an authentication attempt.

Obviously, the dynamic evaluation of fraud indicators and decision on resulting countermeasures for a specific logon attempt must be left to machines. However, it seems advisable for an organization to understand, and be able to control, the scheme that is being employed to make the decision on when and how to step up authentication measures, rather than just accepting vendor-recommended default settings.

The table below provides a simple example of defining parameters for fraud indicators that can then be used as decision points for stepping up authentication during an authentication attempt or not. For example, if no fraud is suspected, simple authentication proceeds. If fraud indicators show values that are suspiciously outside of the norm, we might step up authentication with additional means. And certain conditions might cause us to flat-out reject an authentication attempt.

Fraud Indicator	**Fraud Result**			**Justification**
	Unsuspicious	**Suspicious**	**Unacceptable**	
Origin IP	Consistent with user’s frequent geographical range.	Inconsistent with user’s frequent geographical range, or listed as Internet Café IP.	Blacklisted IP (TOR, known malware hosts, …)	We only have users in North America. Specific users are typically confined to specific regions.
Time	Logon during office hours or consistent with user’s frequent logon timeframes.	Logon during other hours.	None.	Users typically only log on during the day.
Remote Device Trust – browser cookie	Trusted cookie presented.	No cookie presented.	None.	Using the user’s browser for an attack would require sophisticated exploits.
Remote Device Trust – system fingerprint.	Consistent with user’s frequently used operating system type.	Inconsistent with user’s frequently user operating system type.	Blacklisted toolkit frequently used by attackers.	Our users use a limited number of personal devices to connect.
Consecutive Failed Authentication Attempts	< 3	3-5	> 5	Consistent with company policy. ;-)

Table 1: Evaluating Fraud Parameters – Example Scheme

Table 1 only deals with the fraud detection side of the solution. It basically assumes that there are only three options for authentication responses [no step up, step up, and rejection] that can be matched to

the fraud result values [Unsuspicious, Suspicious, and Unacceptable]. Depending on the capabilities of a solution, this can get more complex: different types of responses might be available for different types of fraud detection results, etc.

As discussed previously, a risk-based authentication system needs to be *trained* – authentication histories for users need to be collected to build reference profiles. And once implemented, initial settings (and definitions as provided in the example above) will likely have to be tuned further to achieve a system behavior that strikes the desired balance between reducing the likelihood of a successful attack and putting off too many users with increased authentication requests.

Measuring Risk

We understand now that risk-based authentication is a combination of fraud detection techniques based on circumstantial evidence (context) of an authentication attempt, and a dynamic response based on the results of that fraud detection analysis.

But how do we measure the *risk* in risk-based authentication?

When we are talking about risk-based authentication, this typically means that we need to establish a baseline (authentication profile) for the factors discussed above in order to determine whether a login attempt from a user, without collection of additional information, represents:

- Acceptable Risk – Parameters are in the range that we expect them to be in. Let the user in.
- Increased Risk – We'd like to ask additional questions to assert the user's identity before granting them access.
- Unacceptable Risk – We'll deny access to the user until they go through increased out-of-band scrutiny. (For example, calling them on the cell phone we have on file and asking questions about whether they really just tried to log in from a remote island on a computer they hadn't used before.)

Risk is measured based on some sort of metric. It might be a quantitative or qualitative approach, and it might be formally documented or just based on gut feeling. Table 1 above already provides input to this decision by suggesting a classification of risk for simple systems that is, basically, based on gut feeling.

It behooves an organization to determine the existing risks and the level of risk that is acceptable to the organization in a reproducible and documented fashion. In other words, a risk assessment needs to be performed.

An Example Risk Assessment

The following table shows an example risk assessment that might help an organization to decide which authentication measures are acceptable or favorable for certain systems, based on the assets they protect and the risk to these assets.

The intent of this example is to demonstrate what an organization's assessment of the particular risk associated with account impersonation based on stolen authentication credentials (let us assume passwords) might look like, depending on additional authentication measures that are enforced. We consider these scenarios:

- No authentication measures are enforced in addition to the compromised password.
- Two-factor authentication is enforced, with a second hard factor.
- Risk-based authentication, as introduced above, is enforced.

The table is followed by an explanation of its content, and then by an analysis of conclusions that can be drawn from it. Various assumptions have been made that would usually be determined in an organization-specific light, and an ad-hoc methodology is defined rather than using any of a number of fine standards available to define risk assessment methodologies, reducing our example to a pseudo-assessment tailored to illustrate the point of this white paper.

User Account Type	Asset Value	Attack Path	Existing Measures	User Acceptance	Attacker	Likelihood	Risk
Public Consumer	low	Stolen authentication credential used to impersonate user and steal user's information / assets (e.g., bank details).	none	high	unsophisticated	very high	high
					sophisticated	very high	high
			two-factor	low	unsophisticated	very low	low
					sophisticated	low	low
			risk-based	medium	unsophisticated	low	low
					sophisticated	low	low
	medium (aggregated, regulatory)	Stolen authentication credential used to impersonate user and steal user's information / assets (e.g., accounting details).	none	high	unsophisticated	very high	high
					sophisticated	very high	high
			two-factor	low	unsophisticated	very low	low
					sophisticated	low	low
			risk-based	medium	unsophisticated	low	low
					sophisticated	medium	medium
Corporate Employee	high (intellectual property)	Stolen authentication credential used to impersonate user and cause organizational havoc / steal IP / etc.	none	high	unsophisticated	very high	high
					sophisticated	very high	high
			two-factor	medium	unsophisticated	very low	low
					sophisticated	low	medium
			risk-based	high	unsophisticated	low	medium
					sophisticated	high	high

Table 2: Pseudo Risk-Assessment for Different Combinations of Authentication Measures

Explanation of Table 2:

- User Account Type: We are differentiating here between public consumers, i.e. private users who may have signed up for a service that an organization offers and – by design – have access to only the information concerning themselves; and corporate employees who are assumed to have access to organizational assets (intellectual property, aggregate data, etc.). In either case, we are assuming that the systems they access are facing public networks, i.e. they are available for remote access through the Internet.
- Asset Value: This is the value of the assets on our system that a user is authorized to access. Note that this represents the value to the organization, not for the individual user. (For example, an individual user's retirement money in a bank account is – in terms of organizational risk – likely not of extraordinary value to a bank.) However, the value of personal users' assets in our scenario can increase if we assume that compromise of the assets has additional financial consequences to the organization, for example regulatory fines or reputation loss. Lastly, in a

more complex scenario it might also be possible for an organization to value the accounts of certain users more than those of others (think bank accounts with significantly large amounts of money in them), further influencing the assessment outcome for certain user populations.

- Attack Path: In our example, the threat we are considering is essentially always the same. An attacker uses stolen authentication credentials (let's assume passwords) in order to impersonate a valid user and gain access to the data that the account in question is authorized to access. Note that we are not concerned with how the attacker obtained the password. How (malware, social engineering, etc.) is obviously relevant to our overall security posture and should be subject to a risk assessment, but it doesn't matter when we talk about risk-based authentication: Here, at the time of an authentication attempt, we are already past the point of the exploitation that lead to the knowledge of those credentials and it has turned into an assumption. (Or, maybe more accurately, we are really just looking at a tiny piece of a larger attack scenario.) We also assume that there is no vulnerability in the implementation of our system – the system accepts valid authentication credentials and grants access, and the risk that the system isn't implemented properly is yet another one to be considered somewhere else.
- Existing Measure: This takes into account whether we have existing countermeasures (risk treatments) against the threat already in place. Our example shows the risk related to no such measures (i.e., only the password mechanism that has already been compromised), a second factor for authentication that is always enforced (let's assume a hardware token that spits out time-based secrets), or risk-based authentication measures.
- User Acceptance: We are making assumptions here about the user acceptance of the additional authentication measures. We assume that corporate users are more accepting of an additional authentication burden about whose benefits to the organization they have been educated. In real life, these values could/should be qualified by some research rather than just assumed. Also note that this value is actually not relevant to our calculation of risk to our assets. (It could be, in a more sophisticated methodology.) It will help us in analyzing our overall results later, though.
- Attacker: We differentiate between two types of attackers. Unsophisticated attackers perform untargeted attacks, i.e. they might send phishing attempts to huge user populations or sniff passwords on infected PCs, and then use the obtained passwords to make some money off the assets of whoever ended up in their nets. Sophisticated attackers perform targeted attacks. They select their targets specifically based on the assets they are after, and have significant resources available to execute the attack. In particular, we assume that sophisticated attackers are able to spoof whatever conditions would trigger a risk-based authentication system, and that it would be difficult for them to obtain the hardware token used for two-factor authentication. (In real life, it might or might not be in order to differentiate further based on attacker capabilities.)
- Likelihood: This is our perceived likelihood of an attempted attack to actually succeed. It takes into account the sophistication of the attacker and the effectiveness of existing countermeasures. We are using gut feeling here, or maybe we surveyed our vulnerability specialists' best guess. In real life, this could be qualified further by various means, for example an in-detail assessment of how easy it would be for an attacker to spoof the individual fraud indicators used by our risk-based authentication solution. Another thing that is implied in our

example, rather than well defined, are the potential vulnerabilities in the Existing Countermeasures: In order for an attacker to be successful and circumvent the additional measures, they need to be able to exploit a vulnerability in our system. (Besides implementation or logical errors in our countermeasures, examples include the (temporary?) extraction of a physical authentication token from a user and the ability to spoof input for our risk-based authentication algorithm.) Note that this value is somewhat biased by the likelihood that a specific type of attacker will actually attempt the attack in the first place: While they would be likely to succeed, sophisticated attackers may be unlikely to waste their costly resources on low-value targets.

- Risk. For our particular example, we decided to assign risk levels based on asset values and likelihood of successful compromise according to the following table:

		likelihood				
		very low	low	medium	high	very high
asset value	low	**low**	**low**	**low**	**medium**	**high**
	medium	**low**	**low**	**medium**	**high**	**high**
	high	**low**	**medium**	**high**	**high**	**high**

Table 3: Calculation of Risk

- Note: This whole methodology is somewhat incomplete. We never define what a low, medium, etc. value means to us (for example, this could be dollar-value ranges for asset values, and frequency estimates for the likelihood). Certain details have been omitted or simplified. These details, however, are essential pieces of what makes risk different for different organizations, and would have to be defined well in an organization-specific risk assessment. (Ideally, we would be able to tell our board how much money to budget to deal with successful attacks based on the risk we calculated.)

Assessment Analysis

In the introduction to this white paper, we have proposed that risk-based authentication is primarily a substitution of more rigorous authentication measures in favor of usability, or user acceptance: We don't want to burden users with more of an authentication hassle than we need to. If the risk of impersonation is low, no questions will be asked. Let's not annoy our customers and employees more than necessary. Secondarily, time and resources might be of consideration as well. More authentication factors require more first-level support, for example. If it wasn't for those considerations, why not issue a hardware dongle to every user anyway and make them go through mandatory two-factor authentication every time?

The pseudo risk-assessment above qualifies these thoughts by actually demonstrating the risk for our imaginary example organization as it pertains to different authentication solutions, user populations, and asset values. The addition of a User Acceptance column takes into account additional considerations that must be made in the larger context of managing IT and business risk. Technology with low acceptance rates generally prohibits business.

Some particular conclusions we can draw from the assessment above include:

- Let's assume that our organization is generally averse to "high" risk, however it may be defined in the specific organization's context:

  - In our example, this obviously means that risk-based authentication for corporate users is not acceptable, since the risk of sophisticated attackers being willing and able to circumvent that measure and gain access to our high-value data would be "high".
  - It also means that – if we consider public user data to have medium rather than low value – we need one of the two additional countermeasures (two-factor or risk-based authentication) to support our password-based authentication of public consumers.

- It appears that risk-based authentication for public consumer accounts that hold assets of "medium" value might be acceptable to us, constituting a "medium" risk that we are willing to accept. From an IT administrator's perspective, two-factor authentication with a second hard factor (the hardware token in our example) might be more ideal. But there are other considerations to be made. Issuing hardware tokens to thousands of users that produce low revenue might be cost-prohibitive. And the fact that the user acceptance for those tokens is low might actually mean that the overall business risk of losing customers over it might be too high.
- While already stated as an assumption above, it is worth re-iterating that sufficiently motivated attackers are likely to both have the means to spoof fraud indicators in an authentication attempt to the extent that the attempt does not look suspicious to our systems, and to obtain soft authentication factors that might be employed to step up authentication if an attempt is actually considered suspicious. (If they would actually go through the expense of doing this to obtain low-value assets is another question.)

Qualifying Risk-Based Authentication Measures

Another thing that remains unspecified in our pseudo-risk assessment above is how we came to the conclusion that risk-based authentication as a countermeasure would actually reduce our likelihood of exploitation by a certain amount.

What we are really looking at/for in risk-based authentication are anomalies when compared to a particular user's or user community's typical authentication behavior. How does this translate into risk?

One way to look at this is as a dynamic qualification of our risk assessment above for individual logon attempts. Our risk assessment could be augmented to consider the fraud detection results of our solution when determining the overall likelihood that a specific logon attempt might actually be a potentially successful attack under way.

The following table is derived from the three risks listed in Table 2 above for public consumers, low value assets, and an unsophisticated attacker. It takes the risk-based authentication measure and augments it into three separate risk considerations, based on the result of the fraud detection as defined in Table 1:

User Account Type	Asset Value	Attack Path	Measure	Attacker	Fraud Result	Likelihood	Risk
public consumer	low	impersonation	none	unsophisticated	not available	high	medium
			two-factor		not available	very low	low
			risk-based authentication		unsuspicious	low	low
					suspicious	high	medium
					unacceptable	very high	high

Table 4: Estimating Risk-Based Authentication Risk :-)

The dynamic response can then be configured to be equivalent to our risk appetite, i.e., if the risk is high and thus unacceptable, an authentication attempt will be rejected; if it is medium, additional authentication mechanisms will be engaged; and if the risk of impersonation is low, no additional authentication is required.

More fine-tuning is obviously possible here, depending on the employed solution. It may also be advisable to consider risk for individual fraud indicators separately or in combination, depending on how effective they are perceived to be. The following table illustrates this by just looking at the "time of day" fraud indicator. In this example, it is assumed that attackers can easily discern the typical time of day for a user logging in, and hence the risk of impersonation to our particular organization is still considered "medium" even if the authentication attempt happens in an unsuspicious context, as far as our fraud detection mechanisms are concerned. (Practically, this means that we might as well disable fraud detection based on the time-of-day indicator in our system, since we do not consider it a reliable indicator even for unsophisticated attacks.)

User Account Type	Asset Value	Attack Path	Measure	Attacker	Fraud Result	Likelihood	Risk*
public consumer	low	impersonation	risk-based authentication: time-of-day	unsophisticated	unsuspicious	medium	medium
					suspicious	high	medium
					unacceptable	very high	high

Table 5: Differentiating Between Individual Fraud Indicators when Evaluating Risk

Disqualifying Risk-Based Authentication Measures for Sophisticated Attacks

In the interest of qualifying our statement that risk-based authentication is inappropriate for high-value assets, and/or attackers that invest sufficient resources in subverting them, the following table provides some consideration on how sophisticated attackers can spoof or otherwise re-create the fraud indicators introduced earlier.

Fraud Indicator	Vulnerable To			
IP Origin	Targeted infection of user's personal device allows to directly observe user behavior, access authentication cookies, originate requests that actually come from the user's system, etc.[8, 9]	Active man-in-the middle attack allows attacker to trick user into validating new IP address.	Spoofing of individual addresses.[10] (Admittedly very hard.)	Attacker obtains address in the same subnet / IP address range as user.
Time of Day		Passive man-in-the-middle attacks allow attacker to establish a user's normal logon patterns.	Guessing may achieve reasonable results.	
Logon Frequency				
Pre-Established Browser Cookie		Cryptographic weaknesses, timing based attacks[11], etc.	Man-in-the-browser attack[12], Cross-Site Request Forgery[13], etc.	
System Fingerprint		Use same system stack as user, or emulate user's system fingerprint.		

Table 6: Circumventing Fraud Indicator-Based Detection

[8] https://www.sans.org/reading-room/whitepapers/malicious/bypassing-malware-defenses-33378 [PDF]
[9] http://searchsecurity.techtarget.com/feature/Antivirus-evasion-techniques-show-ease-in-avoiding-antivirus-detection
[10] http://www.symantec.com/connect/articles/ip-spoofing-introduction
[11] http://arstechnica.com/security/2013/03/new-attacks-on-ssl-decrypt-authentication-cookies/
[12] https://owasp.org/index.php/Man-in-the-browser_attack
[13] https://owasp.org/index.php/XSRF

In Table 6 above, it is assumed that attackers are able to figure out which combination of indicators a solution works with – security by obscurity is not a valuable information security paradigm (Kerckhoffs, 1883). No attempt to provide an exhaustive list is attempted.

Should sophisticated attackers trigger an authentication step-up when trying to impersonate a user, they may also be prepared to provide additional authentication factors. Obtaining soft authentication factors based on PII appears trivial for resourceful attackers. Out-of-band emails as a soft authentication factor may be harder to intercept, but attacker who are able to obtain a user's password to our system are likely able to obtain the user's password for their email account as well (if it isn't the same one in the first place).

Conclusions

By providing an actual (pseudo) risk assessment for the attack scenario that risk-based authentication mechanisms attempt to address, we have shown that those mechanisms are in fact able to detect conditions that might indicate increased risk of an authorized user being impersonated by an attacker, and to implement a corresponding risk treatment by potentially stepping up the amount of identity verification mechanisms for a particular authentication attempt. This active, dynamic response can be likened to mechanisms in a network-based intrusion prevention system that automatically alters router or firewall configurations if certain network-based attack patterns (for example, a denial of service attack) are detected.

Risk-based authentication works well in particular contexts, i.e. for mass-market-facing systems that are frequently exploited by untargeted attacks. The primary goal is to keep user acceptance rates high, which is achieved by prompting users for (additional) authentication factors only when an increased chance of impersonation is perceived, while lowering the impact of frequent, automated, and primarily less sophisticated attacks at reasonable expense.

As has (hopefully) become obvious as well, risk-based authentication solutions are not appropriate for scenarios where targeted attacks against high-value targets may exist. This probably includes almost all scenarios in corporate environments, where intellectual property, trade secrets, and/or accumulated data may be at risk. If attackers are determined and resourceful enough, it has to be assumed that they can reproduce or emulate the factors that risk-based authentication systems are looking for in order to determine whether "everything looks normal". Organizations are well advised to enforce increased authentication measures based on two strong authentication factors (two-factor authentication) for all authentication attempts on high-value systems. Acceptance of increased authentication measures in user populations that are governed by corporate goals can be raised by education, and by implementing solutions that function as seamless as possible even when requiring a second factor for authentication every time.

In the context of risk management, which needs to deal with many more scenarios of unauthorized access and involved authentication mechanisms, the term *risk-based* for a specific authentication solution that attempts to address a specific threat scenario appears somewhat unfortunate. It might suggest to some that by employing a risk-based authentication solution, all your authentication-related risks are taken care of. Risk is different for every organization, though, and an organization's risk management needs to be informed by organization-specific risk assessments. An example for a more descriptive term for the solution discussed might be *authentication-fraud prevention*.

Bibliography

Federal Financial Institutions Examination Council. (2005). *Authentication in an Internet Banking Environment.* Retrieved 12 06, 2013, from FFIEC: http://www.ffiec.gov/press/pr101205.htm

Garfinkel, S., & Spafford, G. (1992). *Practical UNIX & Internet Security* (Second Edition ed.). Sebastopol, CA: O'Reilly & Associates, Inc.

Kerckhoffs, A. (1883). La cryptographie militaire. *Journal des sciences militaires , IX*, 5-38.

Revision History

Date	Author	Changes	Resulting Status
2013-12-06	David Ochel	Initial Version	Draft
2013-12-09	David Ochel	Addressed review comments.	Released

