Sexy Defense

Maximizing the home-field advantage

Iftach Ian Amit
July, 2012

Cover Image Credit: IDF Spokesperson

Table of Contents

Ian Amit

Sexy Defense

Maximizing the home-field advantage

Abstract

Offensive talks are easy, I know. But the goal of offensive security at the end of the day is to make us better defenders, and our defense stronger. And that's hard.

Usually after the penetration testers (or worse - red team) leaves, there's a whole lot of mess of vulnerabilities, exposures, threats, risks and wounded egos. Now comes the money time - can you fix this so your security posture will actually be better the next time these guys come around?

Another example would be a failed audit – you are now looking not only at fixing things that do not make sense to you as a business (or a technology organization) but also at fines and potential legal actions that need to be addressed.

This paper focuses mainly on what should be done (note - not what should be BOUGHT - you probably have most of what you need already in place and you just don't know it yet). The focus here is to be able to play to your greatest advantage that a defender has – the home-field advantage. Most organizations fail to realize this advantage, and fail to comprehend how to use it. Knowing your network, organization and processes will help you in both fending off attacks, as well as in containing them. Knowing your enemy will take you further and allow you to stop attacks before they even start.

Methodically, defensively, decisively. Just like the red-team can play ball cross-court, so should you!

Background

A red-team test is a full-scope engagement that simulates a real-world attacker materializing a threat on an organization. Such a test is very different than a traditional penetration test, and as such, if you thought you were ready after "passing" previous penetration tests, you are probably surprised by the findings that the red-team found: full compromise, physical intrusion, stolen intellectual property, and pivoting through different elements (technical and social) inside your organization.

If that is the case – then you are in luck, as someone took the time to map out your real vulnerabilities (and not just the ones that show up on an automated scan on some random part of your network). This should be your starting point, and a wakeup call to start handling your defensive strategy a little better.

If on the other hand you have only had some basic penetration tests run on your organization, you have a lot of additional work ahead of you, but do not fret, as there are actually some useful elements you can dig out of that 2-inch report.

Interpreting bad penetration test reports (or: working with tainted information)

Reading between the lines of most badly written penetration test reports requires some creativity, and an understanding of what your organization really cares about. On the same note – information doesn't necessarily come from penetration test reports, and can show up in the form of just misinformation, or even good information in the wrong context or correlation. For example – audit reports, capability assessment reports, performance evaluation reports, etc…

However, this is an opportunity to make sure that you DO know what you are dealing with in terms of your patch management and updates for the relevant systems.

A couple of things to note about reading badly written reports:

a. Don't even try to figure out the "severity" of the findings from the report. More often than not, the tester lacks the business understanding, and as such, their analyzed severity will not necessarily reflect your actual

situation. You have to do your own homework and figure out how the exposure relates to your business.

b. Try to "collapse" multiple issues that can be resolved in a single action into one issue. This is a common practice when the tester uses a scanning tool and copy-pastes the report into a report template (get more findings/pages).

How to identify a good report – look for the money. If the report provides a business impact analysis (and if it's in $ even better) you know that someone was actually looking at your organization and was trying to figure out what would hurt it. Pay attention to the vectors used to get to the assets, and map out additional relevant vectors that you know of that may have the same issues as the ones portrayed in the report!

Now off to some of the terminology you will encounter in reports and how to interpret it:

Vulnerability

This is what the report will usually detail. Lots of vulnerabilities – usually associated with specific software versions of products used across your IT infrastructure.

A vulnerability is an issue with a software component that, when abused (exploited) can lead to anything from the software crashing, to compromising the system on which the software is installed so that the attacker can have full control over it. Additionally, vulnerabilities also refer to logic and operational issues – whether in computing systems, in processes and procedures related to the business operations, patch management, or even password policies.

Exposure

Usually you'll see references to exposures in more methodical reports where vulnerabilities would be detailed in a more technical part of the report. Exposures would relate to some threat model that has been created to represent the kinds of threats the organization would face.

Threat

A threat can generally be defined as anything that is capable of acting against an asset in a manner that can result in harm[1].

A threat would then need to be broken down into its elements - a threat agent or a threat community, their capabilities, and accessibility to the assets in question.

If the only reference to the threat is a generalized one, you know you are dealing with a badly written report. Look for the right terminology (otherwise referenced to as "threat modeling") when dealing with reports in order to get the most value out of them.

Risk

This is what you are looking for. Most reports won't have it. Here, risk is expressed in more mathematical terms and is usually the right to tool to use when discussing security with management (and/or budgeting/planning security for you organization).

Risk is the probability of something bad happening to your organization's assets. It should be expressed in some form of potential impact that represents the loss that would be incurred from such an event. In order to calculate a risk, all the elements that compose a risk should be expressed coherently - the exposures, the threats and their relevant components, the likelihood of the threats materializing and using the exposures while bypassing the controls, and finally the potential impact on the organization. There are multiple frameworks that deal with how to represent risk. Find one that works for you and try to make sure that everyone "speaks" in that language.

Methodology

In order to turn the table on the information security practice (i.e. attackers have all the initiative and information, defenders are left to react to actions from

[1] Based on the definition of a threat from the FAIR methodology (see bibliography).

attackers and patch things up as they come along) we need to get some homework done before we start practicing proactive defense.

First things first – the methodology here should not be that different than the one used by attackers. The same work that goes into the preparation of a well-planned attack should also take place at the defending side.

Second – the notion of things being "not fair" should be thrown out the window. Just as attackers have no scope limitations, so does the defensive methodology. Anything that limits a defensive strategy will be used as an attack vector once identified by the threat community (and expect it to be identified).

Third – it's not about technology or fancy products. It's about using what you have (and you probably have more information that you are using right now for your defensive strategy). It's about taking the initiative and making sure that you are using everything at your disposal to be proactive, expect attacks, identify them and focus your defensive means on lowering your assumed risk. The scope of the defensive strategy should extend as far as possible to match the scope of the attacker. This is exactly where "traditional" defense fails – by limiting the scope to exploitation or post-exploitation. This in turn limits the budget, the skillsets sought after to man the positions, and the kind of communications that are being sent back to management.

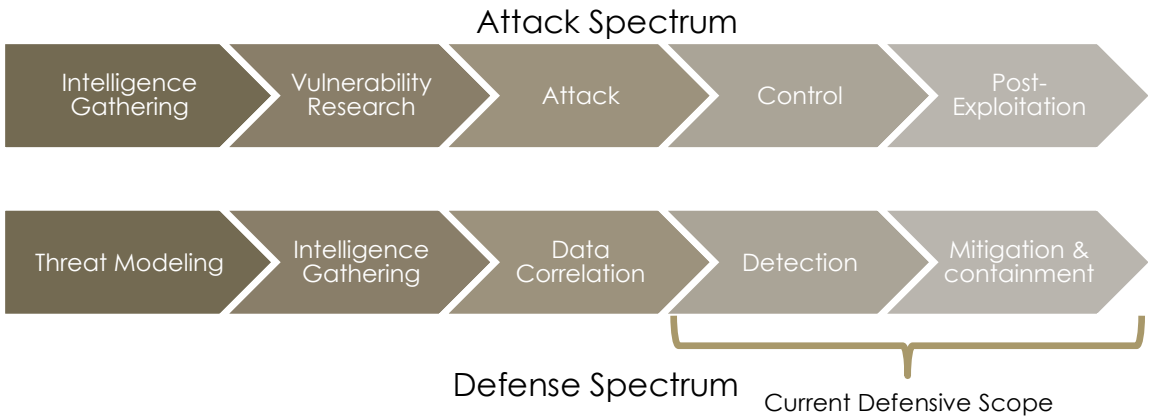

Figure 1: Defensive Security vs. Offensive Security Scope

One thing to remember though – is that technology should be used effectively and in context. The context is that attackers will breach some defenses (i.e. does

your strategy/network takes into account that some of your assets are compromised? Can you still operate in such a situation?). Effectively means that there is no one solution that has the same effectiveness to different organizations. Your task is to figure out what works best for YOUR scenario, and not who else bought a particular technology.

And finally – it's not about egos, not about people, and not about skills. Defenders have to assume the mindset of constant improvement. This also means that they need to understand that there will almost always be gaps in their defenses (hence the assumption that breaches will occur, and they have to design their environment in a resilient way). True strategic defense has to be able to identify gaps (usually by using an outsider's view such as an auditor, a penetration tester, or even a colleague from another department), and remediate it in the context of risk management. This doesn't reflect about the skills of the security personnel one bit. Not doing so and avoiding the criticism does.

Mapping your information assets AND your security assets

Mapping out the assets which you are actually protecting is the first step in creating a defensive strategy. Asset mapping should be derived from the business. What is the business doing, how does it operate and make money, what are the core and supporting processes, and what are the assets related to such processes. Additionally, all relevant elements involved with such processes should be mapped as well in order to identify not only the technologies involved, but also the people, 3rd parties and procedures.

One thing about asset mapping is that more often than not, the human assets are forgotten. While businesses are there to create products or protect money, without the people running these processes no products would be made, and no money could be saved. The human assets need to be factored into the equation – especially where critical processes exist, and where such human assets could reflect on the integrity of other information assets.

Additionally, your security and intelligence assets need to be identified in order to be able to make use of them in the greater scope of a defensive strategy (i.e. beyond using them simply their basic/core features).

Mapping your actual exposures and issues.

You can start off based on a penetration test report if you have one available (good or bad) to map your current technological issues. If a penetration test report is not to be found, even a simple vulnerability and exposure mapping using some automated tool could be used as a head start on the mapping process.

This however does not mean that it would have anything to do with the business you are protecting. In order to address all the issues and exposures, an equivalent of the view from the perspective of a red-team engagement should be sought after (which would focus on getting to the critical assets). In cases where a red-team test was never performed, a dry-run (tabletop simulation) of a full-scope attack can be run in order to map the missing components (usually the less technologically related elements).

Once both the asset mapping is achieved (and updated to reflect any changes in them), as well as the vulnerability mapping discussed above, they can be "overlaid" to get a coherent picture of the defense field.

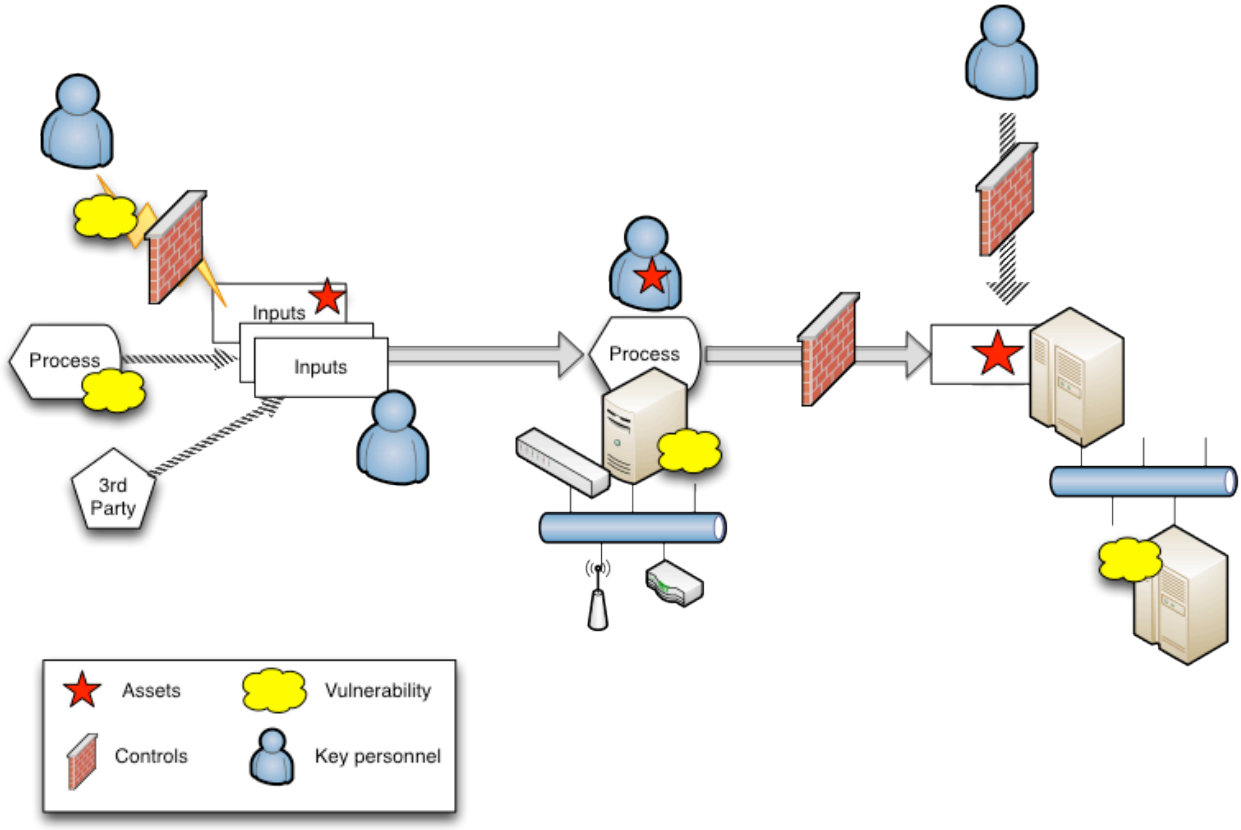

Figure 2: Simplified mapping of assets, processes, technologies and vulnerabilities

Logs

Logs are generated from ALL of your equipment. Workstations, servers, routers, switches, applications, turnstiles, websites, vending machines, cameras, sign-in systems, etc.… Remember that some of these would show up in physical format rather than electronic (yes – papers… scan, OCR, and feed back to your log management).

Usually, these logs are treated in a narrow context of assurance that the generating entity is functioning well, and track it's performance or issues in that specific context. However, as we are trying to get a more holistic view of security in an organization, it is imperative to make sure that we get access to ALL the logs (even those that were deemed irrelevant before). This is the exact equivalent of intelligence gathering at the attacker side, at which ALL information is collected, and then correlated to figure out its relevancy.

Raw intelligence

There's nothing wrong with gathering intelligence on your threat agents/communities. It's the same practice that attackers engage in, as well as large companies – especially on the competitive and financial sides. However, for some reasons the information security element of the defense strategy is left unattended too often.

Raw intelligence can come from many places – either actively looking for signs in forums where you know that perpetrators would hang out (or use for their own intelligence gathering and capability building), or even from 3rd parties that specialize in gathering such intelligence.

Geopolitical, financial, technical, recreational – anything relevant to the threat landscape that has been identified during the mapping phases should be collected and archived for analysis. This is what tips the scale slightly back towards a more balanced playing field, and enables us as defenders to know that an attack is coming even the first probes are sent in. This section may sound a bit like military operations, because it reflects the same tactics used to protect military assets. Nevertheless, there's nothing wrong with ethically and diligently collecting intelligence related to your business operation as long as it does not violate any relevant laws or privacy regulations for your employees.

Early warning signs

Sometimes, logs and alerts show up too late in the game, while signs of a compromise, or of targeting a specific element of your organization might still be visible beforehand. Increased call volume in the internal call-center, increased incidents of mis-behaving PCs, and other commonly overlooked events that are not classified as security incidents should be viewed in a broader context and married back to the defensive intelligence landscape.

Additionally, standard non-security "IT" events can also show early signs of a targeted attack, or an ongoing one. Permissions on files, general storage activity on shares and NAS devices, network activity around specific segments, etc.… These by themselves usually do not raise the classic security alert, but when combined with other activities, will show a different picture that should be analyzed more closely by the security teams.

People

One of the critical elements of defensive security is the same critical element of offensive security – people. Ask any security practitioner and he/she will tell you that they could secure a network airtight if it had no users on it. Also – ask them how many of the technical threats could not have started without a malicious insider, a rogue employee, or an employee clicking on a link or an attachment in an email.

However, the same element that opens up so many attack vectors into an organization is also one of the more attentive, which can provide an element of early warning of suspicious activity before any of the security products installed on the network.

Starting from awareness where employees notify regarding stalkers, tailgaters, new people joining the smoker's circle, strangely behaving applications, unknown devices in the office, and the list goes on. Any input such as this could be critical to realizing that the organization is either under attack, or is being surveyed or profiled before an attack takes place.

Correlating

The actual practice of defensive security starts when one has all the required feeds as discussed above (mapping of assets, exposures, and issues, logs, raw intelligence, general warning signs, and human intelligence). At this point, a correlation process is initiated that weights, and cross-references the different elements of the feeds in order to get a clear picture and context for those events. One important guideline for this process is to not throw away the raw data from the different feeds, and to keep the correlation information handy. As a correlated event may not have a context when it is being analyzed, it may be later added to another event/correlation to provide additional meaning and insight.

Additionally, on top of the raw feeds and correlations of events, external information should also be added to the timeline of correlations. Usually, attacks are aligned to use environmental elements such as holidays, sporting events, industry-related events, geo-political events, etc.… The context of a singular

event or a correlation of events may be crucial based on the timeline in which they occur.

From theory to practice

Implementing a proper defensive security strategy is not a one-time process where at the end of it a finite state of security (i.e. risk management equilibrium) is achieved. It is an ongoing process that involves constantly learning both the organizational elements of the business such as personnel, processes, finances, technology and business development, as well as the intelligence study of the threats and capabilities of the adversaries to the organization.

When creating a defensive security strategy, several important points must be taken into account:

- Assessment of current status - especially for awareness and security controls. In order to correctly address issues with the human factor of the equation, it's critical to get an accurate understanding of how well each department and individual are informed and trained on security issues. It will be a challenge as different departments in the organization will have different capabilities from a defensive posture, and others will need to be compensated with controls and additional monitoring. On the same note, security controls do not offer perfect capabilities, and should be constantly assessed against the threats and their capabilities to detect or mitigate such attacks.
- Constant development. Expect changes all across the process. Not just with the technological elements that are in play, but also with the organizational ones (processes being added, amended, or replaced), and business related changes (3rd parties, suppliers, partners, business lines, internal departments, etc.…). The defensive strategy is not a finite document, but a living document that reflects the current security stature of the organization and the combination of strategy and tactics for securing it. As the state of the threats against your organization will constantly change, so too should your defensive security strategy.
- Always align outwards. Measure the security strategy against peers in the industry, against foreign companies in the industry, and against the latest developments in offensive security that apply to your organization

(probably most of it!). Keeping the scope of the strategy to a local regulation or a past threat that materialized is a recipe for failure.

- Achieving strong defensive security is not about a specific tool or a specific person with certain security skills. It is about the combination of all the moving pieces in the organization, and how they can be utilized to manage the risk associated with doing business. It's about expanding outwards and not "hogging" the data, sharing it, learning from other people's mistakes, and sharing responsibilities across the organization. Management buy-in can lead to a bigger security budget. Peer buy-in can lead to actual better security.

Counterintelligence

As part of the outward reach to collect intelligence on potential threat communities and capabilities, the defensive security strategy should also employ counterintelligence efforts. From the basic honeypots to more sophisticated information manipulation (sometime referred to as information warfare), counterintelligence will allow a defensive strategy to take a proactive stance and expand its scope beyond the reactive "detect and mitigate" one.

Counter-Intelligence is also an opportunity to turn the tables on the attacker-defender relationships, and assume a more aggressive approach when attackers are identified. This approach yields many legal and ethical considerations that should be taken into account, but when constructed properly can be used for anything from proactively obtaining more information on the attacker, to actually counter-attacking[2] as part of the mitigation process. Many opportunities are present in this field of information security – especially when coupled with smart implementations of honeypots[3] and correlated detection data.

There are many ways to perform counterintelligence and this paper cannot assume to cover them in this context, especially as such practices should be

---

[2] Note again to consult with legal regarding the legitimacy of counter-attacking or the extent in which a counter-reaction can be taken against an identified attacker.

[3] Honeypots here is used as a general term for systems, services, networks and data used as a trap.

highly customized to the kind of environment in which the organization operates in – politically, financially, technologically, legislative and geographical.

Training people to identify, report, react

As mentioned before in the "theory to practice" section – training and awareness of employees is key to a successful defensive security strategy. Having the capability of your own employees to identify abnormal behaviors, report on them and sometimes even proactively react (without waiting for support/IT/security to address the issues) provide a great leap in terms of defensive capabilities.

There is no replacement for the human factor in the defensive strategy. It is the cheapest and most "fuzzy logic" solution that can be brought into the security field. Tools and automation can be used to minimize the grunt work and to bring interesting aspects of the data to the table – but people are still needed to bring in the "ah-ha!" factor, which can be then fed back to an algorithm that will learn how to automate that for the next time. It may be a pattern of some sort, an anomaly (or lack thereof), or anything that can be deduced from the data brought in.

Combining technology into the mix

As portrayed so far – the technological elements are one of the last to be added to the defensive strategy. This is because technology by itself cannot solve issues, and in most cases cannot even provide a solid event feed without properly running a mapping process of assets. Technological solutions should be first recognized for their actual added value, and in which fields, before being added into the defensive mesh.

One notable example are "security solutions" that provide a very narrowly scoped protection of a certain medium – be it firewalls for networks, WAFs for web applications, or Anti-Viruses for known samples of malware on PCs. Such solutions cannot be considered as baseline elements of a security strategy, but only as temporary add-ons to a comprehensive layered security strategy that *includes* technology as part of its overall scope.

The same goes for log correlation, SIEM/SOC products – these are fantastic to work with, once a strategy has been designed, and all the elements are working

in sync. Simply acquiring a product from this category, and having an integration process (as long as it gets) does not provide security, nor intelligence. The age-old principal of "garbage-in, garbage-out" applies here more than ever, and the inputs to such systems should be comprehensive, and should constantly be tuned to reflect changes in the threat model and the organizational security posture.

Working with others

So far this paper discussed defensive security in the context of an organization, but one of the greatest benefits of defenders is that there are a lot more like them. Having common threat communities with other organizations means that one defenders work (especially on the threat modeling and intelligence analysis) can benefit other defenders. Working with peers to share attack vector history, recon attempts, detection information, forensics, and other elements that would be used to create a better understanding of the threat and its capabilities can only mean better defense as a whole.

Additionally, there are organizations that can also assist in sharing information (securely and anonymously of course) such as CERTs, vendors and other companies that are exposed to either offensive or defensive capabilities and the relevant information that surrounds them.

Data sharing has been a long standing issue within some communities, but examples such as FIRST (Forum of Incident Response and Security Teams), and other trust-based models show that the benefits outweigh the efforts that are put into the sharing process (anonymization and sanitization of data mostly). Also – several governments offer public-private partnership approaches that allow bidirectional sharing of data, which on a strategic level could also add a layer of information to a corporate defensive security strategy on both the threat landscape, as well as the response and detection capabilities.

Conclusions

Defensive security has long been pigeonholed into a reactive mentality by a combination of products, and an approach that was led mainly by an IT-centric leadership. At the same time, offensive security has gained the spotlight as it kept a large gap ahead of the defenses, while combining techniques and strategies well beyond the traditional IT security realm.

Additionally, most defensive methodologies (again – driven usually by product vendors) focus on the reactive element of the spectrum, and have recently reached the point of assuming compromise and exploitation, thus providing solutions for that end of the spectrum. Most notable of these are the anti-virus, intrusion detection and prevention, firewalls, network access controls, and other reactive solutions.

Figure 3 - the whole is greater than the sum of its elements

Nevertheless, there is still hope – turning the tables on the reactive security approaches, although viewed sometimes as radical, provides an effective mitigation strategy that covers a wider spectrum of the threat. This allows organizations to implement a much more effective risk management operation that is based on more informed decision-making rather than knee-jerk reactions to perceived attacks.

Looking forward

When looking forward at such an approach to defensive security, it is easy to find many opportunities for improvement – both in the methodology presented here, as well as in providing products and services to fill in the voids (especially on the analytical side, intelligence gathering, counter-intelligence, and general integration of the different elements proposed herein). This is a clear call-for-

action to vendors to start providing such solutions, and probably more importantly drive a paradigm change in the way they see security products that should play a more proactive role and cover additional areas of the defense spectrum.

And finally – testing your security is not simply about finding new holes or gaps in the strategy, it is also (and probably mostly) about being able to test for preparedness. When testing shies away from the compliance motivation it can be leveraged to a point where organizations can place themselves in the heat of the battle without having anything at stake to actually lose. This is a priceless opportunity to see how they would fare against an actual attack rather than a theoretical one, and examine not only the technical elements of the security strategy, but also its actual execution, the processes and how the organization behaves.

Bibliography

Assadorian, P; Strand, J. (June, 2011). Bringing Sexy Back (Source Boston 2011, slideshare: http://www.slideshare.net/SOURCEConference/paul-asadoorian-bringing-sexy-back).

Wisner, Frank G.. (22 September 1993). On the craft of Intelligence (CIA Historical Review Program) https://www.cia.gov/library/center-for-the-study-of-intelligence/kent-csi/vol8no1/html/v08i1a07p_0001.htm

Jones, Jack (2005) an Introduction to Factor Analysis of Information Risk (FAIR) http://www.riskmanagementinsight.com/media/docs/FAIR_introduction.pdf