# Measuring DDoS Risk with FAIR

Not too long ago, I had the pleasure of presenting at the F.A.I.R. Institute in __. It was a great pleasure to have presented at this conference, with such amazing conference organizers and generous hosts. I genuinely enjoyed being a presenter. During this conference, I presented a case study on measuring distributed denial of service risk with FAIR (Factor Analysis of Information Risk). I used techniques that had been developed and refined over the years by Jack Jones and other FAIR advocates, as well as other techniques from Doug Hubbard and various others I incorporated into my analyses. I started my presentation on a brief summary of myself and my qualifications as well as how I got into FAIR and quantitative analysis. I felt this explanation would further explain a great deal about the approach that I use.

I then proceeded to the agenda and the objectives for my presentation. The objective was to give a hands-on look at how we can measure DDoS for a typical retail bank. This was going to be structured by creating a fictitious bank that competes with the top ten banks in the United States. I would then walk through each step in FAIR, describe what the step is, the result, and how exactly I got it. It was very detailed and hands on, and my goal was that everyone at the conference would be able to take some of the tools and techniques back to their offices and start using them the following Monday.

The first step that we took was to understand the threat landscape; what exactly is a DDoS attack? This way we are able to know what the threat event we are talking about is, who perpetrates it, how it happens, and what has happened in the past. In short, a DDoS attack is when an attacker uses multiple systems, usually compromised with malware but not always, in order to target a system with the intention of affecting the availability of the system. It’s not a hack and it’s not a breach. It’s essentially the real world equivalent of standing in front of a store entrance with the intention of not letting customers in. DDoS attacks are nothing new, as they have been used for over twenty years now. They’ve been used for a variety of reasons from protesting governments, petitioning policy changes in companies, influencing elections, etc.

After defining the threat landscape, we moved to looking at the threat communities and what their methods are. Of all of the possible threat communities out there, I narrowed it down to four probable communities; hacktivists, foreign governments, cyber criminals, and cyber vandals. The first threat communities are Hacktivists who tend to occupy Wall Street as anonymous, and different sectors have different versions of hacktivists. As an example, I stated that I had worked in retail cyber security briefly and we were concerned about PETA. With this threat community, you need to take a look at who doesn’t like you and who protests you in a meat space, and then look to see if those groups have any cyber capabilities. Their capabilities have a very wide range, from relatively low power attacks to some of the most powerful attacks imaginable. It all depends on how many people they can get to attack your website. The second threat communities are foreign governments. I gave the example that the Chinese government was linked to a DDoS attack on GitHub, and that Russia attacks Ukrainian businesses on a constant basis. It does happens. The third threat communities are cyber criminals. This is probably the most common threat community as far as DDoS attacks against the financial services sector. In prior years it was hacktivists that were the most common, and now the pendulum has swung over to cyber criminal activity. Cyber criminal attacks are subdivided into two separate categories; cyber extortion and attacks using DDoS as a smokescreen. Cyber extortion attacks consist of the attacker hitting your website along with a ransom note stating that you will need to pay them or they’ll continue. Based on my research, they tend to stay away from large FIS, as they can’t overwhelm DDoS defenses long enough to make the bank want to pay, therefore making smaller companies more vulnerable to these types of attacks. The attacks using DDoS as a smokescreen are where a criminal has taken over the account of one of your customers and is about to wire transfer funds out of the country, and can launch a DDoS attack to tie up security personnel while they transfer funds out.  The fourth, and final, threat communities are cyber vandals. Cyber vandals are not typically too sophisticated. Essentially it is as if they have a shiny new DDoS toy that they try out on your site to see if it works. They don't typically have a strong agenda, so they are easily distracted or dissuaded and will move on.

We then briefly discussed Threat Agent Library's. I'm a big proponent of developing a Threat Agent Library for organizations, which is essentially a list of threat agents (FAIR calls them threat communities). In this, you list out the threat agents, a description, common tactics, methods, objectives, capabilities, and any past actions against your company or other companies in your sector. I explained that I have a threat agent library for the company that I work at and am constantly updating it with new information. It really helps me with getting risk assessments done quickly. A lot of the data is re-usable and it is something that I will constantly use when trying to ascertain the TEF when doing a risk analysis. I gave an example that I had of a Threat Agent Library from Intel that had been released publicly.

I displayed the Anatomy of a FAIR Risk Assessment at this point, showing the taxonomy that we would use to apply to our risk analysis. It was the taxonomy that we would refer to, as well as the exact steps of information gathered to get the data, that we would use for the remainder of the presentation.

The first stage we then discussed was to identify the assets at risk. DDoS attacks are typically pointed against any public facing service. Meaning that the more high the profile, the better for the attack. For assets, we identified website availability as our "asset". So how do you identify the asset's at risk? There are several options in order to do this such as talking to the people in the IT Department, talking to the application owner/data owner/data custodian/etc., looking at asset lists, reports from DLP/SIEM/etc, and talking to business continuity managers. We discussed these options at length and came to the conclusion that our asset was the website, and that the loss type was availability.

We moved onto the second stage, which was to evaluate the loss event frequency. FAIR defines this as the frequency within a given time frame that loss is expected to occur. This is made up of two inputs; threat event frequency and vulnerability. We started with TEF, breaking it down in the taxonomy to talk about Threat Event Frequency and how we derived this number. The FAIR definition of TEF is the frequency within a given time frame that threat agents are expected to act in a manner that could result in loss. This is almost always articulated in the form of a range. An example of this would be a threat analysis of a database with PII that would show that cyber criminals will act against in a way that causes a loss event between 1-5 times per year. FAIR is flexible and will let you use a number or a range, but it works best with a range. Why are ranges better? There a reasons as to why, and it's what separates FAIR from some of the other risk analysis methods in information security.

I pressed that ranges are utilized because they are really important. It is something that is covered in the FAIR material and it's a cornerstone of Doug Hubbard's book, and in addition it is something that I think is a really important concept. First, ranges help you build credibility and trust. One of the biggest challenges that you're going to have in your risk career is people arguing with you about quant analysis methods. Where you got your data, how you got your data, or how you could possibly predict the future. Most of this is going to come from within information security and IT. Business folks, especially those in fin services, are accustomed to quantitative analysis techniques and will understand this right away. Using a range will help your credibility and build trust both inside and outside of your team.

We followed this with Monte Carlo simulations. This is a huge unlock for FAIR risk assessments. It's not required, but using ranges with confidence interval allows you to run your assessment through more Monte Carlo sims. I stated that it is a complicated subject and with just an hour presentation time, I was unable to dig deeper into this aspect. The first of the Monte Carlo Sims was calibrated probability estimates was next, which I covered later on, but is a key component of an accurate FAIR risk assessment. It enables you to take imperfect and incomplete information and still do a credible analysis. Last of the Monte Carlo Sims was expressing uncertainty. This ties back to the previous three items. It's okay to have uncertainty about something, it's expected. You have to articulate it with a range of numbers. If you say one DDoS attack will happen in 2018, you seem so certain about it, but how could you know? While on the other hand, stating that the number of attacks will be between one to five a year allows you to express your level of uncertainty.

Returning back to TEF, we clarified that threat event frequency is "the probability frequency in a given time frame and that a threat agent will act in a manner that could result in a loss". To help us measure the TEF, we have Contact Frequency and Probability of Action. CF is the probable frequency that an agent will come into contact with an asset, and the PoA is the probability an agent will act once it comes into contact. You usually don't use this when assessing threats that are always hostile such as cyber criminals. It is very useful, however, when ascertaining accidental human threats. An example that I gave came from the Open Fair review books where you have an extension cord powering a server that is lying across the floor in a public place. You can measure the TEF by getting the CF with how many people walk down the hallway and step over the cord every day. The PoA is the probability of someone tripping over the cord and powering off the server.

We then focused on measuring TEF. The question for our bank is the probability frequency within a given time frame that a threat agent, such as a cyber criminal, will launch a significant and noteworthy DDoS attack. To find this, we broke down how to measure TEF by steps. Step one is to make a calibrated estimate. I always start with a calibrated estimate, a technique in which someone can express a probability of an event occurring in a way that represents their uncertainty. The question is, what is the probability frequency in a given time frame that a threat agent will act in a manner that could result in a loss. Start with the absurd, such as that the low range is 0 and the upper range is 20. Step two is then to build a list of attacks that occurred within the last two years. I chose two because the threat landscape has changed in the past two years away from hactivism and towards extortion and fraud. Step three is where you update calibrated estimate bases on new information. Finally, step four, you simply repeat steps two and three.

Referring to the previous steps, I further elaborated about building a list of DDoS attacks. I always use external data for this in conjunction with internal data if it's available. This is big difference in how I perform assessments, as Chad and Isaiah had taught me. If you use internal data only, your sample size is too small. The first steps in building a list of DDoS attacks against banks is to get a list of US banks. Any Google search will give you this. You then, once acquiring this list, will want to find the number of DDoS attacks against those banks in a two year period. I cautioned everyone to be careful with vendor studies, as many of them will use questionable methods such as non-scientific online surveys. They can be useful, but be wary of people trying to sell you FUD so you'll buy their product. Using the data you acquired, you can improve your calibrated probability estimate. I had previously stated that the range was 0 and 20, with the upper range as a slightly absurd number. My confidence in that number was pretty low. Given the new data that we found, the low range would still be 0, but that the high range would be 2. Based on our analysis the probability estimate of a significant DDoS attack is between 0 and 2 times every 2 years. We can average the attacks out, which is roughly one attack every two years, so we can use that as our Most Likely number. Based on all of my research presented, nearly all of the DDoS attacks came from cyber criminals, so we focused on that and took it a step further.

I asked how we would take this technique and apply it to DDoS attacks? The same principle applies. Take a sampling of banks, it can be a random sample or it can be a list of banks from the ABA or Fortune 100, it doesn't matter. What does matter is that the list you use should have nothing to do with DDoS attacks. After you have your list, find the occurrence of DDoS attacks for that list utilizing the same techniques such as Google News, VERIS, etc. Lastly, you will constantly reapply new data from external and internal sources.

We finished out the TEF portion of the presentation, which was really the hardest part. It's what most people struggle with, myself included, as you're dealing with so much external data. We then transitioned into Vulnerability, which is the calculation of Threat Capability and Control Strength.

We defined threat capability as the level of force a threat agent is able to apply. In FAIR, this is articulated in the form of a percentage. In our case, cyber criminals are generally smart, well resourced, and motivated by money so they are in the top 50-75% of all threats that launch DDoS attacks. They're not nation-states, which would be higher, but they are motivated and have the money, means, and people. How do we get this number? Talk to the SME's in your organization that perform threat analysis. I personally also get a huge amount of value from information sharing organizations such as FS-ISAC. They share information on threat agents and threat capabilities. Next is resistive strength, also called Difficulty in some older FAIR documents, which is a measure of how difficult it is for a threat actor in inflict harm. We rated our resistance strength at 50-90%. This means our DDoS protection services, such as Prolexix, Akamai, and Firewall Blacklisting, can resist 50-90% of attacks. It used to be 80-90% but the new attacks we are seeing forced us to update the assumptions. These would be the attacks the CIO want to know about.

Stage three was the next and last step, which I consider to be the fun part, as I feel like it's the easiest part. The reason why I think this is because, for the most part, you simply have to go talk to people on the business side. FAIR defines primary loss as the loss that is a direct result of a threat agent's actions against the asset. There are six forms of primary loss; productivity, response, replacement, fines and judgments, competitive advantage, and reputation. Productivity Loss is when there is a reduction in a company's ability to generate revenue.  The range of loss for this is anywhere from $0-300,000. Response Loss is where you are managing a loss event. The range loss for this type is larger, at approximately $1,000-300,000. Replacement loss is the cost of replacing an asset. There is no range of loss for this loss type, as it is not applicable in this particular case. Fines and Judgments Loss is where any fines are levied by regulatory agencies or judgments from people suing you. I couldn't find any case of this happening from DDoS attacks, as they usually end before people start to get angry.  The OpenFair book says this category includes bail you have to pay, which always makes me chuckle as it reminds me of a comment Evan Wheeler had made that a company objective can be to "avoid jail". Competitive Advantage loss is when it diminishes a company's position. It's not too likely to happen nowadays with DDoS attacks because they're so common. The final, Reputation loss, can be rolled up into competitive advantage.

Secondary loss, the next topic discussed, is considered 'fallout' from primary loss. Examples of this comes from the OpenFair book when customers are taking their business elsewhere after repeated outages. In the case study we were discussing, we didn't have a secondary loss. This is most often seen in data breaches or other events that have long term or cascading effects that are different from the losses measured in primary loss.

We recapped what we had discussed thus far at this point for our case study. Our asset is the website and the loss type is availability. The threat community is cyber criminals with the TEF of 0-1 times every year. The TC is 50-75% while the Resistance Strength is 50-90%. After recapping, I jumped over to the RiskLens app to derive the final risk. I taught myself FAIR with the basic risk assessment guide and later used Excel, but I prefer the RiskLens app as you can plug the numbers in and it will run all the Monte Carlo for you. We ran the numbers and came up with our final risk.

I greatly enjoyed presenting at the FAIR Institute, and look forward to future conferences! I hope that all of you reading can gain some value from my presentation as well, and can try utilizing these methods to measure DDoS Risk in the future.