NETSCOUT®

| WHITE PAPER |

How to Analyze and Reduce the Risk of DDoS Attacks

SECURITY

Executive Summary

The reality is that DDoS attacks have been around for over 20 years. And most organizations have some form of protection in place from DDoS attacks. But as this paper will point out; over the last few years DDoS attacks have increased dramatically in size, frequency and complexity.

Using risk methodology called FAIR (Factor Analysis of Information Risk) and a fictitious $50M/year e-commerce scenario, this paper is designed to help you re-assess your risk of the modern-day DDoS attack. The question we pose to you is this:

"Is the protection you may have put in place years ago, still adequate to protect you from the modern-day DDoS attack?"

If not, then you are at risk. Another way of visualizing this using the FAIR terminology is in Figure 1 which essentially asks:

"Is the Resistance Strength (i.e., DDoS attack protection) you put in place years ago, still ahead of the Threat Capability (i.e., the modern-day DDoS attack) or has the Threat Capability leap frogged your Resistance strength? — thus making you vulnerable and at risk."

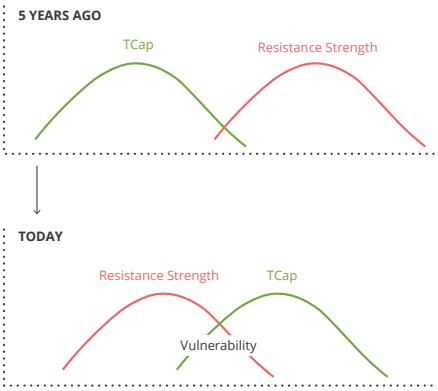

Figure 1: Threat Capability (TCap) continuum.

There are multiple ways to stop a DDoS attack. This paper we'll compare an "as-is" scenario using a firewall to multiple "to-be" scenarios using different NETSCOUT Arbor DDoS attack protection solutions.

The analysis will show, though highly effective in some cases, implementing only an on-premises solution (e.g., To-Be#1a: Arbor APS) is not enough to mitigate all types of DDoS attacks — specifically large attacks which will overwhelm internet bandwidth. And vice versa. Implementing only an in-cloud solution (e.g., To-Be #1b: Arbor Cloud) will be required in some cases, such as when large attacks overwhelm internet bandwidth, but may not effectively mitigate smaller, hard to detect, application layer attacks. Ultimately, the analysis will show that implementing a fully managed, intelligently integrated combination of on-premises and cloud-based solutions (e.g., To-Be #1c: Arbor APS + Arbor Cloud) offers the most comprehensive protection — and thus reduces the most risk and loss exposure from DDoS attacks.

Determining risk of DDoS attack is one thing. Deciding upon and justifying the need for different methods of DDoS attack protection is another. The last section will take our risk analysis a step further as we provide data to help make the business case for DDoS attack protection by analyzing the cost benefits of various Arbor-based DDoS attack protection solutions. Again, the analysis will ultimately show that implementing a layered approach to DDoS attack protection (i.e., a combination of on premise Arbor APS + Arbor Cloud) not only offers the most comprehensive form of protection, but also reduces the loss magnitude by the largest percentage; while still providing a healthy ROI of approximately 150%.

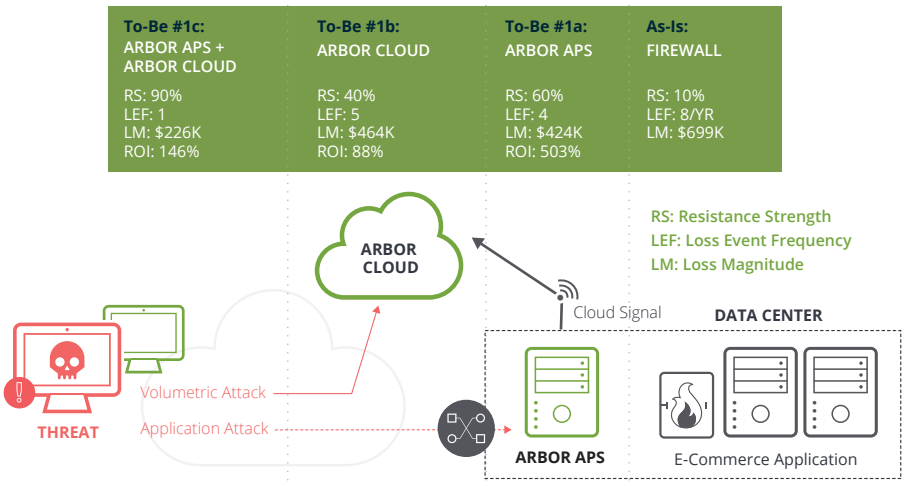

Figure 2: Comprehensive, cost effective risk analysis.

The Modern-Day DDoS Attack

Distributed Denial of Service (DDoS) attacks bring significant risk to organizations that depend on their networks and websites as an integral part of their business.

As organizations that have suffered DDoS attacks will attest — there is no question whether or not DDoS attacks are happening or whether they bring risks. There are however lots of questions regarding how much risk is associated with DDoS, and how much risk is reduced using traditional security controls such as firewalls, and specific DDoS attack protection controls including cloud scrubbing and on-premises, in-line DDoS mitigation appliances.

The modern-day DDoS attack is complex, as Figure 3 depicts.

Today's DDoS attack uses a dynamic combination of multiple vector attack vectors consisting of:

1. **Volumetric** Large bandwidth consuming attacks (e.g., UDP flood) designed to saturate network pipes and internet facing router interfaces.
2. **TCP State-Exhaustion** Attacks (e.g., TCP-SYN) designed to fill TCP State tables in devices such as firewalls, IDS/IPS and load balancers.
1. **Application-Layer** Low and slow application layer attacks (e.g., HTTP header, SlowLoris) designed to slowly exhausts resources in application servers.

According to NETSCOUT® Arbor 13th Annual *Worldwide Infrastructure Security Report*, 59% of respondents have experienced a multi-vector DDoS attack.

Multi-vector DDoS attacks are not a recent phenomenon as they have been around since 2010. What's different today is the ease at which these attacks can be launched by unsophisticated threat actors due to the plethora of Do-It-Yourself DDoS attack tools and DDoS-for-Hire services.

This is exacerbated by the rapid proliferation of inadequately secured IoT devices which are being consumed into IoT-based botnets and weaponized to launch multi-vector DDoS attacks. For example, in the fall of 2015, the IoT-based botnet Mirai was used to launch attacks against well know security blogger Brian Krebs and DYN (a DNS Service provider) which impacted popular sites such as GitHub, NetFlix, Twitter and others.[1]

To further underscore the asymmetry of DDoS attacks; for a mere $5/hour an unsophisticated attacker can rent an IoT-based botnet, launch a multi-vector DDoS attack and potentially cause hundreds of thousands of dollars in damage. Now that's ROI!Add to this the many motivations behind launching DDoS attacks, including geo-political protest, extortion, competitive take out etc. and you can see why DDoS attacks are dramatically rising in size, frequency and complexity.

Mirai botnet is a modern-day multi-vector attack.

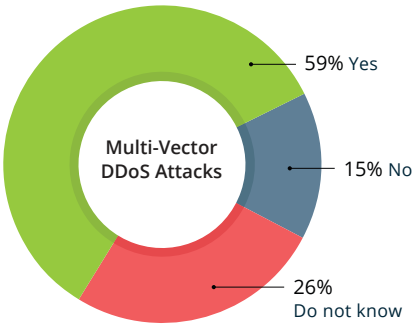

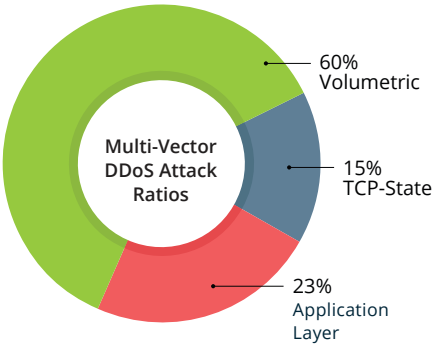

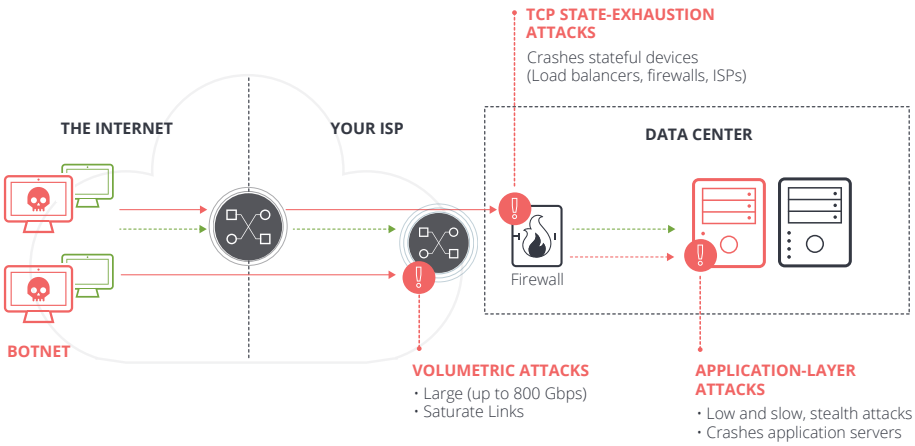

Figure 3: The modern day DDoS attack is complex; dynamic and multi-vector.

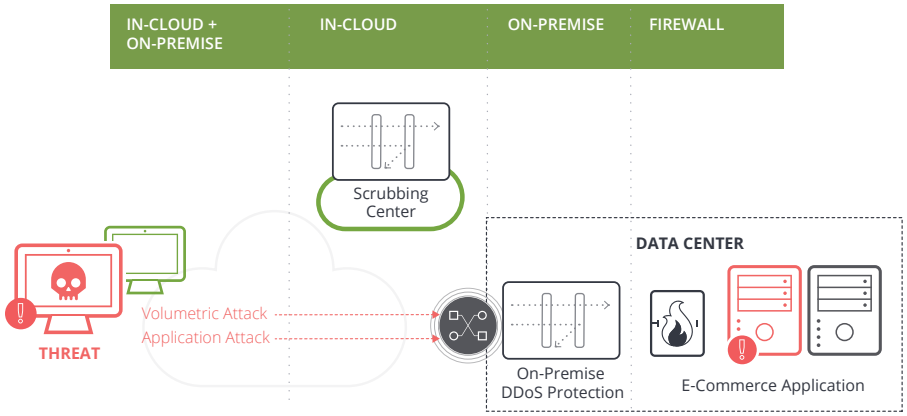

Figure 4: Methods of DDoS attack protection.

Approaches to DDoS Mitigation

Before analyzing the risk, it is useful to understand the protection provided by the various security controls being analyzed, as well as their limitations.

Firewalls

Firewalls have long been deployed at network perimeters in an attempt to keep malicious attackers from penetrating corporate networks. The fact is that firewalls are not very effective in dealing with the modern-day, multi-vector DDoS attack.

In-Cloud Scrubbing Services

Going back to the early 2000's when DDoS attack methods were first being experienced, cloud-based scrubbing services began emerging as a means to inspect large volumes of traffic in ISP networks, in order to remove malicious traffic before allowing it to enter corporate networks. Given that the majority of DDoS attacks are volumetric attacks, cloud scrubbing services deliver a high degree of effectiveness in mitigation. However, cloud-based scrubbing services do have a weakness. Due to the massive amount of traffic they are analyzing, they struggle to recognize the "low & slow" application-layer attack.

In-Line, On-Premises Appliances

With the advent of "low and slow" application-layer attacks, the industry realized that neither firewalls nor cloud scrubbing services, in and of themselves, were effective at mitigating these types of attack. A new technology was required that could be deployed on-premises in the customer network, which could identify and mitigate anomalous traffic including application layer attacks.

In fact, industry best practices dictate that on-premises DDoS protection appliances should be deployed in front of firewalls to help maintain the firewall's availability during a DDoS attack. On their own, in-line DDoS protection appliances are capable of dealing with volumetric DDoS attacks — to a point. That is, on-premises DDoS protection appliances do a fine job with mitigating DDoS attacks that do not exceed the internet circuit size. Once the attack exceeds the size of the internet pipe, the on-premises solution can be rendered useless since the upstream bandwidth is fully saturated — this establishes the need for in-cloud protection. According to Arbor's 12th WISR 41% of Enterprises and 61% of data center operators have experienced DDoS attacks that exceed their internet bandwidth.

With the increase in multi-vector attacks, it has become very clear that reducing the risk from DDoS attacks requires a defense-in-depth or hybrid approach utilizing all of the mitigation approaches described above.

43%

of organizations had firewall or IPS devices experience a failure or contribute to an outage during an attack.

Arbor's 12th Annual WISR

41%

of Enterprises have experienced DDoS attacks that exceed their internet bandwidth

Arbor's 12th Annual WISR

61%

of Data Center Operators have experienced DDoS attacks that exceed their internet bandwidth.

Arbor's 12th Annual WISR

Introduction to DDoS Attack Risk Analysis and Arbor Solutions

In our risk scenario, we will analyze the amount of risk that is present in the current as-is state, with an assumed minimal set of security controls (i.e., a firewall) implemented. After analyzing the risk present in the current state, we will then look at how much risk is mitigated upon implementation of three different Arbor DDoS attack protection solutions; essentially the to-be risk states.

The table below summerizes our risk scenario:

Threat	A modern-day DDoS attack which executes a dynamic combination of volumetric, TCP state exhaustion and application layer attack vectors.
As-Is State #1	Existing firewall.
To-Be State	—
#1a. Firewall + Arbor APS (On-Premise)	Arbor APS is an automated, in-line, on-premises DDoS attack protection appliance (or virtual appliance) capable of stopping all types of DDoS attacks. APS mitigation capacities range from sub 100 Mbps to 40 Gbps.
#1b. Firewall + Arbor Cloud DDoS Protection Service (In-Cloud)	Arbor Cloud is an in-cloud, managed DDoS protection service that has multiple scrubbing centers around the globe offering multi-Tbps of mitigation capacity.
#1c. Firewall + Managed Arbor APS (On-Premises) + Arbor Cloud Service (In-Cloud)	A fully managed, combination of on-premises and in-cloud protection. In the event of a large DDoS attack that will saturate the local internet connection, the Arbor APS; via a feature called "Cloud Signaling"; can automatically and intelligently redirect attack traffic to the Arbor Cloud.

Open FAIR Risk Analysis Methodology

The methodology used to perform the risk analysis is Open FAIR, comprising two open industry standards, the Risk Taxonomy Standard (O-RT), and the Risk Analysis Standard (O-RA).[2] The Open FAIR risk taxonomy used in the analysis is depicted in the graphic below:

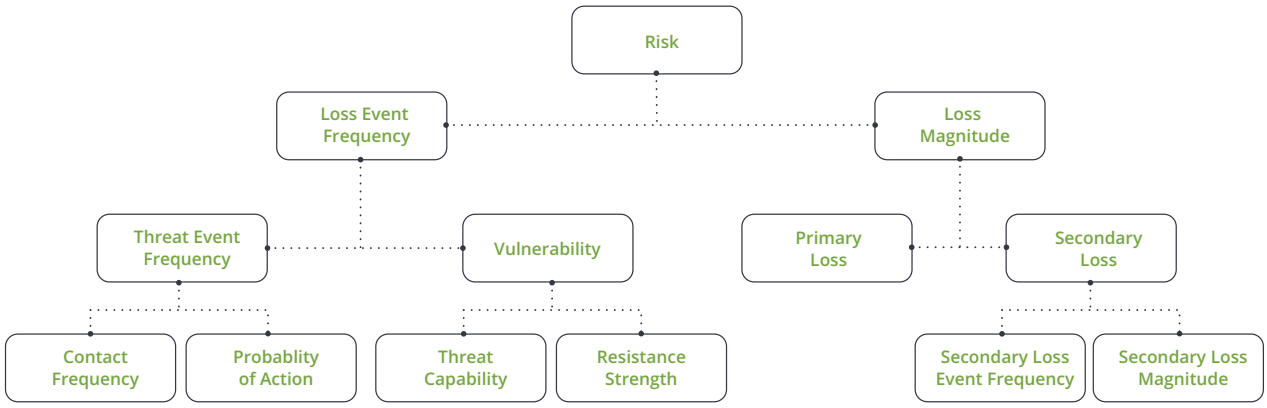

Figure 7: Open FAIR risk taxonomy.

In the Open FAIR taxonomy, it is important to note that risk is a derived value, and that risk is expressed in terms of probability of $ loss in a given time period.

The Open FAIR standards are useful in decomposing risk to describe both impact and frequency in standard, measurable ways, in providing calibrated estimation tools, and in developing quantitative analyses of specific risk scenarios.

Although this risk scenario is specific to multi-vector DDoS attacks on an e-commerce site, the methodology used may be easily applied to other DDoS risk scenarios in other industries.

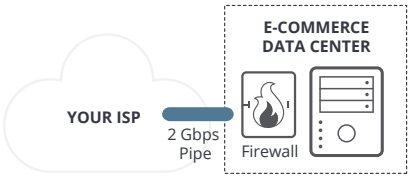

Figure 8: Data center Internet access is 2 Gbps.

The FAIR Risk analysis process of the as-is state and of three to-be states follows this general sequence of steps:

1. Describe the risk scenario including Asset(s) at risk, Threat Community, Threat Type and Threat Effect
2. Analyze Threat Event Frequency and evaluate Loss Event Frequency
3. Evaluate Primary and Secondary Loss factors
4. Determine vulnerability, including assessment of Threat Capability and Resistance Strength of As-Is state
5. Derive risk and produce analysis reports of As-Is state
6. Repeat step 4, with To-Be States
7. Repeat Step 5, for To-Be states; ultimately producing reports comparing all As-Is and To-Be states

Execution of the FAIR Risk Analysis for DDoS Attack

Using the general sequence of steps of FAIR risk analysis as described above, let's conduct the risk analysis for DDoS attacks against a fictitious e-Commerce company.

Step 1. The Risk Scenario

The risk scenario can be summarized by the table below:

Risk Scenario	The risk of a multi-vector distributed denial of service attack on an $50M/year e-commerce website that is connected to the internet with a 2 Gbps of Internet bandwidth.
Threat Actor	A non-skilled hacktivist leveraging a readily available, inexpensive, yet highly effective DDoS-for-hire service.
Threat Type	Malicious
Threat Effect	Availability

The organization for this risk analysis is assumed to be an online retailer, doing all of its business via an e-commerce site. For this risk scenario, the assets at risk include the web servers, databases, and network infrastructure comprising the retailer's e-commerce site. These assets (i.e., their value) generate $50M in annual revenue.

Since network capacity is such a relevant factor in DDoS attack risk analysis, our scenario is further refined by an assumption that the internet circuit feeding the data center hosting the e-commerce site is 2 Gbps.

The risk scenario being analyzed is the risk from multi-vector DDoS attacks, perpetrated by an unskilled hacktivist, utilizing a readily available, inexpensive, but highly effective DDoS for hire service.

Step 2. Threat Event and Loss Event Frequency

When determining Threat Event Frequency, one can reference a plethora of publicly available reports from industry leading vendors and researchers. For example:

Arbor's 12th Annual Worldwide Infrastructure Security Report (January 2017)

Attacks Per Month	Attacks Annually	% Experienced
1–10	12–120	55%
11–50	132–600	30%
51–100	51–100	9%
100+	1,200+	6%

Neustar's 2016 Worldwide DDoS Attacks and Protection Report

Attacks Annually	% Experienced
1	11%
2–5	29%
8–10	14%
12	10%
48	5%
50+	3%

We can even refine this information further by looking at specific industries. For example:

Arbor's 12th Annual Worldwide Infrastructure Security Report (January 2017)

When internet service providers were asked "*What types of industries are targets of DDoS attacks (check all that apply)?*" respondents replied with:

41% Government	13% Law Enforcement
41% Financial Services	10% Healthcare
40% Hosting	10% Energy/Utlities
36% eCommerce	9% Gambling
35% Gaming	7% Manufacturing
31% Education	7% Other

Neustar's 2016 Worldwide DDoS Attacks and Protection Report

Provided some analysis per specific verticals:

	Respondents suffered a DDoS attack	Attacked multiple times a year	Attacked at least once a month
Retail	**73%**	**83%**	—
Tech	81%	85%	25%
Financial Services	76%	87%	33%

Verizon's 2017 Data Breach Investigations Report

Looked at the number of DDoS Incidents different industry verticals experienced on a yearly basis:

Industry	# of Incidents (Per Year)
Public	617
Information	508
Finance	445
Education	228
Retail	180
Manufacturing	10
Accomodation	4
Healthcare	3

Column: Column - Body - Eostruptatum facipsu ntoriassin prae eosserum, conet omnimodi conecep eresequos qui volorereiur.

Considering the data given above; specifically for our e-commerce scenario, we are assuming the following for frequency of attacks per year:

Threat Frequency

Minimum/Year	Most Likely/Year	Maximum/Year
3	24	120

In order to determine loss event frequency, we will need an understanding of what a loss event looks like. DDoS attacks can create a number of different types of loss events. For the purposes of this analysis, we are using a commonly seen loss event, where the DDoS attack aims to prevent customer access to a retailers online ordering system, thereby causing monetary loss to the primary stakeholder (the retailer) from loss of availability and thus revenues from the e-commerce site.

Not all threats events will result in a loss event. For our analysis we are making the assumption that 1 out of every 3 threat events will result in a loss. With that in mind are number look like:

Loss Event Frequency

Minimum/Year	Most Likely/Year	Maximum/Year
1	8	40

A final piece that will influence loss event frequency is to consider the duration of the attacks. Again, referencing data from multiple sources:

Arbor's 12th Annual Worldwide Infrastructure Security Report (January 2017)

Duration of attacks:

Duration	% Experienced
7 Hours or Less	72%
7–24 Hours	17%
1 Day to 4 Weeks	11%

This was also corroborated with data from Arbor's ATLAS® which shows that **90% of attacks last less than 1 hour.**

Very similar stats were also shown in Verizon's 2017 Verizon Data Breach Investigations Report which indicated that a **majority of attacks last less than 5 hours.**

Kapersky Labs Q4 2017 DDoS Analysis Report also indicated that **69% of attacks were less than 4 hours.**

For our analysis, we are assuming:

Average Duration of Attack

Minimum/Year	Most Likely/Year	Maximum/Year
1 hour	2 hours	48 hours

Step 3. Determining Primary & Secondary Loss Factors

FAIR defines 6 categories of Primary and Secondary Loss:

Material Areas of Loss

Primary	Secondary
1. Productivity*	4. Fines and Judgement
2. Primary Response*	5. Reputation*
3. Replacement	6. Secondary Response

*Loss we will analyze in this paper.

In this risk scenario, the most significant form of Primary Loss is Productivity, defined as the ability of the e-commerce website to take orders and produce revenue. To determine what the impact of website unavailability is in terms of revenue, we're making an assumption that the e-commerce site being analyzed produces $50M in revenue/year.

In terms of revenue productivity loss, for our $50M e-commerce site, this translates to $95.13 per minute (or $5,708 per hour) of unavailability, which we are using as a starting point to calculate the impact of the e-commerce site being unavailable during DDoS attacks.

This amount could be considered conservative when compared to the following:

Arbor's 12th Annual Worldwide Infrastructure Security Report (January 2017)

Supports this estimate, with the cost of internet downtime reported by respondents as:

Cost of Internet Downtime/Min	% Experienced
$500	41%
$500–$1,000	23%
$5,001–$10,000	18%
$20,001+	18%

Neustar's 2016 Worldwide DDoS Attacks and Protection Report

49% of respondents would lose $100K or more per hour during peak periods due to a DDoS attack.

The assumptions yield the following loss calculations due to unavailability (per successful attack or Loss Event previously estimated at Min=1hr, Most Likely=2hrs, Max=48hrs):

Primary Loss: Productivity

	Minimum	Most Likely	Maximum
Average Duration	60 min (1 hour)	120 min (2 hours)	2,880 min (48 hour)
Primary Loss Impact Due to Unavailability $95.13/Minute ($5,707.08/Hour)	$5,707.80	$11,415.60	$273,974.40

It is also worth noting that for many e-commerce sites, not all days are created equal. The busy holiday season, for example, can represent a disproportionate % of annual revenues, such that DDoS attacks occurring during this time can create greater losses. We have not attempted to model this "peak period" effect into the risk analysis, although the impact of it could show up in both reputation damage and in productivity loss.

Another Primary Loss factor we will considering for this analysis is Response Costs. We are assuming between 1 FTE and 10 FTE are tasked to respond, at twice the duration of the event, at an average annual cost/FTE to the organization of $120K or $60/hour.

Attack Response Costs

	Minimum	Most Likely	Maximum
Response Cost/Attack	$120	$240	$5,760

For Secondary Loss this analysis will include Reputation loss. We used a study published by the IEEE, Analyzing the Impact of a DDoS Attack Announcement on Victim Stock Prices[2], which evaluated the impact of DDoS attacks on share prices. The study concluded "in some cases there is a noticeable negative impact on the stock prices of the victim firm whenever the attack causes interruptions to the services provided by the firm to its customer." Impacts noted in the study ranged from 0 to -6% share price drops, depending on the company and time period assessed.

For our analysis we'll use assume estimates for minimum, most likely, and maximum impacts are based upon share price drop estimates of 0%, -.005%, and -2% respectively, which are applied to an estimated market capitalization for our e-commerce company of $50 Million (Note: e-commerce companies are commonly valued at from .5x – 6x revenues, and we're valuing this fictitious retailer conservatively at 1x revenues).

Secondary Loss: Reputation

	Minimum (0%)	Most Likley (-.005%)	Maximum (-2%)
Response Cost/Attack	$0	$250,000	$1,000,000

In order to determine Secondary Loss, we also need to determine how frequently secondary loss is experienced by secondary stakeholders (Secondary Loss Event Frequency, or SLEF), as well as minimum, most likely, and maximum values for this loss magnitude. The objective here is to determine for all DDoS attacks experienced, how many of them result in publicly noticed outages that might impact market capitalization. Our assumptions are in the table below:

Secondary Loss: Reputation

	Minimum	Most Likley	Maximum
Secondary Loss Event Frequency/Year Estimated as a % of the Loss Events that become publicly known.	0%	1%	2%

To simplify our scenario analysis for this paper, we have elected to not calculate any loss in the areas of: Replacement, Fines/Judgments and Secondary Response. However, depending upon the type of organization, industry, etc. these losses can be significant and should not be overlooked.

Step 4. Determine Vulnerability, Including Resistance Strength of As-Is State

It is commonly known that firewalls alone are not very effective at dealing with multi-vector DDoS attacks. Per Arbor's 12th Annual Worldwide Infrastructure Security Report, 43% of respondents witnessed their firewalls or IPS/IDS devices experience fail or contribute to an outage during a DDoS attack. With this data in mind, the Resistance Strength of the As-Is state is estimated in the table below. Note: The higher the resistance strength the more effective the protection.

Resistance Strength	Minimum Effectiveness	Most Likley Effectiveness	Maximum Effectiveness
As-Is #1: Firewall Only	0%	10%	25%

Step 5. Derive Risk and Produce Reports of As-Is State

Our analysis of the risk to the e-commerce organization in the as-is state, with the set of assumptions as described above produces the summary below of the amount of risk that is present when using only a firewall for DDoS attack protection. The software used to produce the analysis and reports, provided by RiskLens[3], captures the inputs as described, and uses Monte Carlo analysis to simulate 1,000's of possible outcomes. The first chart shows the loss exposure at various places in the Monte Carlo distribution, while the second chart depicts where the loss comes from, by form of loss category.

Loss Exposure

Maximum	$4.8M
90%	$1.6M
Average	$699K
10%	$112K
Minimum	$12K

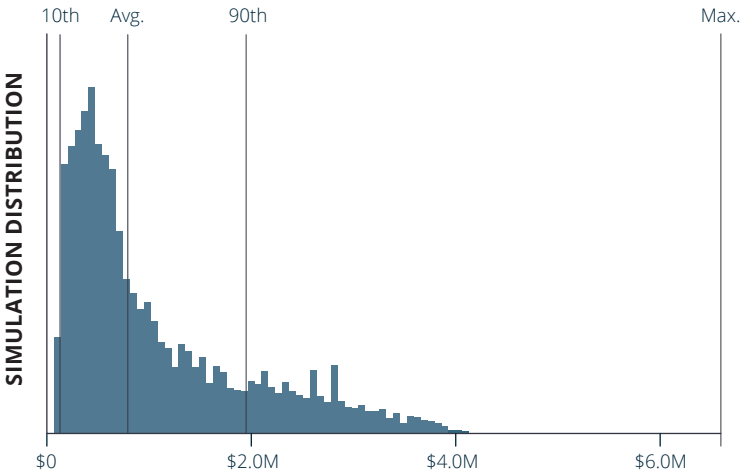

Figure 9: Loss exposure using Monte Carlo Analysis.

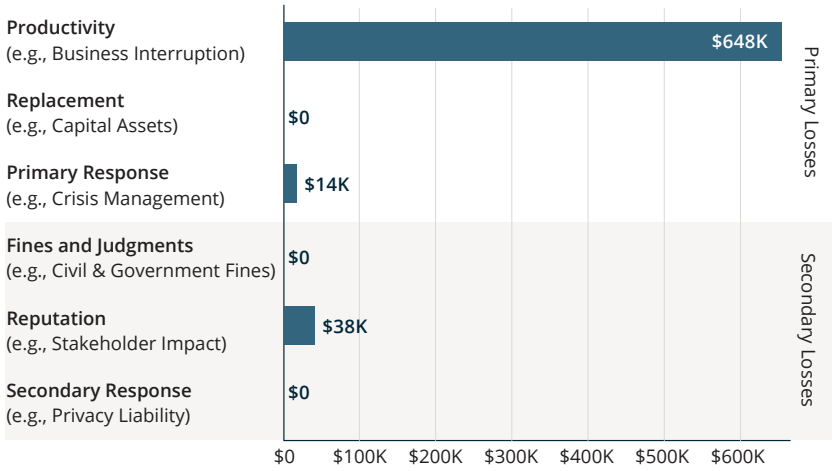

Figure 10: Materialized Areas of Loss.

Step 6: Determine Vulnerability, including Resistance Strength of To-Be States (Arbor APS, Arbor Cloud, and Managed Arbor APS + Arbor Cloud)

Having analyzed the risk exposure in the current as-is state, we'll now evaluate the following Arbor-based to-be scenarios in terms of their potential to reduce risk from DDoS attacks. Our analysis of the different to-be scenarios generally follows the process previously described in the as-is state. The significant differences in each of the three to-be scenarios is the degree of resistance strength provided by the added Arbor-based security controls, which is reflected in lower frequencies of loss events. We have also reduced the response costs for the to-be scenarios, to reflect the fact that the additional controls reduce overall impact, and recovery efforts.

In our next set of analyses, the improvements in risk reduction derive from two changes to the scenario inputs.

First Change In Scenario Input

The degree of Resistance Strength provided is enhanced due to the addition of each security control. *Recall: The internet circuit size for the e-commerce site is 2 Gbps.* Some statistics to consider for analysis of resistance strength:

Attack Types

As previously mentioned, Arbor WISR indicates that 59% of respondents have experienced a multi-vector DDoS attack. The attacks type ratios are:

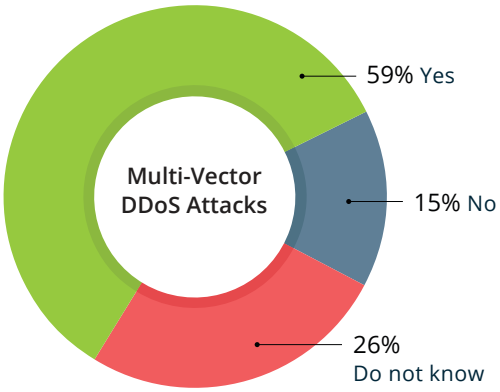

Figure 10: Materialized Areas of Loss.

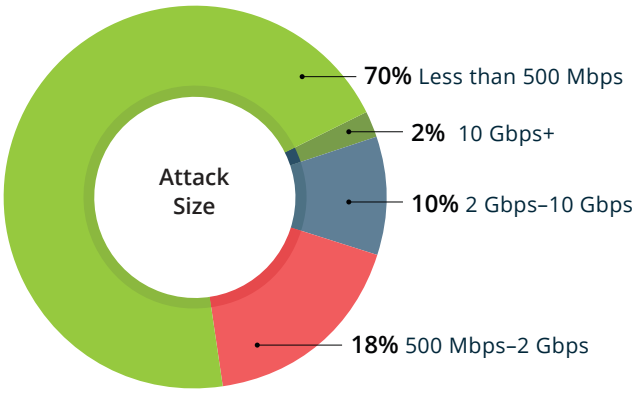

Figure 12: DDoS attack size.

Attack Size

According to Arbor ATLAS data, 70% of attacks are less than 500 Mbps.

Arbor's ATLAS data also showed the average attack size is now 931 Mbps, up from 760 Mbps in 2015, a 23% increase. Arbor predicts that by the end of 2017, the average size DDoS attack will be over 1 Gbps.

Why is this a concern? Most organization's internet bandwidth is less than 1 Gbps — which means the average size DDoS attack can overwhelm them – thus the need for cloud-based protection.

Data from Verizon and Neustar corroborate this:

Verizon's 2017 Verizon Data Breach Infographic Report indicates the median size DDoS attack was 4.97 Gbps.

Neustar's 2016 Worldwide DDoS Attacks and Protection Report DDoS attacks are 1-4.9 Gpbs.

With these numbers in mind let's assume the majority of attacks are under 1 Gbps and the average size DDoS attack is 3 Gbps. Which means that in most cases our e-commerce sites' internet bandwidth of 2 Gbps would not be saturated, and on-premise DDoS attack protection would be effective. This also means that in fewer than 50% of cases our e-commerce sites' internet bandwidth of 2 Gbps would be saturated, and in-cloud protection will be required.

With the above-mentioned assumptions in mind, the table below summarizes the % of resistance strength provided in each scenario. The higher the resistance strength the more effective the protection. Refer to the Notes column for additional details including Strengths and Vulnerabilities.

Resistance Strength	Minimum Effectiveness	Most Likely Effectiveness	Maximum Effectiveness	Notes
As-Is #1: Firewall Only	0%	10%	25%	Firewalls are not very effective in dealing with multi-vector DDoS attacks. Per Arbor's 12th annual Worldwide Infrastructure Security Report, 43% of respondents witnessed their firewalls or IPS/ISD devices experience fail or contribute to an outage during a DDoS attack.
To-Be State				
#1a: Firewall + Arbor APS Inline DDoS protection appliance	30%	60%	100%	**Strength:** Since most DDoS attacks are less than 1 Gbps and Arbor APS has ability to automatically detect and stop all types of DDoS attacks; this will reduce time to mitigation and thus cost of response (FTE). **Vulnerability:** APS will become ineffective when the attack is larger than the size of the 2 Gbps internet circuit.
#1b: Firewall + Arbor Cloud Cloud DDoS scrubbing service	30%	40%	100%	**Strength:** If we assume that in some cases, the attacks will be larger than the 2 Gbps internet circuit; then Arbor Cloud will be the appropriate form of protection. Also since Arbor Cloud is a 24x7 managed service it can also reduce response (FTE) costs. **Vulnerability:** Though it can stop them, Arbor Cloud may be delayed in mitigating low and slow application layer attacks. This could increase attack duration and loss.
#1c. Firewall + Managed Arbor APS + Arbor Cloud	50%	90%	100%	**Strengths:** The intelligently integrated combination of Arbor APS and Arbor Cloud offers the most comprehensive form of protection from dynamic, multi-vector DDoS attacks. Also, since the entire solution is fully managed, response time (FTE) costs are further reduced.

Note: The resistance strengths in the table above are greatly influenced by the internet circuit size. The numbers in this table only reflect our scenario and could be different for other scenarios (e.g., internet bandwidth is 1 Gbps).

Second Change In Scenario Input

The second area where a difference in analysis inputs occurs is in the % of threat events that become loss events. Recall that in the as-is analysis, we estimated that for every three threat events (DDoS attacks), one resulted in an actual loss event. The addition of each additional security control further reduces the number and percentage of attacks which become actual loss events. Using the Resistance Strength from the previous table, our revisions to ratio of threat events to loss events for each additional mitigation scenario are in the table:

Resistance Strength	Minimum Number of Loss Events *Max Effectiveness of Resistance Strength*	Most Likely Number of Loss Events *Max Effectiveness of Resistance Strength*	Maximum Number of Loss Events *Max Effectiveness of Resistance Strength*
As-Is #1: Firewall Only	1	8	40
To-Be State			
#1a: Firewall + Arbor APS Inline DDoS protection appliance	0	4	28
#1b: Firewall + Arbor Cloud Cloud DDoS scrubbing service	0	5	28
#1c. Firewall + Managed Arbor APS + Arbor Cloud	0	1	20

Step 7. Deriving Risk and Producing Reports for All As-Is and To-Be States

Utilizing these new assumptions reflecting the resistance strength of the added Arbor DDoS protection options, we conduct our analysis of the risk present in each of the four scenarios. As see in the table charts below (which is sorted by Average Loss Exposure), the amount of risk that is present in these three to-be state scenarios is obviously much lower than the original as-is state. What's interesting to note is the different Loss Exposure for each Arbor solution. It's important to note that the resulting analysis is based upon the depicted scenario. Results could be different with a different scenario (i.e., size of internet pipe, frequency of loss events etc.) The analysis clearly shows implementation of the complete Arbor DDoS solution (combination of Arbor Cloud + Arbor APS), in addition to the existing firewall, provides the best level of protection and lowest level of risk.

Analysis	Minimum	10th%	Average	90th%	Maximum
To-Be #1c: Arbor APS + Arbor Cloud	$0	$17K	$226K	$572K	$2.6M
To-Be #1a: Arbor APS	$0	$51K	$424K	$1.0M	$3.7M
To-Be #1b: Arbor Cloud	$0	$62K	$464K	$1.1M	$4.3M
As-Is #1: Firewall	$12K	$112K	$669K	$1.6K	$4.8K

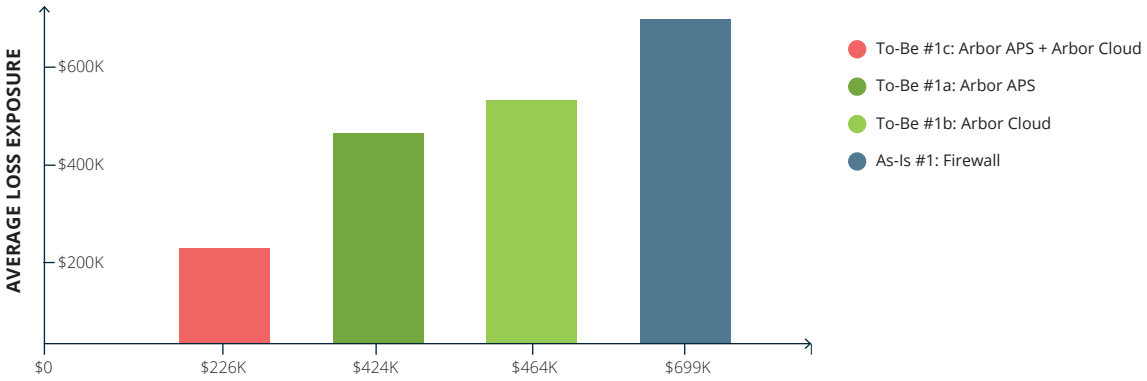

Figure 13: Loss exposure of as-is and to-be states.

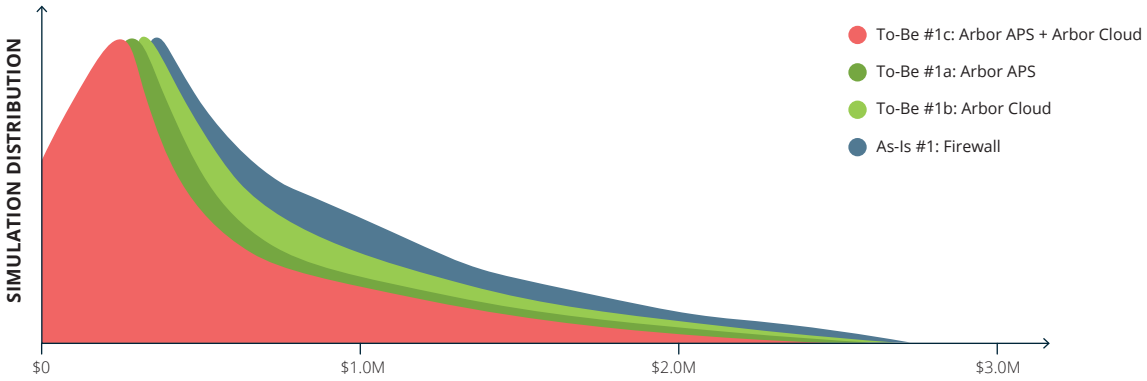

Figure 14: Monte Carlos analysis of loss exposure for as-is and to-be states.

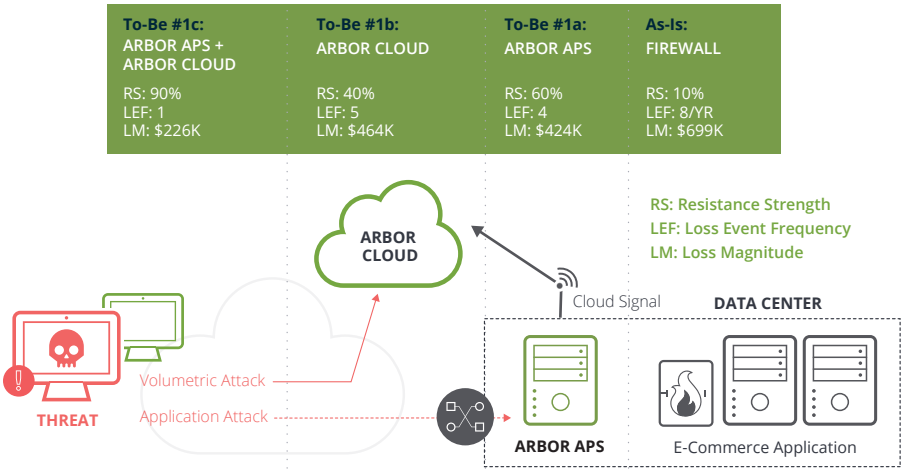

Figure 15: As-is (firewall) and multiple To-Be Arbor DDoS protection solutions with Risk analysis (Resistance Strength (RS), Loss Event Frequency (LEF), and Loss Magnitude (LM)).

Yet another way to visualize is shown in the diagram above. The analysis in this paper and industry best practices have shown that the most comprehensive form of protection from the modern-day DDoS attack is to take an automated, layered approach.

In other words, deploy a fully managed combination of:

1. Arbor APS on-premises to stop all types of DDoS attacks; including application layer attacks, that do not exceed internet bandwidth.
2. In the event of a large attack; via Cloud Signaling, automatically and intelligently redirect traffic to Arbor Cloud for mitigation.

Making the Business Case

This paper could have ended at the previous section as we showed how implementing various to-be scenarios reduced the risk exposure for our e-commerce scenario.

In this section we're going to take the analysis a step further. It's one thing to see the reduction in risk exposure; but at what cost? In other words, how much is your organization willing to invest in protection to reduce the risk and ultimate loss. In this section we compare the cost of added control strength (i.e., different Arbor solutions) to risk reduction benefit. We are using a simple formula of:

$$\textbf{Return on Investment (ROI)} = \frac{\textbf{Reduction in Risk – Cost of Arbor Protection}}{\textbf{Cost of Arbor Protection}}$$

The table below is a summary of our analysis:

Risk Comparisons (Most Likely)	Resistance Strength	Loss Events	Loss: Productivity	Loss: Response	Loss: Reputation	Total Loss Exposure	Loss Reduction (from As-Is)	% Reduction	ROI
As-Is: Firewall Only	10%	8	$648,000	$14,000	$38,000	$700,000	—	—	—
To-Be #1a: Arbor APS	60%	4	$382,000	$8,000	$34,000	$424,000	$276,000	39%	503%
To-Be #1b: Arbor Cloud	40%	5	$431,000	$5,000	$27,000	$463,000	$237,000	34%	88%
To-Be #1c: Arbor APS + Arbor Cloud	90%	1	$212,000	$439	$14,000	$226,439	$473,561	68%	146%

It's interesting to compare the Loss Reduction/% Reduction columns to the ROI column. For example; the To-Be#1a: Arbor APS scenario which reduces risk exposure by 39% and has a high ROI of 503%. Versus the To-Be#1b: Arbor Cloud scenario which reduces loss exposure by 34% but has less ROI of 88%. Ultimately, as in previous analysis we see that the To-Be #1c: Arbor APS + Arbor Cloud is a justifiable solution as it provides the highest % loss reduction of 68% and a healthy ROI of 146%. It's clear how analysis such as this can help with the proper business justification for DDoS attack protection. Our final risk and ROI analysis can be depicted as below:

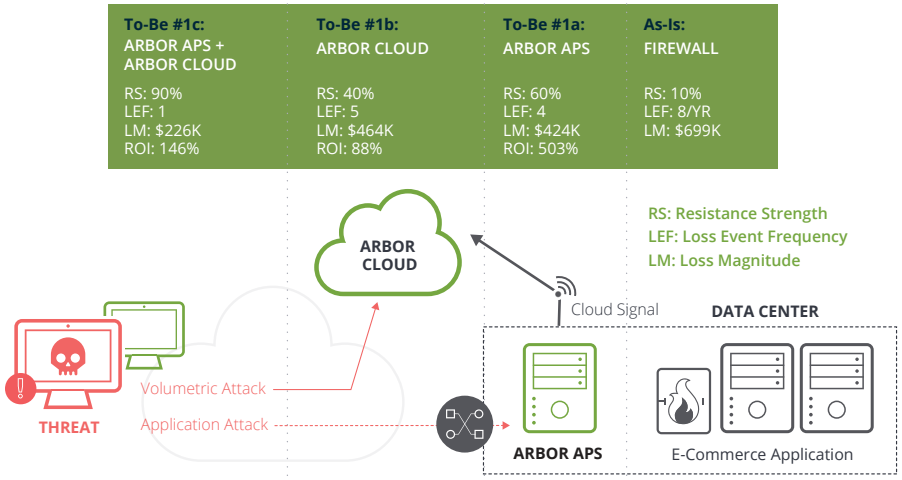

Figure 16: Risk analysis and Return on Investment (ROI) of Arbor DDoS attack protection.

Summary

As with many areas of security, today's DDoS threat is not the same as it was ten or even five years ago. Threat actors have become much more capable, and with the advent of hacking for profit services, highly skilled attack services are available to hacktivists and others. Simple security controls such as firewalls are incapable of addressing todays threats. Even cloud-only scrubbing services that were previously effective against smaller scale volumetric type DDoS attacks are not capable of mitigating the hybrid attacks that are becoming common today.

As the threats evolve, so too do our defense strategies need to evolve, moving toward a defense-in-depth approach to DDoS attacks.

As the risk analyses of the current state and the various future states shows, addition of additional security controls in the form of the Arbor Cloud and Arbor APS solutions brings significant additional risk reduction to organizations dealing with the impacts of DDoS attacks.

LEARN MORE

Contact your local Arbor sales representative to learn more about Arbor DDoS Attack Protection Solutions.

Appendices

[1] www.opengroup.org/subjectareas/security/risk

[2] IEEE, http://ieeexplore.ieee.org/document/7912671/

[3] The provider of the RiskLens risk analysis software product is RiskLens. For more information see: www.risklens.com.

NETSCOUT

Corporate Headquarters
NETSCOUT Systems, Inc.
Westford, MA 01886-4105
Phone: +1 978-614-4000
www.netscout.com

Sales Information
Toll Free US: 800-309-4804
(International numbers below)

Product Support
Toll Free US: 888-357-7667
(International numbers below)

NETSCOUT offers sales, support, and services in over 32 countries. Global addresses, and international numbers are listed on the NETSCOUT website at: www.netscout.com/company/contact-us

SECWP_005_EN-1802 07/2018