



Interested in learning  
more about security?

# SANS Institute InfoSec Reading Room

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

## Denial of Service Deterrence

Denial of Service has been a very useful practice for attackers and continues to remain prevalent today. The main issue for organizations is how to implement a denial of service solution conducive to their enterprise environment. The logic behind implementing a DoS solution is simple. These attacks are a monumental portion of an attackers arsenal, require few local resources to execute, and can be the most adverse in terms of infrastructure and functionality of the business. Denial of service will only continue to be m...

Copyright SANS Institute  
Author Retains Full Rights

AD

Veriato

Unmatched visibility into the computer  
activity of employees and contractors



Try Now

# Denial of Service Deterrence

*GIAC (GSEC) Gold Certification*

Author: Ryan William Sepe, [RyanSepe@hotmail.com](mailto:RyanSepe@hotmail.com)

Advisor: Dominicus Adriyanto Hindarto

Accepted: January 24<sup>th</sup>, 2015

## Abstract

Denial of Service has been a very useful practice for attackers and continues to remain prevalent today. The main issue for organizations is how to implement a denial of service solution conducive to their enterprise environment. The logic behind implementing a DoS solution is simple. These attacks are a monumental portion of an attackers arsenal, require few local resources to execute, and can be the most adverse in terms of infrastructure and functionality of the business. Denial of service will only continue to be more of an issue for enterprises who do not employ a DoS mitigation solution. This paper seeks to analyze the landscape and depict environment configuration variables. Finally, current techniques for mitigation and remediation for denial of service will be described and a framework for an organization's implementation plan delineated.

## 1. Introduction

Denial of service attacks have been around since 1989 and may have been incorporated even before that time. By definition, denial of service attacks are utilized to render a network resource unusable. They arise from the way different layers interact with packets.

Packets maintain a practical use of bussing all types of data around. However, these transit cycles are commonly exploited to try and attack the hardware at a given location. It is very important to be able to discern malicious activity resulting from a denial of service attack. Since this type of attack uses an exploitation of an accepted method of transit, it can be difficult to differentiate from a standard outage, especially to the end user. These symptoms include unusually slow network performance, unavailability of a particular website, inability to access any website, and dramatic increase in the amount of spam received in an account (McDowell, 2009).

To emphasize an important concept, there isn't a silver bullet when it comes to mitigating a denial of service attack. The intent of this assessment is to delineate where these attacks are being seen most frequently and the precise method of attack leveraged. As well as, provide current mitigation options that are being used in corporate environments.

## 2. Landscape

To analyze a threat, the enterprise must first become familiar with the landscape. The landscape of denial of service is risk centric. Risk will vary based on the industry sector. In reference to a report performed by RadWare in 2013 entitled, "Global Application & Network: Security Report", risk can be categorized and quantified in three major categories (Gadot, Alon, Rozen, Atad, Shulman, Shrivastava, 2013). The major categories are high, medium, and low and are indicative of the percentage of attacks directed at each respective sector. High risk organization verticals include government and the recently re-categorized financial vertical. This vertical will act as the analytical data due to its abundance. Please see Figure A for a complete representation.

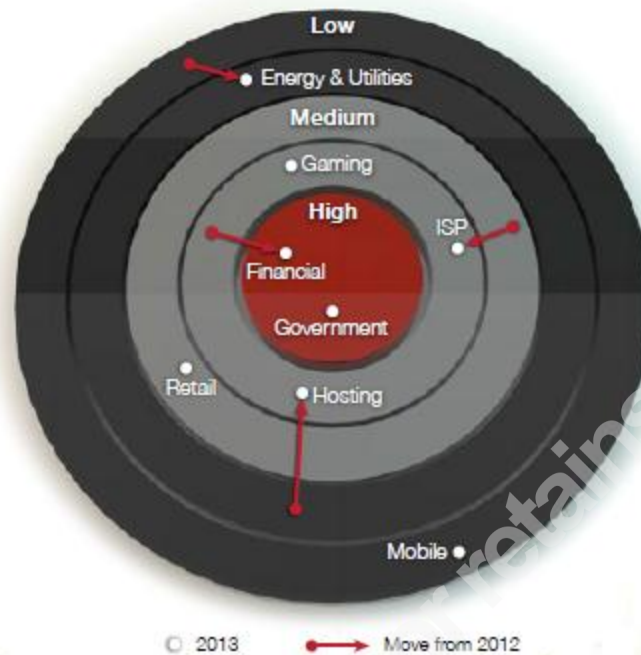


Figure A: Risk Centric Breakdown by Industry Vertical

Why are the financial and government verticals more susceptible to denial of service attempts than other verticals? A good postulation is in the nature of the data that is housed within the organization and in the manner of accessibility derived for their employees and customers. The sensitivity of the data housed can be classified as extremely sensitive including Personally Identifiable Information and Credit Card Data. Data sets of PII can include items such as Name, Address, and even items as sensitive as a social security number. For this reason, data classification is a common practice among security professionals so that appropriate safeguards may be applied.

This data will need to be accessible for business purposes in certain cases. Data that is required to traverse an external medium increases the risk of data compromise substantially. Critical infrastructure such as energy/utilities companies do not subject themselves to external threats because a breach on their end could be catastrophic and potentially endanger lives. Attacking the government and private sectors could have major detrimental effects towards national defense and the economy. For these reasons, among others, the financial and government verticals have become the most heavily attacked verticals.

### 3. Frequency of Attacks

Ryan W. Sepe, [RyanSepe@hotmail.com](mailto:RyanSepe@hotmail.com)

Denial of service attacks are happening on a daily basis and have varying levels of damage. Following the 2013 trends on denial of service compiled by RadWare, 60% of companies that were targeted by denial of service attacks experienced a service level degradation and 27% experienced an outage. Statistically, organizations that experience a denial of service attack more frequently have shorter durations of outage and degradation. These organizations are being attacked by lower level denial of service attack types. Most attacks of this ilk are not using multiple vectors and because of this are easy to remediate quickly. The frequency behind the attacks is due to the increase in detrimental attacks types towards the organization. Volumetric attacks, including packet flooding methodology, may not affect an organization that has safeguards in place to abate that method of attack. Organizations that don't employ safeguards will be subject to a higher frequency of attacks due to the ease of execution of the volumetric floods, allowing less experienced hackers to attack the network more effectively. Organizations that are subject to these attacks more should typically expect a degradation of service between 1 and 60 minutes, whereas organizations that are subject to more advanced attacks can expect a degradation of service of 12 hours or more (Gadot, Alon, Rozen, Atad, Shulman, Shrivastava, 2013)(Figure B).

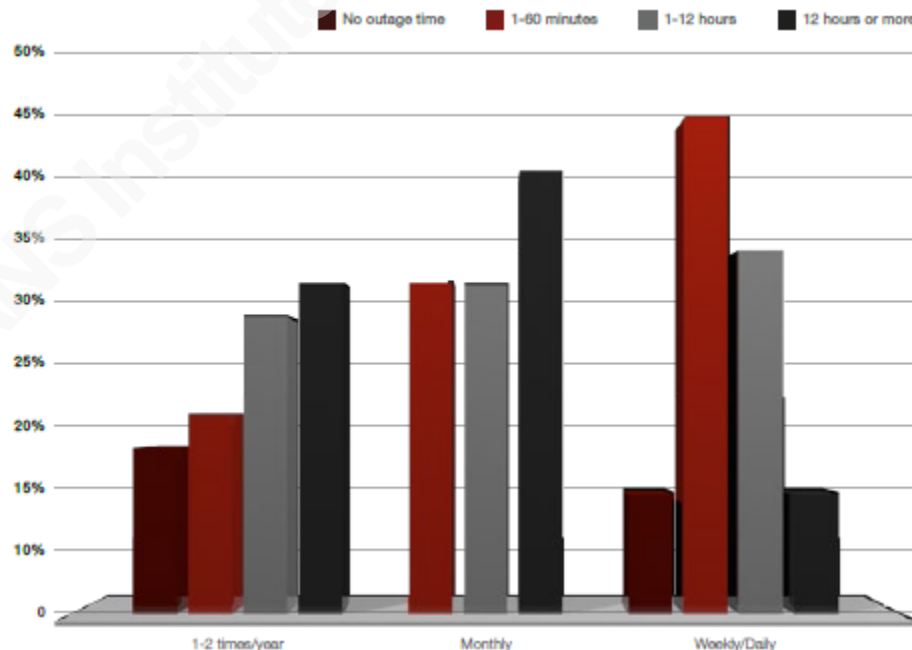


Figure B: Outage Duration in conjunction with Attack Frequency

## 4. Denial of Service Attack Techniques

Denial of service attacks have been prevalent for a long time and will continue to be prevalent due to their nature. They are versatile and in many cases very easy to execute with minimal resources. There are many methods of denial of service. It is important to understand what attacks are being executed and how they are being executed; otherwise cyber security professionals will not know why their solution is mitigating the denial of service attacks. Sections 4 and 5 will delineate recent as well as past techniques used to execute a denial of service attacks. This will to help provide a better understanding of what you would typically see as a cyber security professional and what potentially could occur in the future.

### 4.1 Volumetric Attacks

Volumetric floods can be launched by even the most inexperienced of hackers. UDP and SYN packet floods exploit the dynamic nature of packet communication to try and overwhelm enterprise hardware. User Datagram Protocol packets don't perform packet verification and because of that can be sent quickly and in excessive quantities. In 2013, UDP floods represented 7% of the attack vector landscape for denial of service attacks. This has been the case every year since 2011 and should continue this trend until other mitigation techniques are employed. SYN packet floods can leverage ACK sequences by guessing the next sequence number to spoof legitimate traffic. A volumetric attack with SYN packets can cause an extreme amount of 3-way handshakes to initiate. This interaction between the packets is very effective in overwhelming network resources and causing them to crash if the proper safeguards are not in place. SYN packet floods have been fluctuating over the span of the past 3 years. In 2013, SYN floods represented 16% of the attack vector landscape. These types of volumetric attacks can be performed with relative ease and try to pervade the first layer of a defense in depth enterprise model (Gadot, Alon, Rozen, Atad, Shulman, Shrivastava, 2013).

### 4.2 HTTP Floods

HTTP flooding attacks have a higher level of complexity as POST and GET requests can be sent to web service related hardware. The POST requests can establish multiple parameters which will trigger complex processes on these servers. Since GET requests are not able to

specify these parameters, a POST based flooding attack tends to be more effective than GET based flooding attacks as it takes fewer requests to overwhelm network assets. However, a proponent to using GET requests in opposition to POST requests is that get requests are a lot more common. Therefore, involuntary help can be leveraged from other externally facing users unbeknownst to them. This can be thought of as following a similar principle as a botnet. With legitimate GET messages being leveraged on a frequent basis, it is also plausible that these requests can go unnoticed without the proper checks in place because they are exploiting a required functionality.

### 4.3 SSL Renegotiation

SSL Renegotiation is a low-and-slow attack that renegotiates the SSL keys repeatedly consuming network resources. For its many benefits, SSL contains a substantial vulnerability as well. When looking at SSL from an outside perspective, it can be substantiated as a means of obfuscating data that traverses the wire. If authorization is in the form of the proper key, then unscrambling or decrypting the data is necessary for that data to become usable. The principles surrounding SSL or any type of encryption methodology dealing with transit is that it should be encrypting data in a secure location and decrypting from a secure location. Due to this, the risk of outside interference becomes reduced. Although the benefit to this is significant, this principle also represents a significant threat. The ability to encrypt data is not exclusive to users with genuine intent. Malicious users have the ability to encrypt data just as they have the ability to craft packets. The worst part about this is that the malicious traffic is decrypted within the network, circumventing enterprise security safeguards.

### 4.4 Search Engine Floods

Search engine floods and login page attacks represent application attack derivatives. A search engine flood attacks a web site by executing many searches and consuming network resources. This can be executed by even the most inexperienced hackers as they are simple to initiate and do not demand a myriad of physical resources. However, the speed of this type of attack is slow and may take many attempts to yield results.

A login page attack is an HTTPS encrypted resource-intensive denial of service attack targeting the login page. Web based attacks represent 27% of the attack vector landscape in the year 2013. These attacks are an accurate representation of attacks experienced in the government and financial sectors (Figure D).

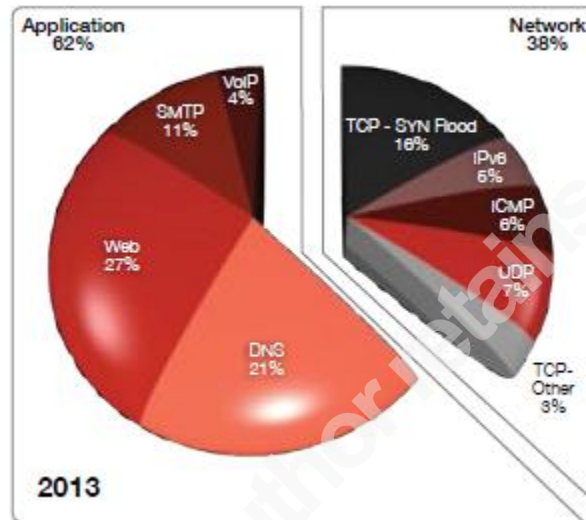


Figure D: Trend Data of Denial of Service Vectors for 2013

## 4.5 Reflective DNS

Reflective DNS attacks were carried out 83% of the time in 2013. The way to generate the most successful reflective DNS attack is through amplification and anonymity. This is leveraged by clogging the Internet pipe with requests by hiding behind spoofed DNS Servers. These requests are replied to by the genuine DNS server, subsequently launching an attack on the target victim with responses to requests that the victim never sent. The target organization does not have to run a DNS server on its own to attack the entire network. Spoofing the IP makes tracing the malicious source near impossible because the denial of service was launched not from the hacker but from a “proxy” that individual chose. This proxy could represent another valid enterprise entity or third party unrelated to the hacker in any way.

Anonymity is reserved with how the query handles certain packets. Typically, UDP is leveraged, which does not provide any type of packet checking or verification. Due to this, IP source validation is excluded. Please see the diagram of Reflective DNS Flood Architecture located in Figure E to visualize why tracing back the genuine source of the attack is increasingly difficult. This is an accurate representation of the previously described procedure.



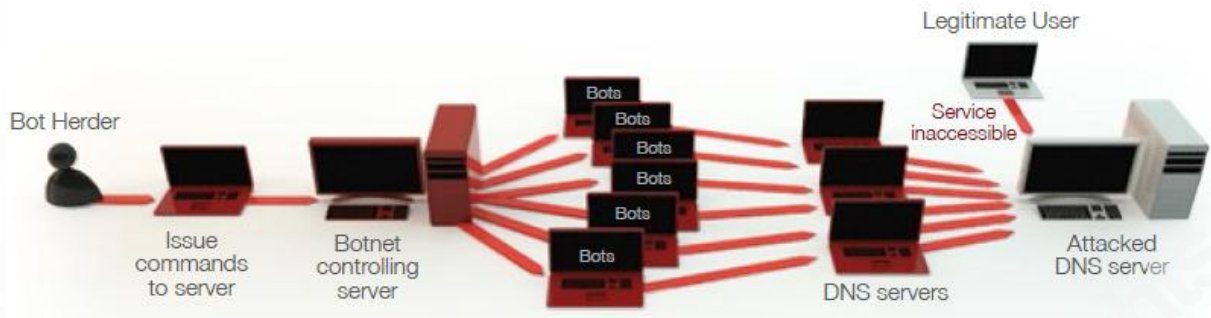


Figure E: DNS Reflective Flood Architecture

Attack amplification is another reason why the reflective DNS attack has become so prevalent. This process is performed through the use of extensions. Extensions can be thought of similar to parameters for an HTTP Post request. The specification of extensions can correlate to different amplification multipliers involving the DNS request size. The space restriction of an extension allows the attacker to send a small number of extension requests, which when sent to the server are greatly amplified. These attacks can be amplified through the following techniques: Regular DNS replies, Research DNS replies, and Crafted DNS replies.

For regular DNS replies, a normal reply is 3-4 times larger than the request. Consequently, a normal request to a legitimate cached object can result in a reply that is 4 times larger. For researched replies, hackers can study the DNS server and find out which legitimate queries can result in large replies. In some cases, the amplification factor can reach up to 10 times of the original request. For crafted replies, an attacker can compromise a poorly secured DNS server and ensure that his requests are answered with the maximum DNS reply message (4096 bytes). Using this approach an attacker can reach an amplification factor of up to 100 times (Gadot, Alon, Rozen, Atad, Shulman, Shrivastava, 2013).

## 4.6 Researching Amplification for DNS

The last two methods are a more researched method of generating amplification. By scrutinizing the DNS server with these techniques not only can a higher level of amplification be attained but a higher level of anonymity as well. Thinking of this instance from a security event management perspective, an overflow of DNS requests to a server sets up a red flag when analyzed against standard network DNS trends by the Security Operations Center. A valid point to touch upon throughout this analysis is that denial of service methodologies seek to exploit a valid means of network traffic. Reflective DNS does not differ from this principle. By

performing a more researched reflective DNS technique, a hacker can acquire crucial information in regards to the victim's DNS structure. By the time these bits of information are compromised, the attacker could launch an amplification attack that consumes network resources quickly. The Security Operations Center would be hard-pressed to efficiently mitigate an attack launched upon an enterprise that was amplified 100 fold.

## 5. Previous Attack Techniques

The above attacks focus on prevalent attacks happening in the present but it is also important to denote attacks that have been detrimental in the past. As technology changes so do the trends related to denial of service attack vectors. In the future, it is plausible that one of these attacks will resurface to cause service level degradation and even worse, outages.

Some of these attacks will encompass a technique known as Distributed Denial of Service attack. As its name denotes, distributed denial of service attacks leverage multiple systems to flood network bandwidth. This is a distributed means of methodically attacking a network with various resources. These resources may include botnets, which are a collection of zombie agents that allow an attacker to employ multiple resources to launch different attacks.

### 5.1 Smurf Attack

An ICMP (Internet Control Message Protocol) flood or smurf attack occurs when a ping command is leveraged against the broadcast address. However, this is reliant on a misconfigured device that allows packets to be sent from this address. To combat misconfigured networks, many vendors offer services that offer the ability to discover misconfigured devices on the network.

### 5.2 Crafted Packets

This next attack deals with crafted or mangled packets. Data can get quite large when being transported over the wire. This is why packet assembly is a key component to data transit. The tear drop attack is a denial of service attack that seeks to exploit packet assembly. The teardrop attack sends mangled IP fragments that have oversized and overlapping payloads to the

network resource. During packet reassembly, this attack is known to take down both Windows and Linux based operating systems (Windows 7, Vista exposed to ‘teardrop attack’).

### 5.3 P2P Attacks

Peer-to-peer servers pose a denial of service threat. For this reason, many networks block peer-to-peer connections through web security. Peer-to-peer differs in that botnets are not leveraged to carry out the attack. The hacker acts as the master controller by instructing P2P client hubs to disconnect from their P2P networks and connect to the target’s website. This results in a network, which could be several thousands of host, connecting to the target website all at once. This represents an issue because a typical web server can only a few hundred requests before the service begins to degrade. Several thousand connections will cause an instant outage.

### 5.4 Phlashing

Phlashing, also known as a permanent denial of service attack, is an attack that is detrimental to the system to the point that it requires replacement or reinstallation of hardware. In contrast with a DDoS, the security flaws exploited involve remote administration on the management interfaces of the targets hardware. The devices firmware can then be replaced with a package of the attacker’s choosing, which is most often a modified, corrupt, or defective firmware image. The process of switching out a devices firmware with a modified version is known as flashing. The intent of flashing is to get a device to be able to perform functions it was not designed to perform. A subsequent effect to flashing a device with a corrupt or defective firmware image is that the image can render to the device unusable. When this occurs the device becomes “bricked”. Another contrast between PDoS and DDoS is that PDoS is a purely hardware driven attack that requires fewer resources and, as a result, takes less time (Leyden, 2008).

### 5.5 TDoS

Telephony Denial of Service is another denial of service technique dealing with the over consumption of telephone based resources. With Voice-Over IP Technologies becoming more

prevalent, a security exploit lies inherently with the ability to create a mass amount of calls inexpensively through automated processes. Scenarios where TDoS has been experienced include false misrepresentation and fraudulent situations.

The Federal Bureau of Investigation denotes 3 examples of TDoS that they have seen. The first is the scammer misrepresenting themselves as a customer to the victim's banker/broker. The scammer will request a transfer of funds to be provided and flood the genuine client's phone with thousands of automated calls rendering the client unreachable for verification.

The second misrepresentation is in the form of a scammer claiming that the consumer has a rather large payday loan that they need to pay. When the consumer refuses the scammer will flood the victim with thousands of automated calls. In some cases, the scammer's id is spoofed to appear to be sourced from a law enforcement agency.

Lastly, the scammer contacts the victim with a bogus debt collection demand and threatens police enforcement. When the victim refuses, the attacker sends thousands of automated calls to the law enforcement agency with the source spoofed to look as if the victim was making the calls. Police will then arrive at the victim's house to investigate. In all three examples, it is easy to see that anonymity is extremely feasible through IP spoofing during TDoS attacks (Woodruff, 2011). Now that the attack techniques have been denoted, the analysis can discuss corporate data governance.

## **6. Analyzing Data Governance before Safeguard Implementation**

Before seeking plausible denial of service safeguards, it's important to note that data governance will be a monumental factor in the decision making process. In the present, cloud technology is becoming more and more prevalent due to its ease of use and lack of on-premise physical infrastructure. This is a crucial area to consider when securing the network from denial of service attacks. Many of the inquiries an enterprise would have can be fielded by the cloud provider. Key items that should be touched upon before moving forward with a cloud provider include how the security safeguards are handled. Many cloud providers do not allow external security safeguards to be placed on their cloud architecture, even when the data being supported in the cloud is owned by the enterprise. It is imperative to ensure that the cyber security team is

aware of what data they are allowed to secure based on the corresponding models. The 5 major models and their typical data governance structures can be below in Figure F.

Dedicated IT	Hosting Provider	Public IaaS	Public PaaS	Public SaaS
Data	Data	Data	Data	Data
App	App	App	App	App
VM	VM	VM	Services	Services
Server	Server	Server	Server	Server
Storage	Storage	Storage	Storage	Storage
Network	Network	Network	Network	Network

Figure F: Cloud Control Layers (Gartner, 2014)

Key
Organization has control
Organization shares control with service provider
Service provider has control

## 7. Denial of Service Mitigation Solutions

Now that each of the attack techniques has been well documented and data governance concerns delineated; the following sections of the assessment will cover mitigation options. Each of the denial of service type attacks launched against the highest risk vertical has developed different mitigation techniques. In addition to a description of these techniques and how they function, vendors who support these options will also be provided in the following sections.

### 7.1 Option 1: Scrubbing Service

The first major mitigation technique to denial of service is provided in the form of a scrubbing service. In this way, traffic is diverted away from the internal network towards a third-party who will scrub the traffic to analyze it for denial attempts. Prolexic is an example of a company that performs these services. Services such as these are beneficial because services typically come with a support/analysis team. These terms will be delineated in the service

contract. Scrubbing center services provide a dedicated Security Operations Center to analyze the traffic along with sensors and detection engines.

A detriment to this method lies in the data architecture. Since this service will sit between the Internet pipe and the border router if the behaviors and anomalies are not detected early enough, then the attack may clog the Internet pipe to the organization. This occurrence is also dependent of the bandwidth provided to the enterprise Internet pipe. The less bandwidth provided the greater risk of outage or service degradation (The best on-demand, cloud-based scrubbing centers for DDoS protection).

## 7.2 Option 2: Mitigation at the Internet Pipe

The second is reliant on the Internet Service Provider. Denial of service attacks are detected at the Internet pipe by the enterprise ISP. Sakura Internet, one of Japan's largest ISP's performs active denial of service monitoring. This allows for quick analysis and mitigation of a potential attack. The developed technology ingests massive amounts of IP traffic and performs in-memory analytics to identify and stop DDoS attacks on the network as they happen. This process will simultaneously enable legitimate traffic to continue. This is done by leveraging a high speed NewSQL database. The databases have the capability to analyze 48,000 IP packets per second. The benefit to this is that inspection is performed in real-time which allows for mitigation techniques to be employed before service degradation or outage. In April of 2014, during the first month of its inception for the NewSQL database named VoltDB, Sakura detected and mitigated 60 DDoS attacks while also successfully restoring legitimate traffic to a majority of targeted websites during the time of attack. Of these 60 cases, 49 were able to have their services restored quickly, providing uptimes of 20 seconds in certain cases (Zorz, 2014).

The mechanics behind analyzing mass quantities of data in real time resides in Big Data principles. Network architecture is monumental in deciding which avenue to pursue for denial of service mitigation. This is why it behooves organizations that are looking to employ a denial of service mitigation option, to not only have a Cyber Security Engineer on the meetings but a Network Architect.

The functionality of packet translation should be embedded within network hardware such as routers. Many types of traffic are transmitted over the wire, but in this configuration it is

crucial that inspection of these packets happens quickly and efficiently. From this principle, packet translation to the UDP protocol is a powerful means of ciphering through data. After this translation occurs, the traffic is then sent to the designator which is essentially a database client. The purpose of the database client is to run processes so that it can identify the flow of malicious traffic on the spot. These flows are aggregated based on their destination address. Another agent will aggregate the traffic flows based on the source IP. With this process, enterprises are able to ensure that genuine traffic is allowed through while traffic with malicious intent is “blackholed” (FortiNet and Radware).

### 7.3 Option 3: Appliance Based Mitigation

The third root mitigation option is appliance based mitigation. This mitigation technique leverages an inline appliance placed in front of the ingress of the Internet or Border router. Companies such as RadWare with DefensePro and FortiNet with FortiDDoS are examples of entities that leverage appliance based denial of service mitigation techniques.

These appliances rely on behavior-based detection as well as DDoS processors to detect threats. This allows expedient inspection of traffic with a latency rate that is calculated in microseconds. Some solutions, such as DefensePro, offer an Out-of-Path option with their appliances that allows for the traffic to be sent to a scrubbing center instead of being inspected inline. A benefit to this is that mitigation can be performed on an ad-hoc basis. When malicious traffic is diagnosed, the out-of-path solution can spring into action to provide mitigation. This provides minimal latency because intense inspection is not performed outside the network before ingress. It is important to note that different configurations suit different needs. It’s up to the enterprise to determine what solution best works best for them.

### 7.4 Option 4: Mitigation with Current Enterprise Hardware

Now that cloud architecture has been considered, the next aspect to focus on is the implementation of current on-premise security safeguards. Unfortunately, it is plausible that a denial of service project may not be on the road map due to other endeavors. For this reason, it is important to know what security professionals can do to leverage their current enterprise hardware to minimize the risk of a denial of service attack.

### 7.4.1 SOC and NIPS

The first major security safeguard that is common amongst enterprises is a Network IPS. The reason that the IPS is equipped to handle volumetric attacks from lower level packets is because network traffic is analyzed beforehand in the refinement process when setting up an IPS solution. A monumental increase in this traffic from a UDP/TCP perspective is indicative of a malicious attempt. The IPS accounts for this and blocks traffic before ingress.

This is not the only benefit to acquiring an IPS solution. One of the strongest proponents for implementing an IPS is the anomaly based analysis and manual signature upload. HTTP floods can vary due to a widespread amount of parameters. The enterprise is able to diagnose the traffic from HTTP floods manually with their IPS and the SOC by creating ad hoc signatures for each attack. In this way, the on-premise IPS can mitigate search engine floods, SSL renegotiation, and many HTTP flood attempts.

However, this is a day to day process. Once the malicious party discovers that their attack has been blocked repeatedly, they can diagnose the event to deduce that a rule was created based on their attack signature. A variation of the attack can be crafted which renders the previous architecture rule-set ineffective. An on-premise IPS differs from a cloud-based IPS in that the on-premise solution is connected to the hardware architecture outside of the firewall. A cloud-based IPS is architected in a manner that supersedes the entire cloud platform.

The login page attacks are mitigated by on-premise challenges. Once challenge technology was implemented in correlation with on-premise IPS, HTTP flood attempts were blocked completely. Most login pages utilize HTTPS because it offers a level of encryption. This encryption method helps keep user credentials secure and promotes site use. To mitigate login page attacks, integration with SSL technology was needed to introduce HTTPS Web Challenges.

Application, Encrypted, as well as Low and Slow type attacks will sneak through the cloud based security safeguards. The mitigation technique here lies in the form of protocol checks and signatures. IPS uses behavior and anomaly based reasoning to search for an influx in traffic that is not base lined in the IPS rules as being standard for the organization. When this occurs, the traffic can be quarantined until it is reviewed or can be set to default deny as all new streams of traffic should be approved by the organization and tested thoroughly before being placed into production. Depending on the frequency at which the enterprise is targeted this can become a cumbersome task. Many organizations will not have enough Security Team bandwidth



to be strictly monitoring traffic flows as there are other aspects that acquire the attention of the Cyber Security Team.

## 7.4.2 Firewall

The firewall is the next piece of infrastructure that an enterprise will typically already have deployed. Organizations that have a well-rounded security posture will set up the firewalls in a Deny/Deny configuration. This signifies that any traffic that is not explicitly whitelisted will be denied ingress. As stated with the IPS, organizations should have an approval process for what traffic they allow. New business will need to go through the change control process and undergo stringent tests to ensure that the traffic has the ability to be routed through the firewall and maintains data integrity. This process is indicative of change management best practices as well as data security best practices. In terms of typical corporate architectures, firewalls are the enterprises final lines of defense.

Any of these mitigation solutions being successfully denial of serviced will essentially render the other safeguards ineffective. The reason it is best to try an incorporate a comprehensive defense in depth approach is because the attack vectors are evolving. The denial of service attacks that are being launched contain multiple attack vectors so having mitigation techniques against as many vectors as possible decreases the risk of a denial of service against the enterprise substantially. Figure G is a representation from Radware of the denial of service data from 2013 in terms of the amount of attack vectors used.

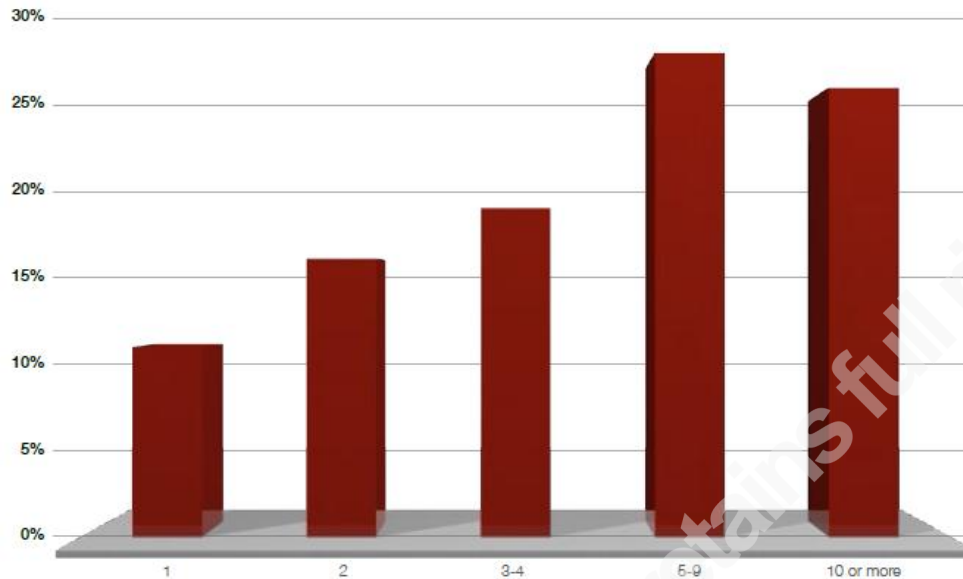


Figure G: Count of Attack Vectors in 2013 Denial of Service Attacks

## 8. How to Set Up a Security Framework to Combat DoS

Thus far, this analysis has been an in-depth review of denial of service landscape, attack vectors, and mitigation techniques. The following is a phased approach applicable to all enterprise sizes looking to implement a denial of service solution.

### 8.1 Phase 1: Network and Cost Analysis

The First Phase of planning out a denial of service solution is to perform a Network and Cost Analysis. This involves the coordination of many teams. A project manager should be charged with the denial of service task and select a team leader from each respective department. A robust network analysis is imperative to determine what devices are vulnerable, as well as what are the most practical solutions conducive to the enterprise network. Once a thorough analysis has been performed, a cost analysis should be conducted. The main reason for conducting a cost analysis is to determine how much capital the business can allocate to the current endeavor. Depending on what the raw number is will depend on if a fully comprehensive solution needs to be implemented in a phased approach as well. The raw number should also be indicative of man hours needed to implement the solution.

## 8.2 Phase 2: Research

The Second Phase is the Research Phase. The project manager, with the team leads, should conduct meetings with a subject matter expert on denial of service trends, attack techniques, and mitigation/remediation methodologies. An example of a Subject Matter Expert provider is Gartner, Inc. Gartner customers have the ability to block off a 30 minute time allotment with an analyst whose expertise is that one subject. They will be able to analyze needs and delineate options specific to the precise enterprise architecture. Being a Gartner customer also provides the ability to peruse a vast archive of materials. Once feasible options have been acquired through research and constructive conversation between the SME and the enterprise team leads, the project manager can then provide the results to upper management. Once upper management decides on a direction then the implementation phase can start. An important point to remember is that different solutions may suit an enterprise differently. A solution that is best for one enterprise may not fit in another. The solution will depend on many variables but this should be definitively worked out between the SME's and team leads before Phase 3 (Gartner, 2014).

## 8.3 Phase 3: Implementation

The Third Phase is the Implementation Phase. This occurs after the Request for Purchase process has been completed and the enterprise owns the solution. The vendor who provides the solution for denial of service mitigation will have a roadmap of installation steps and prerequisites. It would be beneficial to acquire a personal timeline for the enterprise network and, historically, what vendor completion times have been for other customers. Resource allocation also needs to be calculated. It will be the job of the project manager to ensure that the organization has the appropriate bandwidth to complete the implementation in the set time span. The amount of time to complete implementation will vary based on series of factors such as bandwidth, network complexity, and hardware acquisition/ownership.

## 8.4 Phase 4: Refinement

The Fourth and final phase is Refinement. This should be performed proactively as well as reactively. From a proactive standpoint, before the solution is pushed into production it should

be tested with the current enterprise workflow in a test environment. Pushing a non-tested solution into production can restrict functionality.

Once the rules have been refined and it has been determined that the newly acquired solution will not restrict genuine workflow, the solution may then be pushed into production. This phase however is infinite. It should never cease to occur because new business and new solution functionalities are inevitable. The project manager's last task for the project is to ensure that a regular monitoring procedure is instantiated to correspond to business policy. Any changes that need to be made to the solution should be provided in a new change management flow that fits in with the current service structure at the enterprise. This is known as reactive refinement in the context that changes to be made are indicative of other imposed changes.

This four phased approach is opinion based deduced from logical reasoning on how a denial of service project should be managed based on copious research. SME's may provide various methods of performing this approach as each individual prefers certain implementation methods.

Denial of service attacks are a major issue and similar to other issues there are a few things to consider. This analysis seeks to outline the problem by delineating where denial of service is most prevalent, what the detriments are, and how these attacks are manifested. Then, discussing contributing factors that will be crucial when determining how an enterprise would like to proceed. Finally, denial of service mitigation options that are currently viable in the current security space have been outlined. This will provide an informed view of how a security professional can set up a denial of service mitigation framework and what tools can accomplish this task.

With many endeavors in the security realm however it's not a matter of is it beneficial to the company to implement but rather is it feasible and where this endeavor would rank in terms of value. There are other projects that will be on the roadmap and it is imperative that you provide value to security ones such as denial of service mitigation as businesses tend to look at projects from a financial perspective. The last part of this paper will cover the financial benefit to implementing a denial of service mitigation framework and will show statistically why saving money can be just as beneficial as making it.

## 9. Displaying Value to Non-Cyber Security Oriented Professionals

All security professionals understand the fundamental principles behind denial of service attacks. One of the more difficult pieces of implementing the phased approach above is delineating the value to the business side of the corporate environment. As it is not always so clear cut as to what endeavors should take precedence. For this reason, the phased approach above may take time to get on the corporate roadmap if the importance isn't accurately portrayed. The below section denotes strong pieces of data advocating the implementation of a denial of service solution from a financial perspective.

The Ponemon Institute denotes that data center outages in 2013 cost approximately \$200,000 more than in 2012. If this trend continues, 2014 would eclipse 1 million dollars in cost per year that results from the data center outages. The activity based costing model partitions cost in terms of Activity Centers in correlation with Cost Consequences (Ponemon Institute, 2013)(Figure C). The financial detriment can include but is not exclusive to equipment, information technology and user productivity, third party involvement, lost revenue, and business disruption. The activities that correlate with the cost correlation are heavily involved in the incident response process such as detection, containment, recovery, and post-event response. Costs vary based on the degree of the outage which is symptomatic of organizations ability to fend off the attack launched against them and duration of incident handling process.

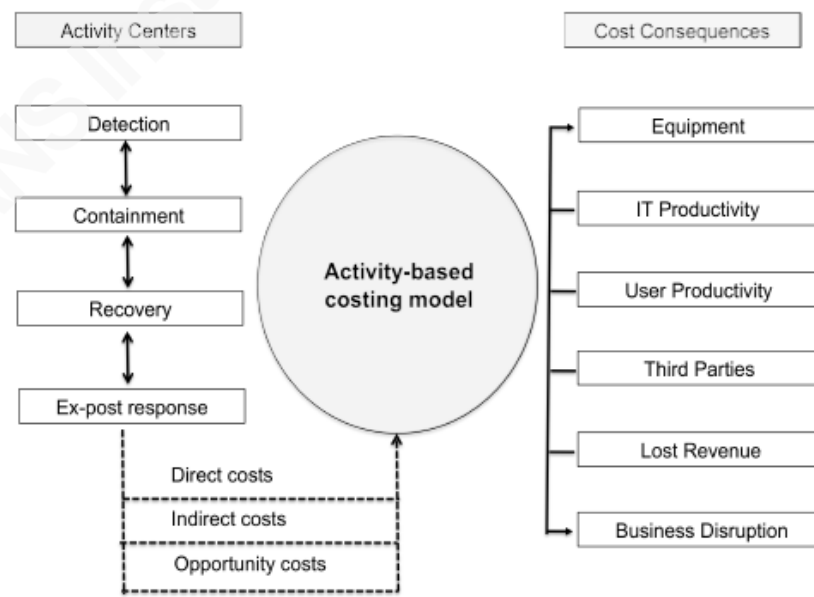


Figure C: Activity-based on account framework

From a business perspective, the detriments can extend even further past direct financial concerns. Denial of service attacks may correlate to reputation loss, direct financial loss due to outage or slowness of Internet, service level agreement compromise, and impact to internal organizational processes. Reputation loss has the potential to hinder financial status, not only in the present but in years to come. Considering publicly known attacks in the year 2014, not exclusive to denial of service, it is apparent that companies are under major scrutiny to amend for these types of events.

Sony's Gaming Platform was successfully attacked via a denial of service attack causing outages for many of the online platforms that it provides on December 7<sup>th</sup>, 2014. This is the second attack they received of this ilk with the first happening in August of 2014. Though the same group, Lizard Squad, also successfully brought down the Xbox Live Network a week earlier than Sony's first attack, Sony's was more publicly chastised. A probable reason for this is because Sony's reputation already had a smudge on their record a few years earlier for having millions of customer's credit card data compromised through their Play Station Network platform. These events cause the company to attain a poor reputation when trying to keep their customers data secure and can directly correlate to a loss of sales from the responsible consumer. These strong statistics help the business side assess the monetary value of implementing a denial of service solution and will greatly increase the chances of getting a solution on the roadmap.

## 10. Conclusion

Unfortunately, denials of service are becoming more complex as different mitigation techniques are unveiled. It is difficult to produce a comprehensive solution because these attacks exploit the genuine function of the network hardware and data transit. It may come down to implementing direct line principles where only point to point connections are available to provide an employee with the assets they need. In a way, this is what scrubbing centers are providing; a "clean" line of transit for data packets between source and destination. The core idea of the Internet, ease of use and availability of data, is its greatest benefit as well as detriment. However, understanding the principles denoted in this analysis and applying them to the corporate enterprise can help put companies ahead of the game.

## References

- The best on-demand, cloud-based scrubbing centers for DDoS protection. (n.d.). Retrieved December 30, 2014, from <http://www.prolexic.com/why-prolexic-best-dos-and-ddos-scrubbing-centers.html>
- DDoS Resources. (2001). Retrieved December 30, 2014, from <http://web.archive.org/web/20100914222536/http://anml.iu.edu/ddos/types.html>
- Fortinet. (n.d.). Retrieved December 30, 2014, from <http://www.fortinet.com/products/fortiddos/>
- Gadot, Z., Alon, M., Rozen, L., Atad, M., Shulman, Y., & Shrivastava, V. (n.d.). Radware: Global Application & Network (Vol. Security Report 2013).
- Gartner, Inc (Analyst Meetings: Precise Contents Confidential)
- Leyden, J. (2008, May 21). Phlashing attack thrashes embedded systems. Retrieved December 30, 2014, from <http://www.theregister.co.uk/2008/05/21/phlashing/>
- McDowell, M. (2009, November 4). Security Tip (ST04-015). Retrieved September 4, 2014, from <https://www.us-cert.gov/ncas/tips/ST04-015>
- Ponemon Institute. (2013). 2013 Cost of Data Center Outages.
- Radware, DDoS Protection and Attack Mitigation. (n.d.). Retrieved December 30, 2014, from <http://www.radware.com/Products/DefensePro/>
- Software Engineering Institute (SEI). Denial of Service Attacks. (n.d.). Retrieved September 4, 2014, from [http://www.cert.org/historical/tech\\_tips/denial\\_of\\_service.cfm](http://www.cert.org/historical/tech_tips/denial_of_service.cfm)
- Windows 7, Vista exposed to 'teardrop attack' (2009, September 8). Retrieved December 30, 2014, from <http://www.zdnet.com/article/windows-7-vista-exposed-to-teardrop-attack/>

Woodruff, B. (2011, June 20). Phony Phone Calls Distract Consumers from Genuine Theft.

Retrieved December 30, 2014, from <http://www.fbi.gov/newark/press-releases/2010/nk051110.htm>

Zorz, M. (2014, September 10). How a large ISP fights DDoS attacks with a custom solution.

Retrieved December 30, 2014, from <http://www.net-security.org/secworld.php?id=17347>





# Upcoming SANS Training

[Click Here for a full list of all Upcoming SANS Events by Location](#)

SANS San Francisco Winter 2017	San Francisco, CAUS	Nov 27, 2017 - Dec 02, 2017	Live Event
SIEM & Tactical Analytics Summit & Training	Scottsdale, AZUS	Nov 28, 2017 - Dec 05, 2017	Live Event
SANS Khobar 2017	Khobar, SA	Dec 02, 2017 - Dec 07, 2017	Live Event
SANS Munich December 2017	Munich, DE	Dec 04, 2017 - Dec 09, 2017	Live Event
European Security Awareness Summit & Training 2017	London, GB	Dec 04, 2017 - Dec 07, 2017	Live Event
SANS Austin Winter 2017	Austin, TXUS	Dec 04, 2017 - Dec 09, 2017	Live Event
SANS Frankfurt 2017	Frankfurt, DE	Dec 11, 2017 - Dec 16, 2017	Live Event
SANS Bangalore 2017	Bangalore, IN	Dec 11, 2017 - Dec 16, 2017	Live Event
SANS Cyber Defense Initiative 2017	Washington, DCUS	Dec 12, 2017 - Dec 19, 2017	Live Event
SANS Security East 2018	New Orleans, LAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SANS SEC460: Enterprise Threat Beta	San Diego, CAUS	Jan 08, 2018 - Jan 13, 2018	Live Event
SEC599: Defeat Advanced Adversaries	San Francisco, CAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Amsterdam January 2018	Amsterdam, NL	Jan 15, 2018 - Jan 20, 2018	Live Event
Northern VA Winter - Reston 2018	Reston, VAUS	Jan 15, 2018 - Jan 20, 2018	Live Event
SANS Dubai 2018	Dubai, AE	Jan 27, 2018 - Feb 01, 2018	Live Event
SANS Las Vegas 2018	Las Vegas, NVUS	Jan 28, 2018 - Feb 02, 2018	Live Event
Cyber Threat Intelligence Summit & Training 2018	Bethesda, MDUS	Jan 29, 2018 - Feb 05, 2018	Live Event
SANS Miami 2018	Miami, FLUS	Jan 29, 2018 - Feb 03, 2018	Live Event
SANS London February 2018	London, GB	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Scottsdale 2018	Scottsdale, AZUS	Feb 05, 2018 - Feb 10, 2018	Live Event
SANS Secure India 2018	Bangalore, IN	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Southern California- Anaheim 2018	Anaheim, CAUS	Feb 12, 2018 - Feb 17, 2018	Live Event
SANS Dallas 2018	Dallas, TXUS	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS Secure Japan 2018	Tokyo, JP	Feb 19, 2018 - Mar 03, 2018	Live Event
Cloud Security Summit & Training 2018	San Diego, CAUS	Feb 19, 2018 - Feb 26, 2018	Live Event
SANS Brussels February 2018	Brussels, BE	Feb 19, 2018 - Feb 24, 2018	Live Event
SANS London November 2017	OnlineGB	Nov 27, 2017 - Dec 02, 2017	Live Event
SANS OnDemand	Books & MP3s OnlyUS	Anytime	Self Paced