



access

////////////////////////////////////
/// **DEFENDING AGAINST** //////////////////////////////////
/// **DENIAL OF SERVICE** //////////////////////////////////
////////////////////////////////////

Denial of Service is a complex and context-specific problem. This document aims to cover some of the common scenarios and provides guidance for prevention and response. It is not designed to mitigate all attacks or be the definitive guide to all instances of Denial of Service. There are new types of attacks every day and many alternative solutions. This document has been peer reviewed by a number of individuals and organizations prior to release and is constantly being updated. We encourage comments, suggestions, or information which will improve this guide and ensure that it is as up to date, relevant and as accurate as possible (email soc@accessnow.org). If mitigation of an attack is beyond your capability or the scope of this document, please contact trusted technical assistance or email soc@accessnow.org. Visit our website at <https://www.accessnow.org>.

January 2012

DEFENDING AGAINST DENIAL OF SERVICE



For more information,
please contact:
soc@accessnow.org
PGP Key ID: 0xF08D380A

INTRODUCTION

Civil Society currently faces significant cyber threats. At the top of the list of those threats are Denial of Service (DoS) attacks. The websites of many organizations and individuals have already come under such attacks, and the frequency of those attacks are on the rise. Civil Society frequently does not have the kinds of resources or technical know-how that is available to commercial enterprise and government websites, and often have to exist in adverse political environments where every avenue available, both legal and illegal, is used against them. Therefore, the threat of DoS attacks is unlikely to go away any time soon.

A Denial of Service (DoS) attack is any attack that overwhelms a website, causing the content normally provided by that website to no longer be available to regular visitors of the website. Distributed Denial of Service (DDoS) attacks are traffic volume-based attacks originating from a large number of computers, which are usually compromised workstations. These workstations, known as 'zombies', form a widely distributed attack network called a 'botnet'. While many modern Denial of Service attacks

are Distributed Denial of Service attacks, this is certainly not true for all denials of service experienced by websites. Therefore, when users first start experiencing difficulty in getting to the website content, it should not be assumed that the site is under a DDoS attack. Many forms of DoS are far easier to implement than DDoS, and so these attacks are still used by parties with malicious intent. Many such DoS attacks are easier to defend against once the mechanism used to cause the denial of service is known. Therefore, it is paramount to do proper analysis of attack traffic when a site becomes unable to perform its normal function.

There are two parts to this guide. The first part outlines preparatory steps that can be taken by Civil Society organizations to improve their website's resilience, should it come under attack. However, we do understand that most Civil Society organizations' first introduction to DoS attacks comes when they suddenly find themselves the victim of an attack. The second part of this guide provides a step-by-step process to assist the staff of NGOs to efficiently deal with that stressful situation.

MITIGATION STRATEGIES AND PHASES TABLE

STRATEGIES	PREPARATORY PHASE	INITIAL PHASE	RESTORATIVE PHASE	LONG-TERM PHASE
HOSTING LOCATION	X		X	X
DoSP	X	X*		
FIREWALLING		X		
OPEN PROVIDER MIRRORS	X	X	X	
LOAD BALANCED MIRRORS	X		X	
BIG BANDWIDTH MIRRORS	X			
DISTRIBUTED CONTENT	X	X**		
DARKNET	X	X**		
FORGOING THE DOMAIN		X		

* Enact here if pre-prepared or determine if provider or provider's provider has DoSP in place

** If set up in advance (Preparatory Phase), then this mitigation strategy is enacted in this phase



PART 1: PREPARATORY PHASE

Most DoS mitigation strategies are best implemented well ahead of time. Some strategies, such as the various types of site mirroring, are implemented as the normal operating procedure whether the site is under attack or not. These strategies are architectures for your website that are good for reach and performance, as well as for resilience to DoS attacks. Other strategies are organized and lie dormant, waiting to be activated if the site comes under DoS attack. Darknets are certainly in that category, and the use of Distributed Content as a mitigation strategy (rather than a strategy for reach) is as well.

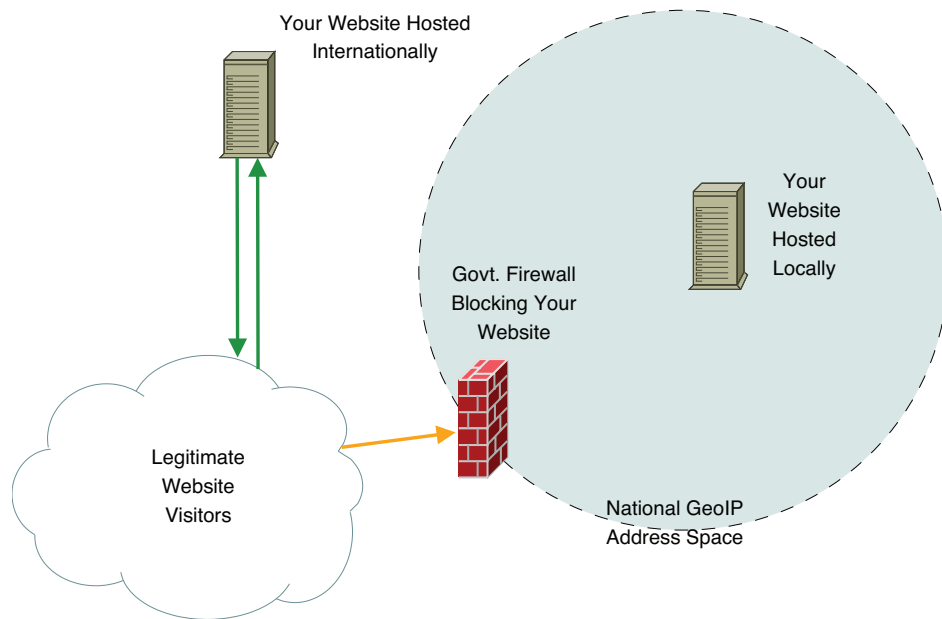
HOSTING LOCATION

The capabilities, geographic location, and the appetite to tolerate attacks by your website hosting provider may be a significant factor in the ability your opponents have in affecting your site. Therefore, it is advised you take careful consideration in choosing the hosting provider or providers of your website carefully in advance.

Denial of Service attacks are effective by exhausting your resources, so it is wise to ask yourself how much resources could you commit to an attack before you got to the point of resource exhaustion brought your site to a standstill. Think of all the resources involved including (but not limited to) hardware and software application limits, network and bandwidth limits, engineering and technician time, and financial costs.

Ask questions of your potential providers: how much bandwidth is dedicated to your site? How much bandwidth could be consumed by an attack, and how many days would the provider tolerate an attack before refusing to host your site? Is your site able to be moved out of a shared virtual environment and onto a dedicated server if it comes under attack (to protect other customer's websites from being affected by the attack against your site)? What Denial of Service Protections do they have in place and/or offer to customers?

This is also true of the provider hosting the DNS for your site. Ask them how many DNS requests per second they would tolerate before refusing to respond to name service requests.



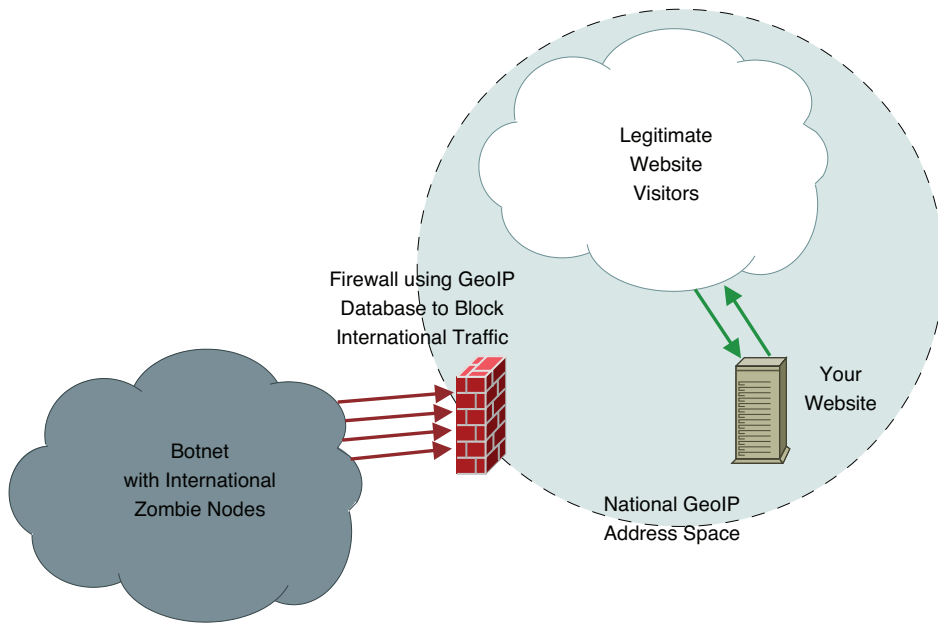
International Hosting Location Strategy

INTERNATIONAL HOSTING LOCATION STRATEGY

If you oppose the government or are highly critical of other sectors within your country, then it is probably not a good idea to host your website within the geographic borders of your country. Some political opposition groups have come under DDoS attack, and when they asked their service provider to assist in mitigating the attacks, they received no assistance whatsoever. It is believed the service providers were under pressure from the government to withhold any assistance to the website owners. It must also be noted that if the government

has control of traffic at key points of the internet infrastructure within the borders of the nation, then it may be possible for them to outright block access to, divert traffic away from, or choke down the bandwidth to your website, creating a very effective denial of service to some visitors to your website (depending on what infrastructure the government has control of, where it is located, where your users are located, etc.).

Another issue to consider in relation to geographic location of your hosting provider is the question of legal jurisdiction. If you host your website within the nation or in a nation with very close ties to the government, you may provide them with an

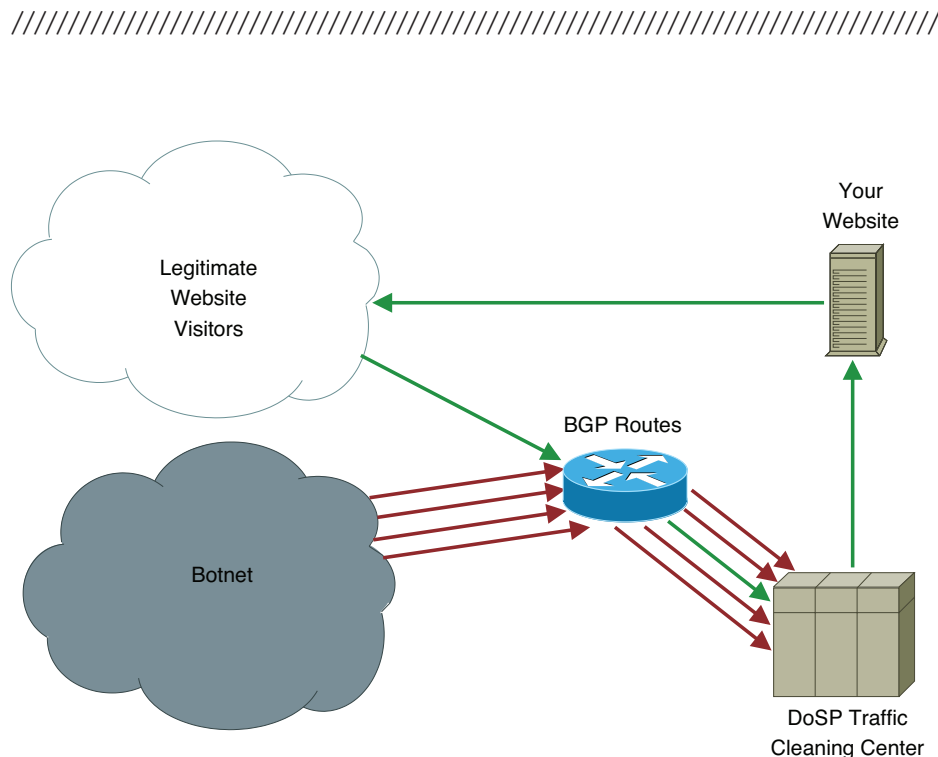


Local Hosting Location Strategy

easy avenue to find or create laws to shut down your site. The solution to these issues is to host your website internationally, away from the infrastructural control of the local authorities, and in a jurisdiction that is politically neutral to the content of the site.

LOCAL HOSTING LOCATION STRATEGY

There may be good reason not to implement the International Hosting Strategy (as per above), so careful consideration should be given to the best solution that meets your specific situation. If your website audience is entirely local to your country, and there is enough incentive for the authorities to hold back from blatant censorship of your website, then it may be a better strategy against DDoS attacks to host your site within the borders of your country. If your opponents engage the services of a botnet on the black market to attack your website, it is highly likely that the majority of the zombie nodes that make up that botnet will be located outside of your country. It is therefore possible to go some way to mitigating a DDoS attack by simply dropping all traffic at an upstream firewall that originates from computers not located within your country. This leaves your website available for your local target audience. To differentiate between national and international connections geolocation will need to be done either at the firewall or a forward proxy. Some firewalls and proxies will come with this functionality out of the box and rules can be written for others to achieve this capability. Note also that the effectiveness of this strategy may be limited if the attack utilizes spoofed source addresses.



Third Party Denial of Service Protection Strategy

THIRD PARTY DENIAL OF SERVICE PROTECTION (DoSP)

When choosing a hosting service provider, take into careful consideration any offerings they may have for the provision of Denial of Service Protection (DoSP). DoSP can take

many forms, but usually the most cost-effective options will be organized by your hosting provider, so choose carefully. Some basic DoSP options are guarantees made about short-term increases in the provision of bandwidth to your site to cope with rate-based denial of service attacks. At the other end of the scale are global providers who specialize in this area and have agreements with most of the world's large Telcos may provide stronger DoSP, as these agreements may allow them the ability to 'clean' incoming traffic or stop attacks closer to their source rather than waiting until the attack traffic reaches the network where your website is hosted. DoSP cleaning of the traffic destined for your site can be done at your hosting location, or with the traffic re-routed to a traffic cleaning centre by re-routing your traffic with altered BGP routes or use of reverse proxy and DNS re-announce. Arbor Networks provides this kind of service (<http://www.arbornetworks.com/>). There are also DoSP 'middlemen': companies that offer DoSP services from a number of higher-level DoSP providers bundled together as a single service, which is then offered to a large number of clients. These middlemen often are able to provide the best DoS protection at a very reasonable price. DOSarrest is one example of a middleman DoSP provider (<http://www.dosarrest.com/>).

PREPARING MIRRORS

The aim of website mirroring is to distribute content across multiple service providers, multiple geographic locations, and multiple allocations of bandwidth, so that the website content has some resilience to disruptive events, whether from localized natural or political events or malicious activity. Building the website so it can be bundled, moved, installed, and reinstalled easily is a great strategy for preparing to mitigate DoS attacks through establishing mirror sites. In order to put your website together like this, you will need to take into consideration the technologies used to build and run your site. Answering the following questions will assist in preparing a site for mirroring:

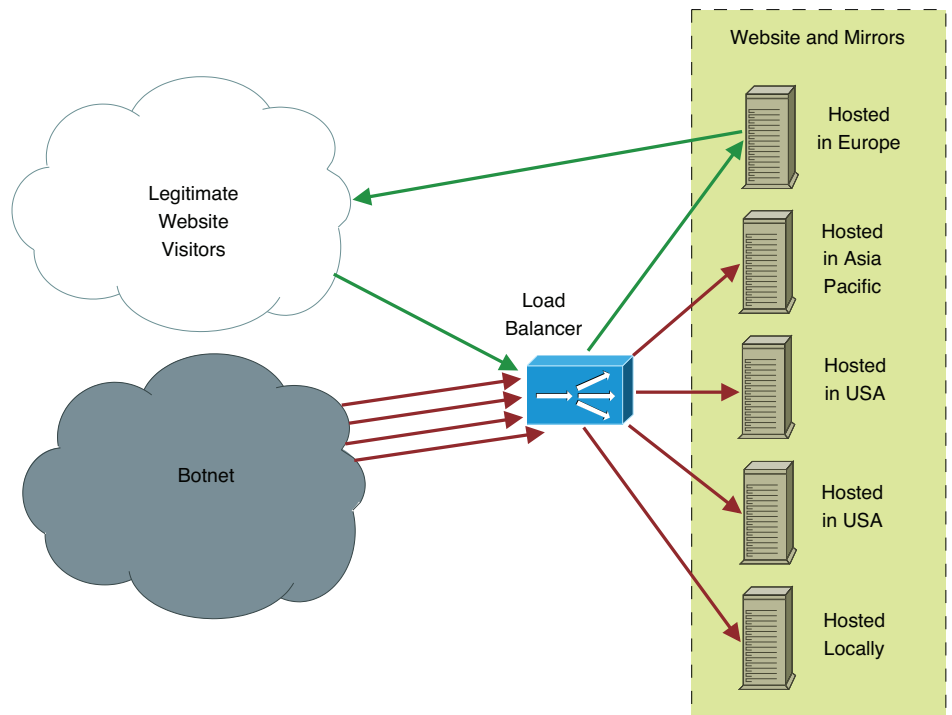
1. Will the technologies run on differently configured platforms?
2. What software dependencies are there?
3. Can all those dependencies be isolated and bundled with the website content?
4. If the site is data-driven, are you able to bundle that data into the same package?
5. Will the site run with snapshots of the data, or is the data so real-time that it cannot be pre-bundled (even if on a nightly basis or similar)?

Another challenge with preparing for mirroring is finding a way to test mirror creation under the kinds of circumstances that would be experienced in the heat of an ongoing DoS attack. Ideally, a site administrator would have accounts with multiple hosting providers (for the purpose of hosting mirrors when required), and

these accounts would be used to test the ability of the site's tech team to deploy, install, and verify the installation of an instance of the site onto the test server. In this way, issues could be ironed out before the procedure has to be done under the more extreme circumstances of a real attack. This test process should not be performed once and then forgotten about. It should be a regularly scheduled event, as the environment, the components that make up the website, and the data that sits behind it all change frequently.

LOAD BALANCED MIRRORS

Load balanced mirroring is possible by re-delegating the DNS for the website domain to point to a device that distributes incoming requests across a number of mirrored instances of the website. Ideally, these instances of the website should be distributed to different hosting providers located in different geographic locations. Load balanced mirroring of a website



Load Balanced Mirrors Strategy

can be set up in preparation of possible DoS attacks, but due to the additional complexity of setting up the load balancer and the expense of paying multiple service providers to host multiple instances of the website, this is often avoided until an active DoS event is initiated. To be able to use this type of mirroring, the credentials and information necessary to modify the DNS records for the domain are required. So is the ability to either host a load balancer technology with administrative rights, or to add configuration to a load balancing device that is provided as part of a service by one of the hosting providers. Examples of such load balancing devices include a whole range of options, from simple, cheap and open source reverse caching proxies (Nginx, Lighttpd, Varnish etc which run on a variety of operating systems such as Linux, BSD, and Microsoft), to proprietary, commercial and expensive load-balancing solutions (F5 routers, etc).

OPEN PROVIDER MIRRORS

One factor that makes Open Provider mirroring so popular amongst Civil Society site owners is that the services offered by the open providers are usually free or low-cost. An open provider mirror for a website is essentially a content publishing framework provided as a service, such as Wordpress or Facebook. These providers are good for defending against DoS because they have access to truly massive amounts of total bandwidth resources, which they seem happy to use liberally when an account in their system comes under significant rate-based DoS attack. The disadvantage is that these frameworks were designed to lower the barrier for average people to publish content, and therefore they provide fairly generic and limited scope for the design and the delivery of dynamic content. The aim of using an open provider mirror is to host the most critical

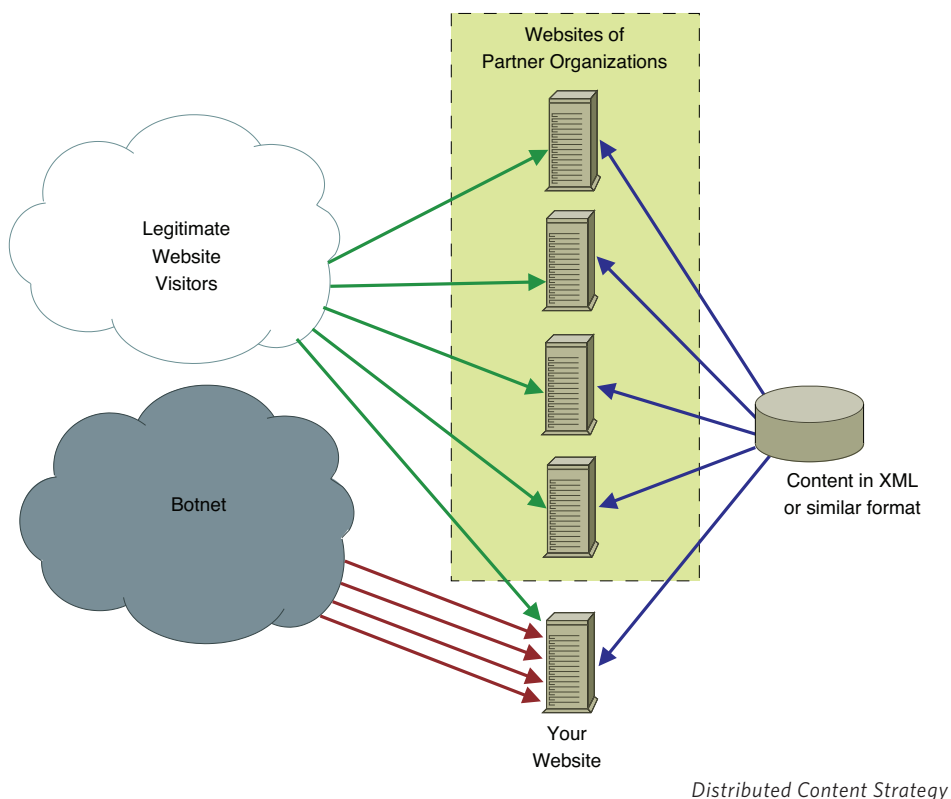
content of your original site elsewhere in a more static form. In preparation for this type of mitigation before an event occurs, the architects of the site can think about ways to automate the process of generating and keeping static versions of critical site content. They can also maintain accounts on these open mirror providers for publishing content or interacting with users, even though no DoS attack is in progress.

BIG BANDWIDTH PROVIDERS

There is another type of big bandwidth provider out there that can be utilized for DoS mitigation. These big bandwidth providers are cloud service providers, such as Amazon and Rackspace. These services are not free and, in the long term, not even necessarily cheap, but they do have gigantic computing and networking resources at their disposal to ensure their clients' content is kept available to their consumers as much as possible. From an architectural and technical point of view, the main consideration of using big bandwidth providers for DoS mitigation is ensuring that the chosen cloud provider/s offer the technical environment necessary to install and run a packaged version of your website.

DISTRIBUTED CONTENT

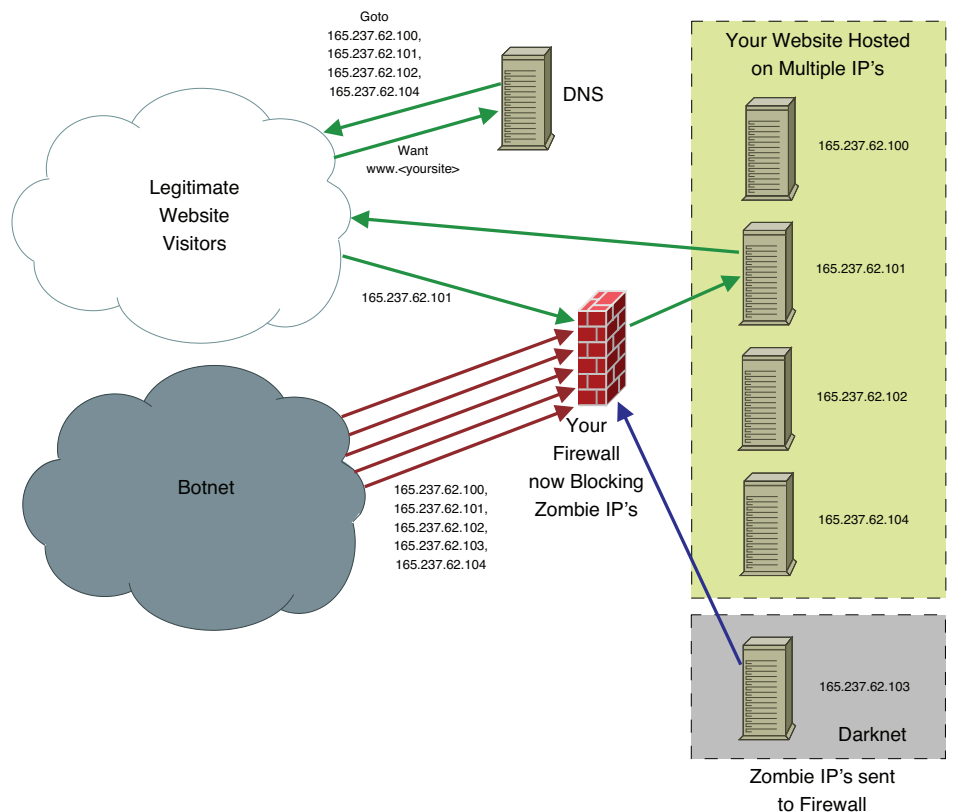
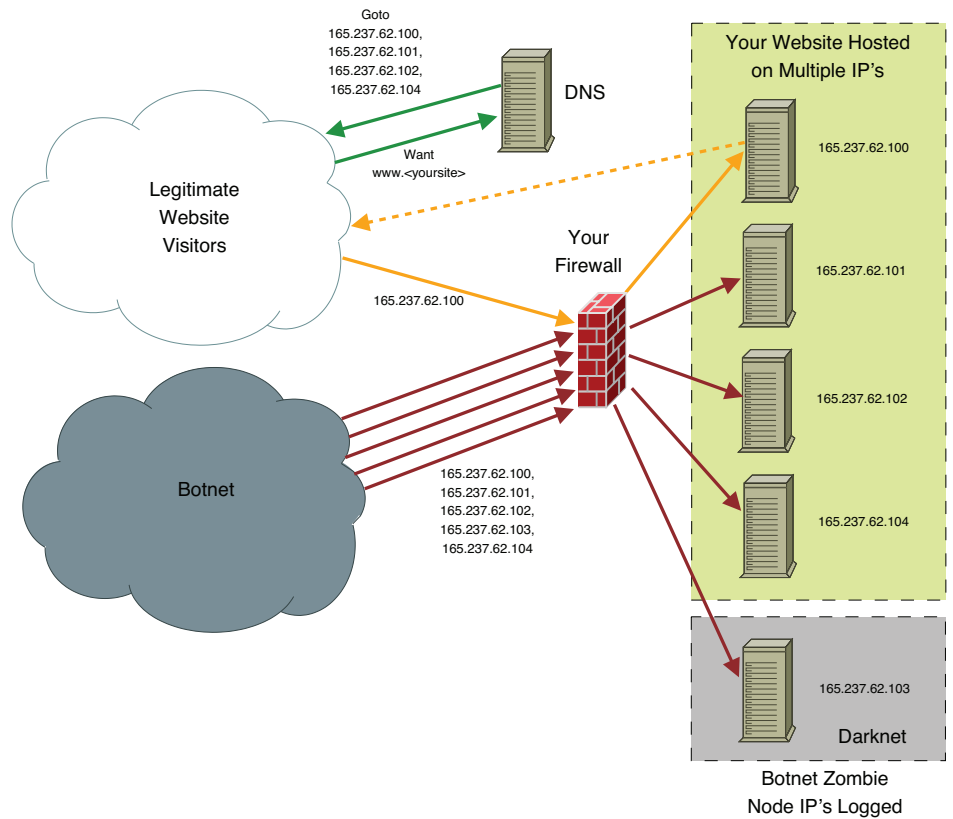
This form of DoS mitigation looks at formatting content in such a way that it is easily consumable via many different protocols and in many different modes of presentation. It is important to adopt the MVC (Model, View, Controller) methodology, which separates the data, the processing of data, and the presentation of data into three very independent functions. In the context of DoS mitigation, this allows the data to always be available for distribution through different media, which helps ensure it is always available. Implementing this mitigation strategy involves designing your website to produce a version of your core content that is distributable via protocols such as RSS. In advance, you will need to organize channels that embed your content into other allied



websites. This gives you the ability to make the content of your site available to your users via those other sites if your own site is unavailable due to DoS attack. The standardization and simplicity of protocols like RSS make it possible for other site owners to quickly set up publishing channels for your content.

DARKNETS

Unlike other strategies, which can be enacted after the DoS has begun, 'darknets' will only work if set up as a preparatory denial of service protection measure. To set up a darknet, you must first purchase a range of IP addresses—for example, 165.237.62.100 - 105. You allocate 100, 101, 102, 104, and 105 to instances of web servers serving your website and put this array of IP addresses into the configuration of your load balancer. Since 103 is not allocated, it is left as a darknet. An attacker will see from DNS queries (etc.) that you have a range of IP addresses allocated to you, and that your website is distributed across the range. When they set up their attack, they may configure it to attack that same range of IP addresses rather than setting the target as the common name or URL for your website (such as 'yoursite.org'). Unknown to the attacker, the range contains a darknet, and when they initiate their attack, you will be able to monitor the machine associated with the 103 address. Any traffic you see hitting that address is malicious traffic. This allows you to take any client IP addresses hitting that server and add them (perhaps even using an automated mechanism) to an upstream firewall. This reduces the number of attacker client IP addresses that get through, leaving more of your web server resources available to serve content to the IP addresses of your real users.



////////////////////
 right: Darknet Strategy,
 Part 1 (top) and Part 2 (bottom)

PART 2: RESPONDING TO A DDoS ATTACK

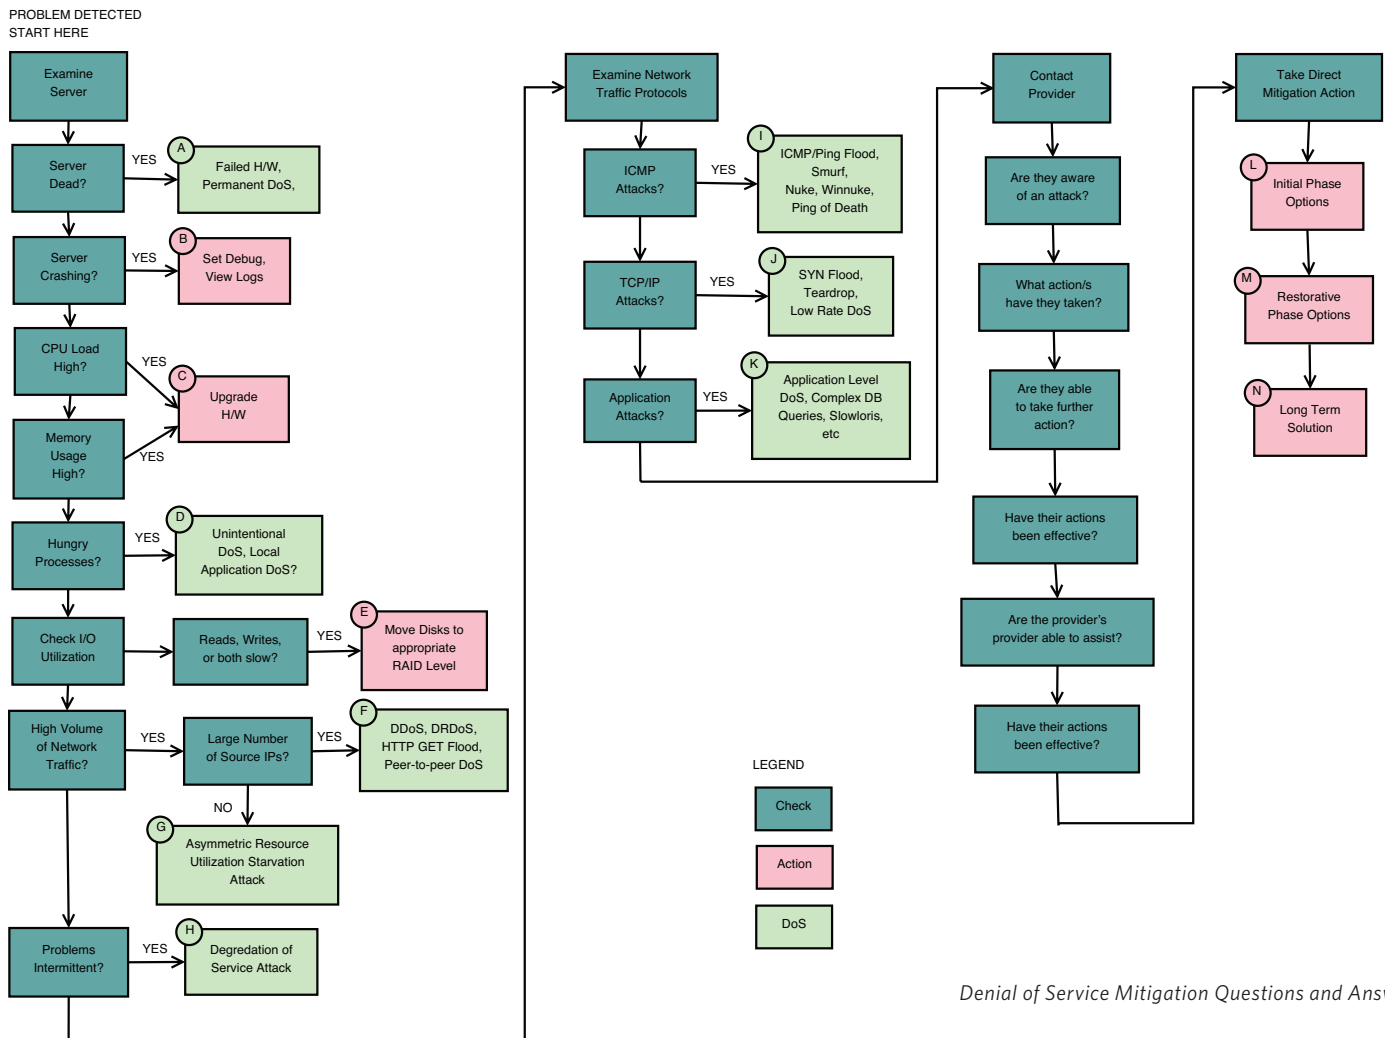
DoS attacks are stressful. They are stressful for the site’s visitors, who are not able to get to the site’s content. They are stressful for the staff of the organization, who knows their message is not reaching their audience. And they are stressful for the technical staff tasked with keeping the website operational. With this in mind, we have devised this simple to follow step-by-step process that will assist in keeping your technical staff efficiently focused on getting the website available to your audience again.

All steps in teal should be worked through, and any answers arising from the questions posed therein should be recorded, as those answers may assist in the diagnosis of the attack.

Steps in pink require action to be taken, whether it be alterations of system configuration, swapping out and replacing hardware, or implementing one or more of the DoS mitigation strategies outlined in this document.

Green steps suggest some possible types of DoS attack that may relate to your situation. These are provided to assist your technical staff in narrowing the focus for further research they may have to do on specific DoS attacks that relate to the situation you are experiencing. Note that while it is unlikely you will actually experience some of the named attacks as operating systems etc are likely to have been patched against the attacks, it is still worthwhile referencing the attacks. The attack vectors utilized by old attacks are likely to come back with new vulnerabilities and researching the old attacks may assist your technicians deal with a new attack utilizing the same attack vector.

Any box with a number reference has a short description of further information following the chart in this document.





RESPONSE PROCESS

A. FAILED HARDWARE, PERMANENT DoS

Check hardware for normal component failures. Cables, hard drives, power supplies, RAM, motherboards, CPU, and network interfaces are all critical components that could cause your web server to go offline. Do be aware that some attacks known as 'Permanent DoS' exist which exploit vulnerabilities in hardware or firmware to 'brick' (irreparably destroy) hardware (usually by writing bogus programs to ROM chips on the hardware device). If you replace a failed hardware component with a new component that subsequently fails similarly to the old one, you may be experiencing a PDoS attack.

B. SET DEBUG, LOGGING FOR CRASH ANALYSIS

If the server is frequently crashing, you may need to set the debug level of some logging processes to a lower level. This may help you analyze the problem by looking at what was logged just before each crash.

C. UPGRADE HARDWARE

If the denial of service you are experiencing manifests itself by the web server running out of compute ability or memory, then upgrading the CPU or RAM may provide immediate temporary relief. Some denial of service attacks are not rate-based, but rather attempt to locally consume all the available resources of components such as CPU and RAM.

D. UNINTENTIONAL DoS, LOCAL APPLICATION DoS

An Unintentional DoS is any denial of service that is accidental. This may be anything from a misconfigured local application, a poorly written custom program, a network technician unplugging the wrong cable at a patch panel, or the 'Slashdot Effect' (named after the site slashdot.org, where articles noting topics of interest to geeks were posted; the resultant immediate interest from Slashdot's readership would usually bring down any website referenced!), so look for evidence of a sudden increase in legitimate visitor traffic due to some unplanned promotion of your website. Also look for evidence of a Local Application DoS, such as a fork bomb. If any such application is found, you will need to subsequently look for evidence of a hacker breach. See also (K).

E. RAID FOR SPEED

A Redundant Array of Independent Disks (RAID) is a technology that breaks data down into segments so that it can be spread across multiple disks, thus allowing parallel writes or reads, which can increase data I/O speed. Refer to Wikipedia for a comparison of the different RAID levels and their theoretical improvements on disk I/O performance (<http://en.wikipedia.org/wiki/RAID>).

F. DDoS, DRDoS, HTTP GET FLOOD, PEER-TO-PEER DoS

Some form rate-based denial of service attack. This could be a Distributed Denial of Service attack using fraggle, a SYN Flood, HTTP GET Flood or similar; a Peer-to-peer DoS through a file sharing network using the ADC protocol or similar; or a Distributed Reflected Denial of Service attack, such as ICMP Echo Request / Smurf attacks or DNS amplification attacks, etc. Depending on the nature of the attack, defending against them can range from simply blocking the ports they use to more difficult measures (if the attack traffic is indistinguishable from legitimate traffic, for example).

G. ASYMMETRIC RESOURCE UTILIZATION STARVATION ATTACKS

Usually such an attack will be perpetrated using a very powerful or capable computer which has strengths that outperform the capability of your web server. The chosen capability could be network bandwidth, compute power, maximum number of concurrent connections, or some other resource.

H. DEGRADATION OF SERVICE ATTACKS

A Degradation of Service attack is any DoS attack performed in intermittent bursts to frustrate any initial attempt to work out what is wrong with the website. Users experience difficulty getting to the website, but by the time they are able to communicate the problem to the website technical team, the attack has ceased, so when the techs look at the server, everything appears normal. This pattern continues until the techs begin to monitor the website at all times.

I. ICMP/PING FLOOD, SMURF, NUKE, WINNUKE, PING OF DEATH

Many ICMP attacks exist, and they form some of the oldest and most frequently used DoS attack methods. If ICMP is not important in your environment, then drop all ICMP traffic at your firewall.

J. SYN FLOOD, TEARDROP, LOW RATE DoS

These are all TCP/IP protocol attacks. Most operating systems now have mechanisms to combat SYN Flood and Teardrop attacks, so ensure those mechanisms are compiled into your kernel. The Low Rate DoS attack exploits the TCP/IP protocol to convince your server to restrict throughput. Occasionally, new attacks that rely on poor implementations of the TCP/IP networking stack are brought to light. Patches will usually be issued by the stack developers reasonably quickly, but you may have to mitigate the attack as best you can in the meantime.

K. APPLICATION LEVEL DoS, SLOWLORIS, INVITE OF DEATH, ETC.

Symptoms of Application Level DoS attacks often include, but are not limited to, system crashes and lockups, CPU and/or Memory at 100% utilization, or abuse of the application protocol. Application attacks may be hard to distinguish from legitimate traffic. For example, they could be complex SQL queries issued to a database backend which can result in high CPU, or database transaction blocks which may prevent other database writes and processes, subsequently resulting in crashing that may corrupt the database. Race conditions and thread starvation are also common mechanisms utilized in these attacks. Other specific examples of

Application Level DoS attacks include HTTP POST DoS, Invite of Death, and Slowloris. Note that for Slowloris there may be a simple mitigation, such as moving your website to a Microsoft platform, as MS IIS is not susceptible to the attack. F5 load-balancing devices also mitigate this attack, so placing one on the network in front of your website would also work. Application attacks perpetrated across the network require the full TCP handshake, which means source IPs cannot be spoofed. You may be able to use this to your advantage by logging those source addresses, using some mechanism to add them to your firewall ruleset, and subsequently dropping all traffic originating from those addresses.

L. INITIAL PHASE OPTIONS

Enact Initial Phase options as laid out in the guide.

M. RESTORATIVE PHASE OPTIONS

Enact Restorative Phase options as laid out in the guide.

N. FINDING LONG TERM SOLUTIONS

Enact Long Term solutions as laid out in the guide. Please refer to the Access website (<https://www.accessnow.org>) for more comprehensive and up-to-date information.



PHASES FOR DENIAL OF SERVICE ATTACK MITIGATION

For a website under a denial of service attack, it is useful to focus the mitigation of the attack into three phases of getting the site back online. These phases are the Initial Phase, the Restorative Phase, and the Long Term Phase. Note that regardless of DoS mitigation phase, one key thing to remember is that it is in the best interests of all your upstream providers to work with you to mitigate the DoS attack. If you are under a rate-based attack, it will be affecting their bandwidth, and may be affecting not just your website but also the sites and networks of other customers as well. Therefore, it is in everyone's best interest to communicate and work together collaboratively. Let your upstream providers know this. Let them know you want to work with them to resolve the issue. Talk to your provider first. Then get them to talk to their provider.

INITIAL PHASE

The aim of the initial phase is to get as much of the content back up and available as possible while continuing to fully diagnose the attack. This may mean compromising about the fullness of coverage available on the site. It may mean the content is not available at the same levels of low latency and high bandwidth, but

it does mean that content is able to be requested and received. It may mean the content is not in its normal full format, or its normal dynamic format. The data may have to be static, simplified or not as aesthetically formatted, but data should be available until a version more resembling the 'normal' site can be restored. Part of the Initial Phase should also be analysis of the DoS attack. This analysis can be done parallel to work being done to get the site content back up, and it will help greatly to determine which

approach to take in the Restorative Phase. To perform the analysis, the web server's network interface will need to be 'sniffed' and packet captures taken. This can be done with tools such as Wireshark (<http://www.wireshark.org/>). Administrative access to the server will be needed as well to see how it is performing. Narrowing down the effect of the DoS attack can help in formulating an effective mitigation response.

INITIAL PHASE STRATEGIES FOR MITIGATING DoS ATTACKS

THIRD PARTY DENIAL OF SERVICE PROTECTION (DOSP)

If DoSP has not been organized beforehand and you find yourself under attack, ask your hosting service provider if they are able to offer any DoSP. If the answer is no, then go to their upstream provider and ask the same question. It may be that a bigger provider upstream has the ability to utilize a DoSP service. The bigger or more 'wholesale' the service provider, the more likely they are to have an agreement to work with a DoSP service. If this turns out to be the case, then you may be able to negotiate with them to clean the traffic destined for your website before it even gets to the network of your service provider.

FIREWALLING

In the mitigation of DoS attacks, we should not discount the usefulness of even a basic firewall. While in some sense it is too late if the volume of a rate-based DoS attack is reaching a device just before your web site, there may be some usefulness here. If the firewall can be configured to drop as much of the malicious traffic as possible, then at least the web server itself may get some of its processing power back, which means that legitimate requests that do get through may actually be able to be served responses. The legitimate incoming and outgoing traffic will still have to fight for bandwidth with attack traffic, but at least you may be able to reduce the number of variables in the equation that are affected by the attack.

The ability to firewall comes down to being able to identify the malicious traffic from the rest of the traffic. There may be certain circumstances in how the attack has been constructed that will help you, or you may be able to utilize tactics such as the use of a darknet (see above) to take an active role in sorting the malicious traffic from the legitimate traffic. If doing a traffic analysis reveals that attack traffic is coming predominantly from a single or small number of countries based on geolocation of the source IP addresses, then all traffic from those countries may be blocked at the firewall if the firewall supports blocking on IP geolocation. This is a fairly harsh measure and will likely also cause some collateral damage, but in the early stages of mitigation, this may be warranted.

If your website is locally hosted within your nation and your audience is almost entirely within the same country, then another drastic measure is to block all traffic coming from any IP geolocated outside the local country. This can be an effective short-term response to a DDoS attack, as the 'zombie' computers that make up the botnet are likely to be widely distributed around the world, not centered in your local nation. By blocking all but local traffic, particularly if it can be done upstream, the majority of your audience who is local will still be able to get to your content, while most of the attack traffic will not.

When thinking of utilizing firewalling as a DoS mitigation tactic, remember that the further upstream the firewalling can be done, the more effective it will likely be in helping your situation. This means talking to your service provider as well as your service provider's upstream provider.

OPEN PROVIDER MIRRORS

If not already in place, having been done in the preparatory phase, then set up an open provider mirror as quickly as you can. This is a very popular option for NGOs and opposition and activist websites that come under DoS attack, as it can be very effective in getting a form of the site content back up and available at an arbitrary speed—the goal of the initial phase. Please refer to the description of the same name within the 'Part 1: Preparatory Phase' section of this document for further information on this type of site mirroring.

DISTRIBUTED CONTENT

If set up in the preparatory phase, then activate this mechanism now and start sending your content to be hosted within portions of allied websites. If no arrangements for the provision of a publishing channel for your content has been made in advance, then a plea for other website owners to embed your content via an RSS feed or similar can be made while the DoS attack is underway.

Darknet

If set up in the preparatory phase, then activate this mechanism now. It may or may not work, depending on whether the attack software used by the attacker does DNS lookups for your web servers or not.

FORGOING THE DOMAIN

This is a very drastic measure, and one likely to have only very short-term benefits in mitigating a DoS attack. This strategy relies on hosting your website either on IP addresses not associated with the website's normal domain name, or hosting it on an alternate domain name. There is likely only a short amount of time before the attacker realizes you have shifted your website and adds the new IP address(es) or domain(s) to their attack configuration. This gives you just a small window where your audience may be able to get to your content. The greatest challenge of this strategy is to find quick and creative ways to let your audience know where to find your content.

