

Air Force Institute of Technology

AFIT Scholar

Theses and Dissertations

Student Graduate Works

3-2005

Analysis of Effects of BGP Black Hole Routing on a Network like the NIPRNET

Michael D. Kleffman

Follow this and additional works at: <https://scholar.afit.edu/etd>



Part of the [Information Security Commons](#)

Recommended Citation

Kleffman, Michael D., "Analysis of Effects of BGP Black Hole Routing on a Network like the NIPRNET" (2005). *Theses and Dissertations*. 3878.

<https://scholar.afit.edu/etd/3878>

This Thesis is brought to you for free and open access by the Student Graduate Works at AFIT Scholar. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of AFIT Scholar. For more information, please contact richard.mansfield@afit.edu.



ANALYSIS OF EFFECTS OF BGP BLACK HOLE ROUTING
ON A NETWORK LIKE THE NIPRNET

THESIS

Michael D. Kleffman, Captain, USAF
AFIT/GIA/ENG/05-01

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the United States Government.

AFIT/GIA/ENG/05-01

ANALYSIS OF EFFECTS OF BGP BLACK HOLE ROUTING
ON A NETWORK LIKE THE NIPRNET

THESIS

Presented to the Faculty

Department of Electrical and Computer Engineering

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the
Degree of Master of Science in Information Assurance

Michael D. Kleffman, BS
Captain, USAF

March 2005

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

Abstract

The Department of Defense (DoD) relies heavily on the Non-secure Internet Protocol Router Network (NIPRNET) to exchange information freely between departments, services, bases, posts, and ships. The NIPRNET is vulnerable to various attacks, to include physical and cyber attacks. One of the most frequently used cyber attacks by criminally motivated hackers is a Distributed Denial of Service (DDoS) attack. DDoS attacks can be used to exhaust network bandwidth and router processing capabilities, and as a leveraging tool for extortion. Border Gateway Protocol (BGP) black hole routing is a responsive defensive network technique for mitigating DDoS attacks. BGP black hole routing directs traffic destined to an Internet address under attack to a null address, essentially stopping the DDoS attack by dropping all traffic to the targeted system.

This research examines the ability of BGP black hole routing to effectively defend a network like the NIPRNET from a DDoS attack, as well as examining two different techniques for triggering BGP black hole routing during a DDoS attack. This thesis presents experiments with three different DDoS attack scenarios to determine the effectiveness of BGP black hole routing. Remote-triggered black hole routing is then compared against customer-triggered black hole routing to examine how well each technique reacts under a DDoS attack. The results from this study show BGP black hole routing to be highly successful. It also shows that remote-triggered black hole routing is much more effective than customer-triggered.

Acknowledgements

First of all, I would like to acknowledge the big guy upstairs for making this all possible. I could not have survived this thesis if it wasn't for the loving guidance of my Father in Heaven. I would also like to thank my wife, without her I would be nothing. She is my rock and I am very lucky to have her in my life. Next, I would like to thank my two girls for being very patient and allowing me the opportunity to fulfill a lifelong dream. I also express my thanks to my advisor, Maj. Robert Graham for his guidance, support, and empathy during this lengthy and taxing process. Next, I would like to thank Dr. Raines for all of his support. Finally, I would like to thank Mr. Timothy Lacey for his A+ support and always being available when I was in desperate need of computer support.

Capt. Michael D. Kleffman

TABLE OF CONTENTS

Abstract.....	iv
Acknowledgements.....	v
List of Figures.....	ix
List of Tables.....	xii
I. Introduction.....	1
Overview.....	1
Research Goal.....	2
Results Overview.....	2
Summary.....	3
II. Literature Review.....	4
Introduction.....	4
NIPRNET Configuration.....	4
Distributed Denial of Service Attacks.....	6
DDoS Attack Strategy.....	6
Classification by Degree of Automation.....	7
Manual Attacks.....	7
Semi-Automatic Attacks.....	7
Automatic Attacks.....	8
Classification by Exploited Vulnerability.....	8
Protocol Attacks.....	8
Brute Force Attacks.....	8
Classification by Attack Rate Dynamics.....	9
Classification by Impact.....	9
DDoS Defense Techniques.....	10
Prevention.....	10
Detection.....	11
Response.....	12
Routing Protocols.....	13
Open Shortest Path First.....	13
Routing Information Protocol.....	14
Internal Border Gateway Protocol.....	15
BGP Black Hole Routing.....	18
Black Hole Routing as a Filter.....	18
Remote-Triggered Black Hole Routing.....	19
Setting up a iBGP Black Hole Routing Enabled Network.....	20

Limitations of Destination Address iBGP Black Hole Routing	22
BGP Black Hole Routing versus Other Techniques	22
BGP Black Hole Routing.....	22
Firewalls.....	23
Intrusion Detection Systems	24
Over-Provisioning.....	25
Review	25
BGP Black Hole Routing Implementations.....	25
Remote-triggered Black Hole Routing	26
Customer-triggered Black Hole Routing	26
Source-based Black Hole Routing	27
Conclusion	27
III. Methodology	28
Problem Definition.....	28
Scope of Problem.....	28
Goals	28
Approach.....	29
Evaluation Technique	30
System Boundaries.....	31
System Services	32
Workload.....	32
Performance Metrics.....	33
Parameters.....	35
Factors.....	37
Experimental Design.....	38
Analysis of Results	42
Summary	44
IV. Results and Analysis.....	45
Introduction.....	45
Effectiveness of BGP black hole routing on a network	45
Effectiveness of BGP black hole routing when not all border routers are dropping attack packets.....	50
Effectiveness of remotely triggering BGP black hole routing on a network like the NIPRNET while it is under attack	61
Effectiveness of customer-triggered BGP black hole routing as compared to remote-triggered black hole routing in defending a network under attack.....	66
Summary	72
V. Conclusions.....	74
Goal Restatement.....	74

Conclusions.....	74
Contributions.....	76
Suggestions for Future Work.....	77
Appendix A. Model Configurations	78
Bibliography	92

List of Figures

Figure	Page
1. Example of Two Bases on the NIPRNET.....	5
2. BGP Weight Attribute.....	16
3. BGP Local Preference Attribute	16
4. BGP Path Selection.....	17
5. Using Static Routes for Black Hole Routing.....	19
6. Remote-triggered Black Hole Routing Scheme.....	20
7. Trigger Router Setup.....	21
8. BGP Activation	21
9. System Diagram	31
10. Queuing Delay of all 4 Simulations.....	47
11. Utilization of 9 Mbps Pipe Defended Against a 12.8 Mbps DDoS Attack	51
12. Utilization of 40 Mbps Pipe Defended Against a 12.8 Mbps DDoS Attack	51
13. Utilization of 9 Mbps Pipe Defended Against a 38.4 Mbps DDoS Attack	52
14. Utilization of 40 Mbps Pipe Defended Against a 38.4 Mbps DDoS Attack	52
15. Utilization of 9 Mbps Pipe Defended Against a 64 Mbps DDoS Attack	53
16. Utilization of 40 Mbps Pipe Defended Against a 64 Mbps DDoS Attack	53
17. Queuing Delay of 9 Mbps Pipe.....	55
18. Queuing Delay of 40 Mbps Pipe.....	56

19. Inbound Latency of 9 Mbps Pipe.....	58
20. Inbound Latency of 40 Mbps Pipe.....	58
21. Remote-triggered Update Router Convergence	60
22. Inbound Bandwidth Utilization of 9 Mbps Pipe	61
23. Inbound Bandwidth Utilization of 40 Mbps Pipe	62
24. Router Queuing Delay of 9 Mbps Pipe.....	62
25. Router Queuing Delay of 40 Mbps Pipe.....	63
26. Inbound Latency of 9 Mbps Pipe.....	63
27. Inbound Latency of 40 Mbps Pipe.....	64
28. Customer-triggered Update Router Convergence.....	65
29. Inbound Bandwidth Utilization of 9 Mbps Pipe	67
30. Inbound Bandwidth Utilization of 40 Mbps Pipe	67
31. Router Queuing Delay of 9 Mbps Pipe.....	68
32. Router Queuing Delay of 40 Mbps Pipe.....	68
33. Inbound Latency of 9 Mbps Pipe.....	69
34. Inbound Latency of 40 Mbps Pipe.....	69
35. Outbound Latency of 9 Mbps Pipe under 64 Mbps DDoS Attack	70
36. Outbound Latency of 40 Mbps Pipe under 64 Mbps DDoS Attack	70
37. Internal Internet Set-up	78
38. Internal Base Set-up.....	78
39. Border Router BGP Parameters	79
40. Border Router BGP Neighbor Information.....	80
41. Border Router IP Routing Parameters	81

42. Border Router IP Processing Information.....	82
43. Border Router IP Quality of Service Configuration of 9 Mbps Pipe.....	83
44. Border Router IP Tunnel Information.....	84
45. Border Router to Internet Link Configuration.....	85
46. Border Router to Base Link Configuration.....	86
47. TCP Settings on All Routers.....	87
48. Border Router Route Map Configuration	88
49. Border Router Static Routing Table	89
50. Trigger-Router Route Map Configuration.....	90
51. Trigger-Router Static Routing Table	91

List of Tables

Table	Page
1. Inbound Utilization Statistics.....	45
2. Outbound Utilization Statistics.....	45
3. Queuing Delay Averages in Microseconds.....	47
4. Inbound Latency Averages in Milliseconds	48
5. Outbound Latency Averages in Milliseconds.....	49
6. Bandwidth Utilization Averages.....	50
7. Queuing Delay Averages in Microseconds.....	54
8. Queuing Delay Averages of Routers in Microseconds.....	56
9. Average Added Queuing Delay of Communication Link in Milliseconds.....	59

ANALYSIS OF EFFECTS OF BGP BLACK HOLE ROUTING
ON A NETWORK LIKE THE NIPRNET

I. Introduction

Overview

As a result of the current Information Age, the Department of Defense (DoD) relies heavily on the Non-secure Internet Protocol Router Network (NIPRNET) to exchange information freely between departments, services, bases, posts, and ships. The primary mission of the NIPRNET is to provide the capability for all departments and services within the DoD to freely exchange information to advance the warfighting capabilities of the DoD.

Although the NIPRNET is composed of hundreds of smaller networks which are geographically separated, the NIPRNET must have the capability to transport information amongst the smaller networks. All types of data are exchanged, from voice communications and data transfer to real-time video and graphics-intensive distributed interactive simulations. The effectiveness and efficiency of this data communications capability will determine in large part the ability of the warfighters to achieve their in-garrison missions, as well as their war-fighting missions.

The NIPRNET is vulnerable to various attacks, to include physical and cyber attacks. One of the most frequently used cyber attacks by criminally motivated hackers is a Distributed Denial of Service (DDoS) attack. DDoS attacks consist of multiple systems

on the Internet sending enough IP packets to a system to effectively stop the system from serving information to legitimate users. DDoS attacks can be used to exhaust network bandwidth and router processing capabilities, and as a leveraging tool for extortion. DDoS attacks are fast becoming a major problem to information assurance. Border Gateway Protocol (BGP) black hole routing is a defensive network technique for mitigating DDoS attacks. BGP black hole routing directs traffic destined to an Internet address under attack to a null address, essentially stopping the DDoS attack from affecting the entire network by dropping all traffic to the targeted system.

Research Goal

Due the vulnerability of the NIPRNET, the DoD is researching various ways to defend it. Consequently, AFIT has been tasked by the National Security Agency (NSA) to assist in this process by using modeling and simulation to evaluate the effectiveness of BGP black hole routing to defend the NIPRNET from DDoS attacks. The goal of this research is to investigate the capabilities of BGP black hole routing in protecting a network like the NIPRNET from a DDoS attack. In addition, this research investigates two approaches to trigger the BGP black hole routing to determine if one is more efficient than the other.

Results Overview

The results of this research proved that BGP black hole routing is effective in defending a network like the NIPRNET from DDoS attacks. This finding alone should assist the DoD community in drafting up defense policies of the NIPRNET against DDoS attacks. Another contribution this research should provide is that it demonstrated

customer-triggered BGP black hole routing isn't as reliable and efficient as remote-triggered BGP black hole routing. Finally, this research demonstrated that BGP black hole routing is more efficient and is a better defense mechanism when all of the border routers are configured to drop attack traffic.

Summary

The remainder of this thesis is organized as follows. Chapter 2 presents a literature review of DDoS attacks and techniques, to include BGP black hole routing, on how to defend networks. Chapter 3 lays out a network performance analysis methodology, which is used to define the problem at hand, describe the approach of this research, define the system, and explain the experimental design. Chapter 4 details the analysis portions of this process by evaluating the simulated network's outputs, which are used in the decision making process to arrive a recommended solution. Chapter 5 concludes this report by summarizing the results and providing recommendations for follow-on work to this research.

II. Literature Review

Introduction

This chapter gives the necessary theoretical background required to devise an effective Distributed Denial of Service (DDoS) attack defense methodology using Border Gateway Protocol (BGP) Black Hole Routing techniques. To do this, the first Section begins with background information on the Non-Secure Internet Protocol Router Network (NIPRNET). The next Section breaks down the different techniques used to conduct DDoS attacks. Then the third Section describes different defense techniques used against DDoS attacks. The characteristics of Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and BGP protocols are discussed in the fourth Section. The technology of BGP Black Hole Routing is covered in Section 5. Section 6 compares BGP Black Hole Routing against other techniques for defending against DDoS attacks. Options for implementing BGP Black Hole Routing are presented in the next to last Section, followed by a summary of the chapter in the final Section.

NIPRNET Configuration

The Defense Information Systems Agency (DISA) is the owner of the NIPRNET. The NIPRNET is a virtual network that is made up of a large compilation of smaller networks throughout the world. Every DoD network comprises the NIPRNET. DISA maintains border routers that connect the NIPRNET to the internet. Each DoD network connects to one of the DISA border routers to enable users to access the internet and each other. The DISA border routers are connected to each other via a Virtual Private Network (VPN). Data exchanged between two DoD networks on the NIPRNET travels

from one border router to another via internet communication channels. A diagram to illustrate how two Air Force bases are connected to the NIPRNET is shown in Figure 1 below.

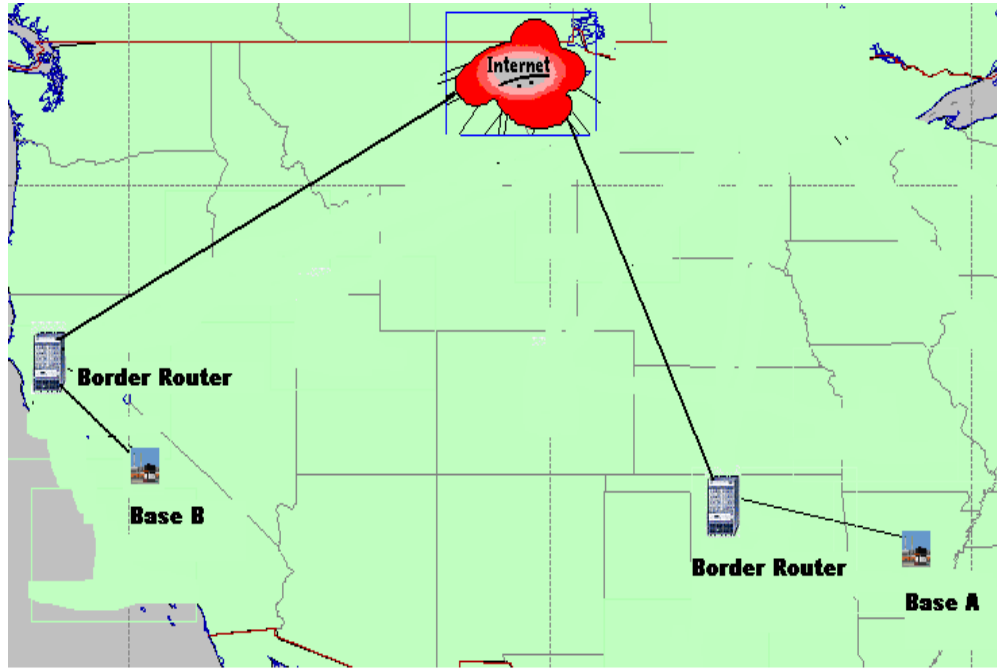


Figure 1 – Example of Two Bases on the NIPRNET

As shown in Figure 1, if Base A wants to send something to Base B the information will travel from Base A to the border router on a dedicated communication line. Once the data is received by the border router, it will transmit the data via a communication line attached to the Internet. The data will flow through the Internet until it reaches the border router that Base B is connected to. The Base B border router will then forward the data to Base B on another dedicated communication line. This illustrates how the NIPRNET is a virtual network and it relies on communication channels on the Internet to transport data from one network to the next.

Distributed Denial of Service Attacks

The goal of a DDoS attack is to inflict damage on the victim, either for personal reasons, for material gain, for popularity, or for political reasons [MMR00]. DDoS attacks have to incorporate hundreds, thousands, or millions of compromised systems to be effective. An attacker will break into these systems to install DDoS software on them [BEN00]. The attacker will then use these compromised systems to attack innocent victims on the internet. These attacks take advantage of limited resources of the victim to exhaust such things as bandwidth, router processing capacity, or network stack resources [BEN00]. The result of such attacks is a victim that can no longer provide a service to its customers. This section will address the following questions:

1. What makes DDoS attacks possible?
2. How do these attacks occur?
3. Why do they occur?

DDoS Attack Strategy.

In order to perform a distributed denial-of-service attack, the attacker needs to recruit the multiple agent machines. This is usually done automatically through scanning of remote machines, looking for security holes. Vulnerable machines are then exploited and become part of the DDoS network. These agent machines are then used to carry out the attack against the victim. Attackers usually hide the identity of the agents during an attack through spoofing of the source address field in packets. This technique allows the attacker to use the agents for a future attack.

Classification by Degree of Automation.

During the attack preparation, the attacker has to locate different agents and infect them with the attack code. To accomplish this, the attacker has three different modes of automation which are: 1) Manual, 2) Semi-automatic, and 3) Automatic [MMR00]. These three modes will be discussed in the following paragraphs.

Manual Attacks.

Only in the early days of DDoS attacks were manual techniques used. These techniques consisted of the attacker breaking into the different machines and loading the attack code. The attacker would then have to manually control each of the agents from his/her own workstation. This technique led the attackers to find faster techniques that led to semi-automatic techniques.

Semi-Automatic Attacks.

In semi-automatic attacks, the DDoS network consists of handler (master) and agent (slave) machines [MMR00]. The attacker deploys automated scripts for scanning and compromise of those machines and installation of the attack code. The attacker then uses the handler systems to specify the attack type and the victim's address and to command the onset of the attack. The attacker can choose to set up the DDoS network with either direct communication or indirect communication. With direct communication, the handlers and slave machines must know each other's identity to communicate. With indirect communication, the attacker uses a service already on the Internet, such as internet relay chat (IRC), and controls the slaves via IRC channels. This technique makes it hard to distinguish between legitimate traffic and DDoS traffic. The use of the service in this case replaces the need for handlers.

Automatic Attacks.

Automatic attacks further automate the attack by programming the time of the attack, attack type, duration, and victim's address into the attack code [MMR00]. This form of attack allows the attacker to have minimal exposure since he/she only has to issue a single command – the start of the attack command.

Classification by Exploited Vulnerability.

Attackers will use different techniques to conduct DDoS attacks to deny clients access to the victim. Two classes are protocol attacks and brute-force attacks.

Protocol Attacks.

Protocol attacks exploit a bug of some protocol installed on the victim in order to carry out the attack [MMR00]. An example of a protocol attack would be a SYN flood. In this type of attack an attacker would repeatedly send SYN packets to the victim without completing the three-way TCP handshake. The victim's queue would eventually fill up waiting for the ACK response from the attacking machines. Thus the victim would not have sufficient processing power to service legitimate traffic requests.

Brute-force Attacks.

Brute-force attacks are carried out by generating a large amount of traffic [MMR00]. The attacker uses the fact that certain services are necessary for a victim and he/she uses this knowledge by crafting what seems to be legitimate traffic. The upstream providers can handle the amount of traffic generated by the attacker, but the victim service has a smaller processing queue, thus a resulting DDoS attack. There are some brute-force attacks that can be filtered at the victim such as ICMP attacks and UDP flood attacks. Other brute-force attacks cannot be filtered and thus the victim is defenseless

against, such as generating a vast amount of HTTP requests to a web server on the victim's network. Brute-force attacks do need to generate a greater amount of traffic to have the same results as protocol attacks.

Classification by Attack Rate Dynamics.

DDoS attacks can be differentiated between continuous rate attacks and variable rate attacks [MMR00]. The majority of known attacks deploy a continuous rate mechanism [MMR00]. This technique is carried out by having the agents generate packets at full force after the attack command is given. The continuous rate attacks are much easier to detect and defend against, since the victim can see the continuous onslaught of traffic coming in. Variable rate attacks can be further broken down into increasing rate attacks and fluctuating rate attacks [MMR00]. Increasing rate attacks start with a low amount of attack packets and gradually increase the number of packets. This technique will lead to a slow degradation of the victim's service, but it does delay the detection of the attack. Fluctuating rate attacks adjust the rate of attack packets depending on the victim's behavior [MMR00]. A form of fluctuating rate attack is pulsing attacks [MMR00]. This is where the attacker will cease sending attack packets, thus the victim will only experience periodic service disruptions.

Classification by Impact.

DDoS attacks can further be classified by the impact they have on the victim. The two different categories are disruptive and degrading [MMR00]. Disruptive attacks strive to shut down the victim's service entirely [MMR00]. Currently, all known attacks belong to this category [MMR00]. Degrading attacks attempt to only consume a portion of the victim's resources [MMR00]. Degrading attacks can have a much greater impact on

victims than disruptive attacks. Degrading attacks could lose the victim customers or even worse have the victim spend more money to upgrade the resources to which the attacker would then just increase the amount of traffic sent to the victim to degrade the victim's service. The victim would never be able to spend enough money to produce a customer friendly environment as long as the attacker went unnoticed.

DDoS Defense Techniques

DDoS Defense Techniques are broken down into three main categories which are: prevention, detection, and response [SEC03]. The goal of a well organized network is to have a good prevention scheme in place to prevent DDoS attacks from occurring in the first place. If a DDoS attack does slip by the preventive measures in place then the network administrators need to be able to detect the DDoS attack. Finally, the part of most well thought out networks that is lacking in today's network oriented world is the response plan for stopping a DDoS attack after it has been successfully launched. The following three sections will discuss some of the techniques used in each of these three areas.

Prevention.

Prevention is the first step to ensure a network will not be affected by a DDoS attack. The first form of prevention would be ingress filtering. Ingress filtering is accomplished by the network routers and would prevent spoofed packets with the same address as the intranet or packets with invalid addresses from entering the network [SEC03]. A clever hacker would use known good addresses to spoof the routers and thus bypass the ingress filtering. Egress filtering, on the other hand, would prevent any

outgoing packets with a source address out of the range of the intranet from leaving the network, thus preventing the systems on the network from participating in a DDoS attack against another network [SEC03]. A good hacker would use valid intranet addresses to send the attack traffic from, thus bypassing the egress filtering. Another form of prevention is to ensure protocol security mechanisms are in place. Examples of protocol security mechanisms include guidelines for a safe protocol design in which resources are committed to the client only after sufficient authentication is done, or the client has paid a sufficient price, deployment of a powerful proxy server that completes TCP connections, TCP SYN cookies, etc. Prevention techniques are a good start to defending against DDoS attacks, but since they can be bypassed, there is still a need for detection.

Detection.

Detection is the phase where either an analyst or a computer system on the network detects something strange about the traffic. One approach is by pattern detection [SEC03]. Using pattern detection, an analyst or computer system can detect most common DDoS attacks, since they usually possess some sort of signature. Another avenue of detection is through the use of anomaly detection [SEC03]. Anomaly detection takes advantage of the fact that most networks have a normal day-to-day network load. Anomaly detection uses this fact and if it notices an unusual load on the network it will raise a red flag. Anomaly detection can be conducted by analysts monitoring their network, or in most cases it is conducted by a system on the network that compares current traffic patterns with a graph of the normal traffic on the network. An advantage of anomaly detection over pattern detection is that anomaly detection has the possibility

of detecting new attacks with no known patterns. A disadvantage of anomaly detection is it could lead to false positives due to an increase in normal network traffic. Now that the DDoS attack has been detected the final thing that needs to be accomplished is responding to it.

Response.

Response is the final action needed to be taken after a DDoS attack has been launched. The main goal of a DDoS attack response is to minimize the impact on the victim while imposing minimal collateral damage to valid customers. There are four general response techniques that can be used in response to an ongoing DDoS attack. The first response technique is called traceback. Traceback determines the addresses (somewhat accurately) of the attackers and informs the victim of their identities [ZAR03]. The victim could then use such things as IP filtering or a personal firewall to prevent packets from the attacking systems from communicating with it. A second response technique involves rate limiting [ZAR03]. Routers on the network could be configured to limit the rate of traffic from addresses sending malicious-appearing packets. The flaw with rate limiting is that it will still allow some attack traffic through and thus a high-scale attack could still prove successful. Filtering is a third response technique. Filtering is conducted with network firewalls or routers by configuring them to deny any traffic on the network from the attacking addresses. A flaw with this technique is the attackers could be using spoofed addresses of valid customers, thus valid customers could be denied access to resources. A final strategy would involve the reconfiguration of the network. The network could be designed to have a system that serves no purpose but to capture attack packets. Once the attack started, the dedicated

victim system could be configured with the actual victim's address and the victim would be assigned a new address. This would involve updating Domain Name Service (DNS) tables and ensuring the update was propagated throughout the internet.

Routing Protocols

The Internet is split into Autonomous Systems (ASes). These ASes correspond to such things as companies, universities, military departments, backbones, etc. Each AS is given a unique 16-bit number (ASN) to differentiate it from other ASes [KAL00]. The administrator from each AS then uses an Internal Gateway Protocol (IGP) to configure each of the routers in his or her domain. Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and internal Border Gateway Protocol (iBGP) are examples of IGPs. External Gateway Protocol (EGP) is used to distribute routing information between ASes. External Border Gateway Protocol (eBGP-4) is the most prevalent EGP used on the Internet today [UCD03]. This section will discuss the three IGPs listed above since BGP Black Hole Routing uses iBGP within an AS. The focus of this section will be to demonstrate the strengths and weaknesses of the three IGPs, as well as the differences between the three. The first protocol to be discussed will be OSPF.

Open Shortest Path First (OSPF).

OSPF functions just as its name suggests. This protocol will use the shortest path between two nodes whenever possible. OSPF is a link-state protocol. The link is an interface on the router while the state is the description of the link as well as the relationship of that link with its neighbors [CIS96]. The link state algorithm used by OSPF begins with either an initialization or a change in the routing information of a

particular router. The router will generate a link-state advertisement to include all of the link-states associated with it. The other routers on the network will then exchange link-states by means of flooding [CIS96]. Each router that receives a link-state update will store it in its database and then forward the update on. Upon update of every router database, the initiating router will calculate the shortest path tree to all destinations using Dijkstra's algorithm. Upon completion, the router will now have a new IP routing table. OSPF only generates network traffic when an update to the network is made, thus reducing the amount of traffic on the network. Some of the advantages of OSPF include: no hop count limitation, allows for better load balancing than RIP, allows for router authentication by using different methods of password authentication, only sends updates when necessary to better utilize available bandwidth, and finally it allows for the transfer and tagging of external routes loaded into the AS [CIS96].

Routing Information Protocol (RIP).

RIP uses algorithms that use distance vectors to mathematically compare routes to determine the best path to any given address. RIP uses a single routing metric (hop count) to determine the distance between the source and destination networks. RIP routers will only store the best route (the route with the fewest hops) to a destination. RIP sends routing-update messages at regular intervals and when the network topology changes [DAT04]. A router will update its routing table when it receives a routing update with changes. The metric value for the path is increased by one, and the sender is indicated as the next hop [DAT04]. After updating its routing table, the router immediately begins transmitting routing updates to inform other network routers of the

change. These updates are sent as soon as changes are made with no regard to when the next scheduled update is due to take place. RIP also prevents routing loops by limiting the number of hop counts in any source to destination path to 15 [CIS96]. RIP routers will classify a network as unreachable if a routing update causes the metric associated with that network to reach 16. A negative to this approach is that RIP networks are limited to 15 hops between any given source and destination. RIP uses timers to regulate its performance. The routing-update timer keeps track of the time between periodic routing updates, which is usually a 30 second window. A route-timeout timer is used to mark routes as invalid when the timer expires without receiving any updates from that route within the timeout window. The route will be retained in the router's routing table until the route-flusher timer expires at which time the invalid routes will be removed. Some of the advantages of RIP include: limiting of hops virtually eliminates loops, timers ensure only valid routes are maintained in the routing tables, as well as mechanisms that prevent incorrect routing information from propagating throughout the network.

Internal Border Gateway Protocol (iBGP).

BGP-enabled routers within the same AS use iBGP when communicating with one another. The core router is the only router that uses eBGP to publish updates to its neighbors of other ASes. eBGP is the most prevalent used EGP on the Internet. For the remainder of this section, BGP represents both iBGP and eBGP. BGP is a very robust and scalable routing protocol. BGP uses many routing parameters, or attributes, to maintain a stable routing environment. The weight attribute is local to the router and is

never advertised to the community [CIS03]. When a BGP router receives multiple paths to the same network, it assigns weight values to each of the paths, and the one with the highest weight is loaded into its routing table. Figure 2 demonstrates this process.




Figure 2 - BGP Weight Attribute [CIS03]

The local preference attribute is used to determine the exit point from one AS to another AS. When multiple routers on the local AS receive advertisements to another AS, they each assign a local preference value and the one with the highest value is then designated as the exit point to send traffic from the local AS to the advertised AS [CIS03]. Figure 3 illustrates this.




Figure 3 - BGP Local Preference Attribute [CIS03]

The origin attribute indicates how BGP learned of a particular route. The three possible values of the origin attribute are IGP, EGP, and Incomplete [CIS03]. An origin attribute

of IGP indicates the route was learned via the internal AS, whereas EGP indicates the route was learned from an external AS. Incomplete states the origin of the route is unknown or was learned via other means. The community attribute provides a way of grouping destinations, to which routing decisions can be applied. A community attribute of no-export states not to advertise the route to eBGP peers. This attribute is used with Black Hole Routing. A community attribute of no-advertise states not to advertise the route to any peers and an internet community attribute states to advertise the route to the internet community. BGP uses these attributes to determine which paths to put in its routing table. The first decision BGP makes with an update is if it specifies a path that is inaccessible, then it will be dropped. If the update is good then BGP loads the path with the greatest weight. In the case of identical weights, BGP will then prefer the path with the largest local preference. When the local preferences are the same, BGP will prefer the path generated by BGP running on the local router. If no route was originated, then BGP will prefer the path with the shortest AS_path. In the case of all AS_path lengths being equal, BGP will prefer the path with the lowest origin type beginning with IGP as the lowest. If all attributes are equal, BGP will prefer the path with the lowest IP address.

Figure 4 summarizes the decisions made by BGP for path selections.

- If the path specifies a next hop that is inaccessible, drop the update.
- Prefer the path with the largest weight.
- If the weights are the same, prefer the path with the largest local preference.
- If the local preferences are the same, prefer the path that was originated by BGP running on this router.
- If no route was originated, prefer the route that has the shortest AS_path.
- If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
- If the origin codes are the same, prefer the path with the lowest MED attribute.
- If the paths have the same MED, prefer the external path over the internal path.
- If the paths are still the same, prefer the path through the closest IGP neighbor.
- Prefer the path with the lowest IP address, as specified by the BGP router ID.

Figure 4 BGP Path Selection [CIS03]

BGP Black Hole Routing

BGP black hole routing is a responsive filtering technique used to defend against active DDoS attacks. This technique uses the router's ability to route unwanted traffic to a Null0 interface (black hole) [CIS04]. Black hole routing takes advantage of iBGP to stop unwanted traffic from reaching the victim.

Black Hole Routing as a Filter.

It is common practice to use the Null0 interface to filter packets to a predetermined destination [GRE02]. The creation of static host routes pointing to the Null0 pseudo interface is how black hole routing is accomplished [GRE02]. The Null0 is a pseudo-interface that is always up and can never forward or receive traffic, much like the host address of 127.0.0.1 on personal computers. The Null0 interface is not a valid interface within the Forwarding Information Base (FIB) [GRE02]. Since Null0 is not a valid interface, packets forwarded to Null0 will be dropped by the FIB. An example of how to set up the interface to black hole traffic destined to a certain address is shown below.

```
interface Null0
    no icmp unreachable

ip route 171.68.10.1 255.255.255.255 null 0 [GRE02]
```

The “no icmp unreachable” command is used to prevent the router from becoming overloaded with numerous ICMP unreachable replies. Black hole routing relies on the strength of the router's forwarding performance to drop the black listed packets. Black hole filtering can be configured to either drop packets based on the destination address or

the source address. Figure 5 below shows how the Black hole filtering technique is executed.




Figure 5 – Using Static Routes for Black Hole Routing [GRE02]

Remote-Triggered Black Hole Routing.

Remote-triggered black hole routing relies on the strength of the routers to route unwanted packets to the Null0 interface at the border routers. The routing update that is sent by a trigger router via iBGP activates a pre-configured static route on all of the border routers that filters traffic to a particular address. Figure 6 below demonstrates how this remote-triggered black hole routing is accomplished.




Figure 6 – Remote-Triggered Black Hole Routing Scheme [GRE02]

Setting up a iBGP Black Hole Routing Enabled Network by Utilizing the Destination Address.

There are four steps that need to be considered when setting up a network to use iBGP Black Hole Routing as a defensive mechanism against DDoS attacks. The first step is to set up a static route to the Null0 interface on all of the routers you want to participate. Included in this step is the allocation of a block of address space that is not used on the internet, for example the 192.0.2.0/24 network space [MOR04]. The second step is to set up the trigger router. The trigger router should be included in the iBGP network with the other routers [MOR04]. The trigger router doesn't have to be a router at all, it could be a workstation with tools such as Zebra/GateD installed on it [MOR04]. Figure 7 below demonstrates how the trigger router is set up to redistribute static routes.

```

router bgp 65501
.
redistribute static route-map static-to-bgp
.
!
route-map static-to-bgp permit 10
match tag 66
set ip next-hop 192.0.2.1
set local-preference 50
set community no-export
set origin igp
!
Route-map static-to-bgp permit 20

```

Redistribute Static with a route-map

Match Route Tag

Set BGP NEXT_HOP to the Trigger

Set LOCAL_PREF

Figure 7 – Trigger Router Set-up [MOR04]

This example matches a static route of 66 and sets the next hop to 192.0.2.1. It sets the local-preference to 50 to override the original advertised route and finally, it sets the community no-export to prevent the route from being advertised outside of the network that this router is on. The third step is activation. In the activation phase a static route for the address under attack will be put into the trigger router with a tag of 66. The trigger router will then advertise this to the other routers in the network, who will see that it has a local-preference of 50 and will put it in the FIB. From this point, all traffic destined to the address under attack will be sent to the Null0 interface. Figure 8 below demonstrates this process.

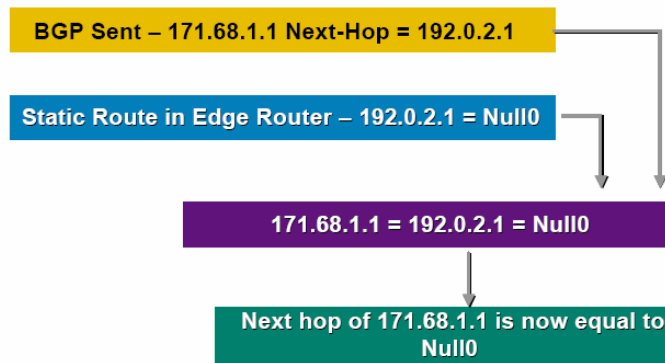


Figure 8 – BGP Activation [GRE02]

The final step involves the removal of the static route from all of the routers. The trigger router will send out an iBGP withdrawal to all of the routers, that will remove the static route from the FIB and thus traffic will then start flowing normally back to the address that was under attack [GRE02].

Limitations of Destination Address iBGP Black Hole Routing.

Black hole routing takes place at the network layer of the TCP/IP model [CIS04]. Thus, it is not possible to use this technique to filter TCP protocols such as HTTP, Telnet, etc. To accomplish this form of filtering, you would have to use a firewall or equivalent. Black hole routing is an all-or-nothing technique for stopping unwanted traffic. Another limitation of black hole routing is it is hard to by-pass or provide exceptions to the filtering. If you wanted legitimate traffic to get to the address under attack you would have to find some way to by-pass the FIB, which isn't a trivial process, or to conduct source-based filtering, which isn't a great option. Finally, routers used to conduct black hole routing and their associated links have to be robust enough to stand up against a rigorous DDoS attack [RAJ02].

BGP Black Hole Routing versus Other Techniques

In reality, there is no known defense technique that will stop all DDoS attacks. This section demonstrates how BGP black hole routing, as a responsive technique, differs from other preventive, responsive, and detection techniques.

BGP Black Hole Routing.

BGP black hole routing takes advantage of stopping the attack traffic at the perimeter of the network. It uses all of the perimeter routers as a defense mechanism. By

stopping the DDoS traffic at the exterior, the BGP black hole routing network is protected from loss of bandwidth within the internal network. By using other tools, such as Intrusion Detection Systems (IDS), BGP black hole routing can be automated to activate the trigger router to send out the iBGP announcements of the static routes to use. This alleviates human intervention during an actual DDoS attack. BGP black hole routing isn't without weaknesses, though. By using destination-based BGP black hole routing the network is dropping all packets destined for the victim machine that might include some legitimate traffic. This is one aspect of BGP black hole routing that network administrators will have to consider before deciding on using BGP black hole routing. Another advantage of black hole routing is it takes advantage of equipment that is already on the network. By using the routers that are already in place it saves money and time by not requiring additional equipment, money, and set-up time to defend against DDoS attacks. BGP black hole routing isn't the Holy Grail when it comes to DDoS defense and that is why this section will address some other defense techniques.

Firewalls.

Another responsive DDoS defense technique commonly used is firewalls. Firewalls are designed to deny any unwanted traffic from entering the network. Firewalls are in-line systems, which mean they are potentially a single point of failure if they quit working. This leads to the first disadvantage of firewalls, in that they are susceptible to DDoS attacks themselves [RIV04]. If a firewall isn't robust enough to handle a large DDoS attack then the firewall system could fail, thus accomplishing the goal of a DDoS attack by taking the entire network off-line. A second disadvantage of firewalls is they can not detect spoofed IP addresses [RIV04]. Once a DDoS attack starts using spoofed

IP addresses, firewalls are useless against defending the network. An advantage of using firewalls is that they operate at layer 4 of the OSI model. Firewalls can effectively block traffic destined for specific TCP protocols, such as HTTP, DNS, and SMTP. On the other hand, if an attack was focused at a public web server that the firewall had to allow traffic to, then the firewall would be rendered useless against such an attack due to firewalls' inability to detect anomalous traffic [RIV04]. Firewalls can shut down flows associated with DDoS attacks, which is another advantage of using firewalls as a DDoS defense mechanism. This flow shut down can only occur when the IP addresses remain constant.

Intrusion Detection Systems (IDS).

IDS solutions are a detection technique used to defend against DDoS attacks. An IDS solution could provide behavioral or anomaly-based algorithms to detect a DDoS attack. The only limitation is that IDS solutions only detect DDoS attacks and they have no capability to mitigate such an attack [RIV04]. This is the reason that the use of an IDS solution would have to be combined with another DDoS defense mechanism capable of mitigating an attack. IDS solutions have emerged through the years as a prime candidate to detect DDoS attacks. A weakness of IDS solutions is they are only as useful as the people who have designed the algorithms. IDS solutions have to constantly be tweaked by experts to stay current with the latest DDoS attack signatures [RIV04]. IDS solutions also rely on storing information in a database for analysts to look at network traffic patterns. A DDoS attack could take advantage of this and send enough traffic to actually crash the IDS by filling up the database and rendering the IDS useless for detecting future attacks.

Over-provisioning.

A prevention strategy in defending against DDoS attacks is to buy excess bandwidth or network devices to withstand such attacks [RIV04]. This strategy has a weakness in that, no matter how much bandwidth is available or how much redundant equipment is available, an attacker just needs to increase the amount of attack traffic to achieve a successful attack. This strategy is not cost-effective either. Most companies, to include the Department of Defense (DoD), don't have the money to invest in excessive bandwidth and network devices just to defend against DDoS attacks.

Review.

As you can see from this section, there is no silver bullet for defending against DDoS attacks. BGP black hole routing definitely seems to be the most capable, but it is not without its weaknesses. To make better use of black hole routing, it stands to reason that combining it with an IDS solution would increase a network's defense posture against DDoS attacks. There are other defense techniques, but this section highlighted the most common techniques used in networks today.

BGP Black Hole Routing Implementations

This section will address black hole routing implementations. Remote- and Customer-triggered black hole routing refers to ways to implement the dropping of packets. Source-based black hole routing refers to what packets should be dropped. Source-based black hole routing is addressed here due to the fact the remainder of this research will be using destination-based black hole routing. Destination-based black hole routing is a safer choice to use, since it is dropping all traffic to a known internal address

and the only affect is that one internal server is not able to serve its customers. Source-based black hole routing is dropping traffic based on an outside source address range. Therefore, there is a potential that customers are not able to conduct business with more than just one internal server if their addresses are within the source address range being dropped by the border routers.

Remote-Triggered Black Hole Routing.

Remote-triggered black hole routing is the dominant implementation in use today. Remote-triggered black hole routing is controlled by the Internet Service Provider (ISP) and in the case of the NIPRNET the ISP would be DISA. Remote-triggered black hole routing was thoroughly explained previously in this chapter.

Customer-Triggered Black Hole Routing.

Customer-triggered black hole routing is controlled by the customer of the ISP. In the case of the NIPRNET, the customer would be the base under attack. Customer-triggered black hole routing takes advantage of the fact that the customer's router speaks BGP to the border routers. If the customer notices his or her network is getting overloaded with a DDoS attack, he or she can send a BGP update to the border routers through his or her local router to have the incoming traffic dropped. As with remote-triggered black hole routing, the border routers would have to be configured to accept and apply the updates from the customer's router. The commercial ISPs do not like to allow customers to update their border routers due to the chance of sending bogus updates that would affect more than just the customer under attack. That is the reason that customer-triggered black hole routing is not widely used in the internet.

Source-Based Black Hole Routing.

This research is conducted to test the effects of black hole routing using the destination address of the attack traffic. Another form of black hole routing is to drop traffic based on the source address. Source based black hole routing incorporates Cisco's unicast reverse path forwarding (uRPF) check. uRPF checks the source address of incoming packets against the routing table to determine whether the interface the packet arrived on is the next hop interface for the source address. Packets are dropped if they arrive on an interface the router wouldn't use to reach the source address. This technique works well for interfaces connected to customer networks and it can work for peers, but it isn't very useful on transit links.

Conclusion

This chapter has presented the theoretical background necessary to devise an effective Distributed Denial of Service (DDoS) attack defense methodology using Border Gateway Protocol (BGP) Black Hole Routing techniques. It began with an overview of the NIPRNET, followed by a description of the different DDoS attacks. Next, a brief review of the different defense techniques used against DDoS attacks was presented. The characteristics of the three major routing protocols were addressed, followed by an in depth review of how BGP black hole routing is currently used to defend against DDoS attacks. BGP black hole routing was compared against other techniques for defense against DDoS attacks. Finally, an explanation of BGP Black Hole Routing options to defend against DDoS attacks was addressed.

III. Methodology

Problem Definition

Scope of Problem.

The Department of Defense (DoD) must have the capability to defend against DDoS attacks to protect information systems. BGP black hole routing has been used successfully in the civilian sector, but certain aspects of it have not been studied to determine the feasibility of using it as a defensive mechanism against DDoS attacks. This research determines the feasibility of using BGP black hole routing to defend against DDoS attacks targeting NIPRNET systems.

Goals.

This research had five goals which were:

1. Determine if BGP black hole routing has any adverse effects on the normal operations of a network like the NIPRNET.
2. Determine if BGP black hole routing is effective in defending a network like the NIPRNET against a DDoS attack.
3. Determine the effectiveness of BGP black hole routing when one or more border routers are not dropping attack traffic.
4. Determine the feasibility of remote-triggered BGP black hole routing on a network like the NIPRNET during a DDoS attack.
5. Determine whether customer-triggered BGP black hole routing is as effective as remote-triggered BGP black hole routing.

A hypothesis of goal one was that BGP black hole routing would not affect Virtual Private Network (VPN) traffic between nodes within a NIPRNET, but it would prevent the delivery of any Internet traffic to the system under the DDoS attack. As for the second goal, BGP black hole routing has been effective in the civilian sector, so it was expected to be successful in a NIPRNET as well. In regards to the third goal, it stood to reason that if one or more of the border routers continued passing DDoS traffic to the target system, then a sophisticated DDoS attack could still be successful by sending a majority of the attack packets through those routers. A hypothesis of the fourth goal was that it should take longer to successfully trigger black hole routing when the DDoS attack uses more attack packets. Finally, customer-triggered BGP black hole routing shouldn't be as effective as remote-triggered black hole routing due to the fact that the base routers communicate with the border routers on communication channels with smaller bandwidth capabilities.

Approach.

To achieve the first goal, the approach was to compare bandwidth utilization, latency, and border router queuing delay in the NIPRNET during an attack to the same data prior to the attack. If the data between the two was the same, then it was determined that BGP black hole routing does not have any adverse affects on the normal operations of the NIPRNET. The approach to achieve the second goal used the same data as the first goal. If the attack traffic was successfully dropped and there were no adverse affects on normal operation of the NIPRNET, then it was determined that BGP black hole routing successfully defended against DDoS attacks. The third goal was achieved by comparing

the bandwidth utilization, inbound latency of the bases under attack, and the queuing delay of the border routers connected to the bases under attack. For this goal, there were three different scenarios of data to compare against the baseline data. If the data was the same then it was determined that BGP black hole routing was effective in defending a network with less than all of the border routers actively participating in black hole routing. The approach to achieve the fourth goal was to measure the router convergence delay. This test measured the amount of time it took each of the border routers to receive the update packet from the trigger router and to start blocking the DDoS attack traffic. To achieve the fifth goal, the approach was to compare the router convergence delay data between remote-triggered simulations and customer-triggered simulations. If the border routers received the updates and started blocking the DDoS attack traffic in the same amount of time, it was determined that customer-triggered black hole routing is as effective as remote-triggered black hole routing.

Evaluation Technique

The evaluation technique used in this research was simulation. The simulation software package used was Opnet Modeler version 10.5. Empirical data was obtained so simulation can test the hypotheses. The model used in this simulation can also be used in future studies of BGP black hole routing. Due to the complexity of setting up a large network and the cost that would have been incurred, direct measurement wasn't considered. Analytical modeling wasn't chosen due to its inaccuracy and the need for numerous assumptions. Simulation is commonly used for network studies and was the most appropriate technique for this research. The results obtained from this research were validated by both empirical data and analytical modeling. The baseline was tested

against the empirical data obtained from the AFNOC to ensure that it is as close to accurate as possible. Analytical modeling techniques were used to obtain the theoretical results of each of the tests to ensure the results from the simulation were similar.

System Boundaries

This section defines what comprises the system under test. This section also defines the component under test as well as any research scope limitations. The System Under Test (SUT) is the NIPRNET beginning with the Defense Information Systems Agency (DISA) border routers. The system diagram is displayed in Figure 9 below. The system consists of the six border routers, the trigger router, the twelve Air Force base routers, a web server at each of the twelve bases, and traffic generators at each of the twelve bases. Only the BGP protocol is used to transmit packets from one router to another. The border routers use BGP to send packets from the Internet to the base routers. The base routers use BGP to send packets to the web servers.

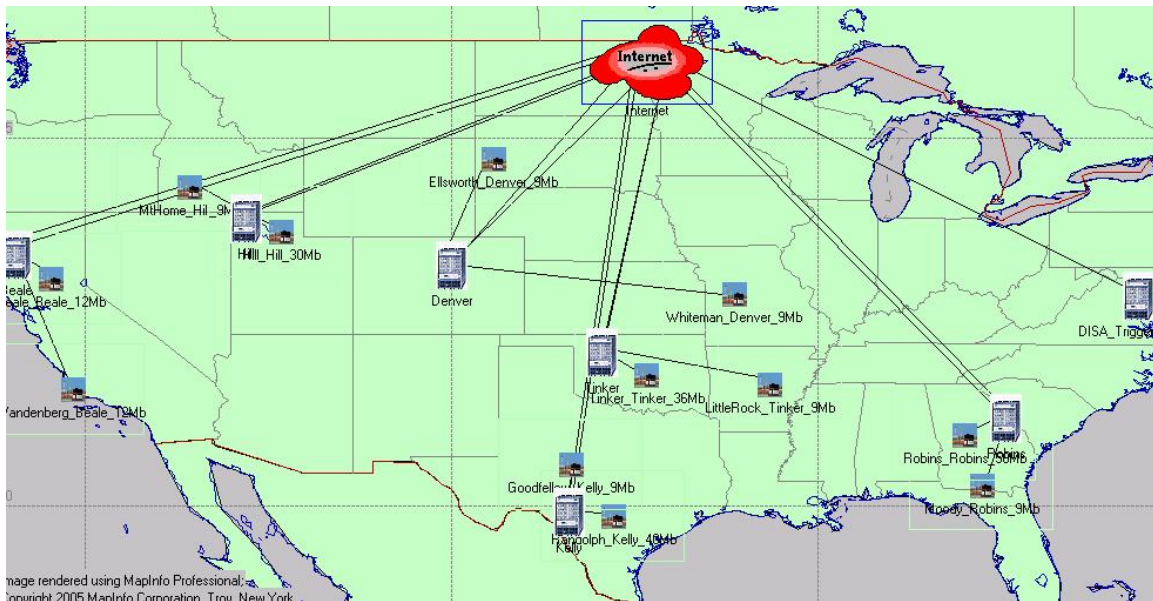


Figure 9 – System Diagram

A non-routing router is configured as the “trigger” router to remotely trigger the deployment of BGP black hole routing to the other six border routers. The trigger router was configured to send the BGP black hole routing update after allowing the DDoS attack to last for ten seconds. The ten seconds was chosen to simulate an IDS identifying a DDoS attack and sending the update to the trigger router within ten seconds. The internet traffic, which includes DDoS traffic, is not a part of this system since it is introduced as part of the workload. The key component under study for this research is BGP black hole routing.

System Services

The service this system provides is defense against a DDoS attack. The border routers are the primary defensive system protecting the NIPRNET from the DDoS attack. The expected outcomes under increasingly severe attacks are that the routers will successfully drop all attack packets and forward all legitimate traffic, the routers will be partially degraded and drop some legitimate traffic while forwarding some of the attack packets, and finally the routers will cease forwarding any traffic, failing to prevent a DDoS attack.

Workload

The workload used to determine a baseline performance of the NIPRNET is based on data obtained from the Air Force Network Operations Center (AFNOC). This data consists of the line speed of each base backbone and the percentage of bandwidth each base utilizes on average. Using this data to establish a baseline allowed comparisons between normal NIPRNET operation and NIPRNET operations with additional workloads introduced.

The second workload introduced to the system was the baseline workload with an additional 12.8 Mbps of attack traffic from six attack systems to simulate a small DDoS attack.

The next workload introduced consisted of a baseline workload with an additional 38.4 Mbps of attack traffic from six attack systems to simulate a medium DDoS attack.

The final workload consisted of the baseline workload and an additional 64 Mbps of attack traffic from six attack systems to simulate a large DDoS attack.

Performance Metrics

The performance metrics chosen for this research were queuing delay, latency, router convergence delay, bandwidth utilization, and packet drop rate (which was later dropped).

The first performance metric is queuing delay. Queuing delay is the time the router holds a packet for transmission before forwarding the packet through the network. The queuing delay is affected by the amount of traffic the router has to process. The queuing delay measurement used for this research was the average delay of each packet over a two-minute time interval. The router's queuing delay measured the effect the DDoS attack had on the router's capability to process packets.

The second performance metric is latency. Latency is the amount of time it takes packets to travel through the network. Latency is measured on the communication channels between the bases and border routers and between the "trigger" router and the border routers. The latency measurement used for this research was the average delay of each packet over a two-minute time interval. The latency measurement was used to

determine how much slower traffic traversed the network due to the introduction of DDoS traffic.

The third performance metric is router convergence delay. Router convergence delay for this research was the amount of time it took the six border routers to update their routing tables with the black hole routing update. Router convergence delay was measured on the pipes between the “trigger” router and the border routers. The router convergence delay measurement used for this research was the average amount of time each of the border routers took to update their routing tables. The router convergence delay measurement was used to determine how long it took to update the routing tables during a DDoS attack.

Bandwidth utilization is a percentage of the network capacity being used. It is measured on the pipes between the bases and the border routers. This measurement revealed how much additional bandwidth the DDoS traffic is consuming. It was also a good measurement to determine how much bandwidth the system can offer and still provide a good quality service.

The packet drop rate is a ratio of the number of packets dropped within the network. It was measured on the border routers. During analysis of the data obtained for the first goal, it was determined that the packet drop rate did not add value to this research. Due to the fact the border routers were suppose to be dropping attack packets, it was difficult to determine whether the packets dropped were actual attack packets or valid packets. Therefore, the packet drop rate performance metric was not used for this research.

Parameters

The parameters for this system include network bandwidth, router configuration, number of routers actively participating in BGP black hole routing during a DDoS attack, the number of routers handling DDoS traffic, length of simulation run, time at which border routers start dropping attack traffic, and length of DDoS attack.

The bandwidth for each individual base pipe was assigned according to the data obtained from the AFNOC and varied between 9 and 50 Mbps per pipe. The sizes of the two pipes that are attacked are 9 Mbps and 40 Mbps. The pipes between border routers and on the internet were all assigned a bandwidth of 155 Mbps. The reasoning behind using bandwidth as a system parameter was simple: sufficient amounts of bandwidth leads to an efficient system. In addition, since this research was focused on the effects of the NIPRNET, the external pipes were assigned excessive bandwidth capabilities to prevent them from becoming the bottleneck in the system. If the system didn't have a sufficient amount of bandwidth, the baseline configuration would not perform adequately, thus the system would not be efficient to begin with.

The router configuration has a major effect on the system performance as well. If the router isn't capable enough, it could become a bottleneck in the system. The routers were configured to process 100,000 packets per second, which is adequate for a border router. The queues for the routers had to be configured to only hold one second of data for each interface connected to a base. This was due to the limitations of the system on which the simulations were run.

The number of routers participating in BGP black hole routing will have an impact on how the system behaves. With just one border router not defending against a

DDoS attack, the system should show an increase in bandwidth utilization. The system may behave differently if a majority of the DDoS attack traffic is traversing through a portion of the routers as opposed to traversing uniformly through each of the routers. This research assumes that the DDoS attack traffic is distributed evenly amongst all of the border routers.

Each simulation run lasted three minutes. The border routers converged with each other within the first minute. The attack traffic was introduced at the one minute mark and lasted for the final two minutes of the simulation.

The system consists of an initial set of IP traffic to simulate bandwidth utilization and an additional set of IP traffic. The workload parameters of this IP traffic that could affect the performance are the packet arrival rate, packet distribution, and packet size. The packet arrival rate was chosen because it affects queuing delays on the routers. The packet arrival rate used for the baseline system was different for each base, due to each base having different bandwidth utilization quotas. The empirical data obtained from the AFNOC determined the packet arrival rate for each base to ensure each base's data was verified against the empirical data. Packet distribution affects bandwidth utilization and latency throughout the system. The packet distribution used for this research was a Poisson distribution. Packet size was chosen because it affects queuing delays at the routers, as well as latency. Due to the nature of DDoS attacks and how the packets can be different between different styles of attacks, the packet size parameters chosen for this research were exponential and 1024 bits. The additional set of IP traffic is used to simulate a DDoS attack. It consists of the same set of workload parameters as the initial

set of IP traffic and is chosen for the same reasons. This set of IP traffic is essential in determining how well the system responds under a DDoS attack.

Factors

This research uses IP traffic offered load, the number of routers participating in BGP black hole routing, and the time set for the border routers to start dropping attack traffic as factors.

The level of the IP traffic factors are six attack systems generating 12.8 Mbps of traffic by sending 2083 packets per second, six attack systems generating 38.4 Mbps of traffic by sending 6250 packets per second, and six attack systems generating 64 Mbps of traffic by sending 10417 packets per second. The three attack scenarios were chosen to simulate 50, 150, and 250 attackers respectively with each attacker possessing a 256 Kbps upload capability. The reason for the six actual attack systems was to evenly distribute the attack traffic among the six border routers. People who launch DDoS attacks do not know how much bandwidth the target system possesses. They simply continue to increase the attack traffic until the service is unavailable. By altering the additional traffic, this research demonstrates appropriately how effective BGP black hole routing is under different attack scenarios.

Based on the pilot study discussed on the next page, the number of routers participating in BGP black hole routing is one, three, five and all. The number of routers actively participating in BGP black hole routing is used to determine how effective BGP black hole routing defends against DDoS attacks when not all border routers are available for defense. The number of routers was altered to simulate certain routers not receiving the routing table update through the network, thus not permitting them to defend. The

system bandwidth utilization should increase proportionally to the number of routers not actively defending against the DDoS attack. As the number of routers not defending increases, the queuing delay at the border routers not defending should also proportionally increase. A pilot study was conducted to determine how many different values of routers should be used in this research. The performance metrics obtained from the pilot study showed insignificant differences between the case of zero or one router, the case of two or three, and the case of four or five routers defending the network against a DDoS attack. The results of this pilot study were used to set the numbers of routers to one, three, and five.

To achieve goals one and two of this research, the border routers were configured to start dropping attack traffic at the one minute mark of the simulation. This was accomplished to ensure all of the attack traffic was being dropped from the beginning of the attack. To achieve goals three and four, the trigger router was configured to send an update to the border routers at the 70 second mark of the simulation to simulate taking ten seconds to detect a DDoS attack. Finally, to achieve goal five, the base router under attack was configured to send an update to the border router at the 70 second mark of the simulation to also simulate taking ten seconds to detect a DDoS attack.

Experimental Design

This research had a limited number of factors, but a full factorial design was not utilized due to the nature of the research. The three factors, the number of routers involved in actively deploying black hole routing defensive measures, the amount of attack traffic introduced, and the time at which the border routers started dropping attack traffic, were included in the design of this research. This research included an

experiment to establish the baseline system. It was determined the system on which these simulations were to be run could not handle all of the traffic being explicit. A pilot study had to be conducted to determine how much of the baseline traffic could be background traffic and still be able to obtain meaningful metrics. The pilot study showed no significant difference in the performance metrics collected between 90 percent background traffic, 75 percent background traffic, 50 percent background traffic, and 10 percent background traffic. Therefore the baseline system was configured with 90 percent being background traffic and 10 percent being explicit traffic to aid in the speed of the simulations. The number of replications necessary to establish a baseline is five. The mean, standard deviation and variance is calculated for the bandwidth utilization, latency, and router queuing delay from these five runs. A 90 percent confidence interval was then calculated for each of the three parameters. The baseline was validated by ensuring the empirical data obtained from the AFNOC falls within the specified 90 percent confidence interval.

After establishing a valid baseline system, three experiments were run to simulate the small, medium, and large DDoS attack scenarios. These experiments were run to determine if BGP black hole routing caused any unusual problems with respect to the performance metrics. In addition, these experiments were used to determine if BGP black hole routing effectively defended the network from the three different attack scenarios. In order to validate that these simulations were actually configured correctly, the packets received by the target under attack was verified to be zero. There are currently no known external studies to which the data obtained from these simulations could be validated against.

The next 24 experiments tested the effectiveness of BGP black hole routing when one or more border routers didn't participate in defending the network. Six of these experiments were run to simulate five border routers actively black hole routing attack packets. Three of these six were run to simulate the small, medium, and large DDoS attacks against a 9 Mbps pipe and the other three were run to simulate the three DDoS attacks against a 40 Mbps pipe. The next six were run to simulate three border routers actively black hole routing attack packets. Three of these simulated the small, medium, and large DDoS attacks attacking a 9 Mbps pipe and three simulated the three DDoS attacks against a 40 Mbps pipe. Six were run to simulate one border router, other than the border router attached directly to the base under attack, actively black hole routing attack packets. Three of these simulated the three DDoS attack scenarios against a 9 Mbps pipe and the other three simulated the three DDoS attack scenarios against a 40 Mbps pipe. The final six were run to simulate the border router attached to the base under attack being the only router actively black hole routing attack packets. Three of these simulated the three DDoS attack scenarios against a 9 Mbps pipe and three simulated the three DDoS attacks against a 40 Mbps pipe. Once again to verify that these simulations were run correctly, the packets received by the system under attack was explored. The simulation was verified as correct as long as the system under attack received the correct amount of packets. There are no known existing studies with data that the data from these simulations could be validated against.

Six additional experiments tested the effectiveness of remotely deploying BGP black hole routing after a network is already under a DDoS attack. A bug was found in Opnet while setting up these experiments. Opnet Modeler didn't support the "tag"

attribute that is used by the trigger router. A workaround was used: the route-map advertised by the trigger router matched on the IP address instead of a tag equal to 66. The border routers received a static route announcement from the trigger router and then received the route-map and since the IP address of the static route matched the IP address of the route-map the border routers updated their routing tables to drop any more traffic destined to that particular IP address. Each of the six experiments simulated allowing the attack traffic to penetrate the network for ten seconds before the trigger router sent the update to the border routers. Three experiments simulated a 9 Mbps pipe being attacked by the small, medium, and large DDoS attacks respectively and three experiments simulated a 40 Mbps pipe being attacked by the small, medium, and large DDoS attacks. The routing tables from each of the border routers and the packets received by the system under attack were looked at to determine if these simulations ran as expected. These simulations were deemed verified as long as each border router had the routing update and the system under attack received no more packets after the update. There are no known related studies to validate the data obtained by these simulations against.

The final six experiments tested the effectiveness of deploying BGP black hole routing from the base router after a network is already under a DDoS attack. The same workaround as discussed above was used in these six experiments. Another problem was discovered while setting up these experiments. BGP updates are sent via Transmission Control Protocol (TCP) packets. In a study by Li Xiao, Guanghui He and Klara Nahrstedt from the University of Illinois at Urbana-Champaign, it was determined BGP failures occur in bandwidth-saturated networks [XHN05]. Since this scenario is dealing with DDoS attacks larger than the communication links under attack, a modification to

the border routers had to be made in order for the simulation to work. Xiao, He, and Nahrstedt research proved by setting the backoff timer to 30 seconds and the maximum retransmission timeout (RTO) value to eight seconds on the routers that the BGP session was improved significantly [XHN05]. Therefore, the maximum RTO for this experiment was also set to eight seconds in order for the simulation to work. Furthermore, to compare like results the maximum RTO was set to eight seconds to obtain the results for the remote-triggered BGP black hole routing experiments, discussed in the previous paragraph. Each of the six experiments simulated allowing the attack traffic to penetrate the network for ten seconds before the base router sent the update to the border routers. The six experiments simulated the small, medium, and large DDoS attacks attacking the 9 Mbps and 40 Mbps pipes, respectively. These simulations were deemed verified as long as each border router had the routing update and the system under attack received no more packets after the update. There are no known related studies to validate the data obtained by these simulations against.

Each of these 39 experiments was replicated five times by using the same five seeds for the random number generator. The mean, standard deviation and variance was obtained from each of the metrics of these 39 experiments to calculate a 90 percent confidence interval. The data was validated once it was verified that four out of five runs possessed metric means within the specified 90 percent confidence interval.

Analysis of Results

Once all of the data was collected, it was analyzed to either support or disprove the hypotheses. An analysis was conducted for each of the metrics.

To determine which metrics were better, comparisons between them were made. The queuing delay of the baseline routers was compared against the queuing delay of the routers under a DDoS attack. The baseline data was used to derive a 90% confidence interval and by conducting a visual test it was determined if the queuing delay of the routers under attack were worse or not. The performance metric is a Lower Better (LB) metric, which means the least amount of queuing delay the better. The latency data from the baseline was compared to the latency data of the system under a DDoS attack. Since the traffic had been validated and latency is a LB metric, a visual test was used to determine which was better. The router convergence delay of each of the attack scenarios were compared to each other. Again this metric is a LB metric; therefore a visual test was used to determine whether the convergence delay increased with the amount of attack traffic. Bandwidth utilization of the baseline system was compared against the bandwidth utilization of the system under attack. Once again, this is a LB metric and a visual test was used to determine which was better.

The need to determine a confidence level is addressed next. A 99% confidence interval (CI) would increase the validity of this research, but due to time constraints a lower confidence level was used. With the number of tests conducted in this research, a 90% confidence level was appropriate.

Since lower confidence levels lead to an increased chance of the CI of two test results overlapping, there needed to be another approach to determining which metric was better. In this research, additional replications were run in the case of overlapping confidence intervals. The data was then recalculated. This process repeated itself until the confidence intervals did not overlap.

Summary

This chapter followed a systematic approach to performance evaluation to define the problem to be solved, define the system, address the parameters and factors of the system, define the evaluation technique used, justify the experimental design, and to address how the analysis of the results will be handled. A model of the NIPRNET established a baseline system to compare the results of 39 different experiments against. The goal of this research is to determine whether BGP black hole routing is an effective defense technique against DDoS attacks, to determine if a large network such as the NIPRNET suffers any adverse affects when using black hole routing, how much time is needed to restore border routers ACLs after a DDoS attack, how effective BGP black hole routing is when 1 or more border routers are not actively defending the network, and how successful a large network like the NIPRNET is in deploying black hole routing ACLs during a DDoS attack. The factors used in the 39 experiments, the amount of DDoS traffic, the number of routers actively participating in black hole routing, and the time set for the routers to start dropping packets, varied between the experiments.

IV. Results and Analysis

Introduction

Chapter 3 explained the methodology used to conduct this research, so this chapter focuses on the results and analysis of this research. To reiterate, the purpose of this thesis is to study the effects of BGP black hole routing on a network like the NIPRNET and to answer the following five questions:

- 1) Does BGP black hole routing have any adverse effects on the normal operation of a network like the NIPRNET?
- 2) Is BGP black hole routing effective in defending a network like the NIPRNET from DDoS attacks?
- 3) Is BGP black hole routing effective when not all of the border routers are participating in the black hole routing?
- 4) Is remotely triggering BGP black hole routing effective on a network like the NIPRNET while it is under a DDoS attack?
- 5) Can customer-triggered BGP black hole routing be as effective as remote-triggered black hole routing in defending a network under attack?

Effectiveness of BGP black hole routing

BGP black hole routing does not have any adverse effects on the normal operation of a network like the NIPRNET and it does successfully defend a network like the NIPRNET. To come to this conclusion the bandwidth utilization, router queuing delay,

and latency data are compared between the baseline system and the system being defended by BGP black hole routing against the three different DDoS attack scenarios.

The first metric to be discussed is bandwidth utilization. Table 1 lists the 90 percent confidence interval of the inbound bandwidth utilization along with the averages obtained from the baseline configuration and the three DDoS attack scenarios. Table 2 lists the 90 percent confidence interval of the outbound bandwidth utilization, as well as the averages obtained from the baseline configuration and the three DDoS attack scenarios.

Table 1. Inbound Utilization Statistics in Percentages

Base	Inbound Utilization 90% Confidence Interval	Baseline Inbound Utilization Avg	12.8 Mbps DDoS Attack Inbound Utilization Avg	38.4 Mbps DDoS Attack Inbound Utilization Avg	64 Mbps DDoS Attack Inbound Utilization Avg
Beale	(32.483, 32.589)	32.536	32.487	32.489	32.506
Vandenberg	(59.149, 59.359)	59.254	59.280	59.293	59.351
Ellsworth	(28.424, 28.532)	28.478	28.505	28.523	28.523
Whiteman	(44.141, 44.326)	44.233	44.299	44.303	44.260
Hill	(48.808, 49.009)	48.909	48.905	48.865	48.891
Mt Home	(54.525, 54.703)	54.614	54.609	54.594	54.617
Goodfellow	(41.002, 41.156)	41.079	41.079	41.070	41.076
Randolph	(50.851, 51.050)	50.950	50.943	50.966	50.958
Moody	(54.668, 54.852)	54.759	54.775	54.752	54.803
Robins	(79.584, 79.811)	79.697	79.736	79.738	79.708
Little Rock	(60.742, 60.988)	60.865	60.756	60.770	60.750
Tinker	(64.682, 64.905)	64.793	64.800	64.799	64.791

Table 2. Outbound Utilization Statistics in Percentages

Base	Outbound Utilization 90% Confidence Interval	Baseline Outbound Utilization Avg	12.8 Mbps DDoS Attack Outbound Utilization Avg	38.4 Mbps DDoS Attack Outbound Utilization Avg	64 Mbps DDoS Attack Outbound Utilization Avg
Beale	(8.400, 8.443)	8.422	8.442	8.439	8.441
Vandenberg	(14.663, 14.722)	14.692	14.672	14.671	14.686
Ellsworth	(8.613, 8.655)	8.634	8.615	8.617	8.631
Whiteman	(9.917, 9.960)	9.938	9.924	9.939	9.939
Hill	(19.413, 19.471)	19.442	19.434	19.435	19.443
Mt Home	(11.897, 11.949)	11.923	11.905	11.902	11.918
Goodfellow	(9.081, 9.120)	9.100	9.090	9.087	9.112
Randolph	(61.119, 61.285)	61.202	61.199	61.181	61.192
Moody	(13.476, 13.529)	13.503	13.493	13.501	13.499
Robins	(28.956, 29.021)	28.988	28.982	28.997	28.972
Little Rock	(13.505, 13.561)	13.533	13.560	13.552	13.552
Tinker	(16.567, 16.631)	16.599	16.607	16.604	16.625

The bandwidth utilization is compared by deriving a 90 percent confidence interval for each base inbound and outbound utilization level and conducting a visual analysis of the bandwidth utilization data. Tables 1 and 2 illustrate that each of the base's inbound and outbound utilization averages in each of the three attack scenarios are contained in the 90 percent confidence interval of the baseline system. Therefore, it is concluded that no observable differences exists among the utilization levels of the three attack scenarios and the baseline utilization levels. For that reason, this research verifies BGP black hole routing does not have an adverse effect on bandwidth utilization within a NIPRNET-like network. In addition, this data proves that BGP black hole routing successfully defended the network bandwidth by not allowing any of the attack traffic to use available bandwidth.

The next metric to analyze is the queuing delay of the border routers. The routers were configured to process 100,000 packets per second (pps) which is the service rate, μ . The arrival rate, λ , varied by base between 440 pps and 3785 pps for the baseline simulation and it increased on average by 1095 pps, 3425 pps, and 5912 pps for the 12.8 Mbps, 38.4 Mbps, and 64 Mbps DDoS attacks respectively. The gateway utilization ρ is equal to λ/μ . The mean time spent in queue is equal to $1/\mu(1 - \rho)$. As λ gets larger in the case of the DDoS attacks ρ becomes larger and as ρ becomes larger the denominator of the mean time spent in queue becomes smaller, thus the queuing delay will increase. Figure 10 plots the relationship between the queuing delays of the six routers and the number of packets arriving at the six routers. As displayed by the graph, the queuing delay of the routers increases with the increase in the number of packets the routers receive.




Figure 10 – Queuing Delay of all 4 Simulations

As demonstrated by Figure 10, even though the queuing delay increases, the increase is in nanoseconds, which is minimal and doesn't have a major impact on the overall performance of the network. The raw data from the four scenarios is displayed in Table 3 below. As shown, the largest increase in queuing delay is 696 nanoseconds at Robins.

Table 3. Queuing Delay Averages in Microseconds

Router	Queuing Delay 90% Confidence Interval in Microseconds	Baseline Queuing Delay Avg	12.8 Mbps DDoS Attack Queuing Delay Avg	38.4 Mbps DDoS Attack Queuing Delay Avg	64 Mbps DDoS Attack Queuing Delay Avg
Beale	(10.072, 10.075)	10.074	10.186	10.427	10.697
Denver	(10.042, 10.045)	10.044	10.152	10.394	10.665
Hill	(10.147, 10.150)	10.149	10.261	10.511	10.793
Kelly	(10.222, 10.226)	10.224	10.341	10.607	10.896
Robins	(10.391, 10.395)	10.393	10.520	10.784	11.089
Tinker	(10.222, 10.225)	10.224	10.337	10.594	10.884

The arrival rate was minimal in this research, but even if λ was increased to 99000 pps the queuing delay would be one millisecond, which is still a minimal increase to the

overall performance of the network. As λ approaches μ , the queuing delay approaches infinity, but since border routers are configured to handle the maximum load capability of the communication links they are connected to, the λ would never reach μ . Therefore, queuing delay would not have a major impact on the performance of the network. This data definitely suggests that BGP black hole routing will not adversely affect the queuing delay of the border routers. The only impact that BGP black hole routing has on the queuing delays of the border routers is related to the increase in attack traffic and not to the actual dropping of the packets.

The third and final metric to analyze is the latency of the packets traveling through the network. Tables 4 and 5 below lists the 90 percent confidence intervals of the baseline inbound and outbound latency as well as the baseline inbound and outbound latency averages and the averages of the systems under the three different attack scenarios.

Table 4. Inbound Latency Averages in Milliseconds

Base	Inbound Latency 90% Confidence Interval	Baseline Inbound Latency Avg	12.8 Mbps DDoS Attack Inbound Latency Avg	38.4 Mbps DDoS Attack Inbound Latency Avg	64 Mbps DDoS Attack Inbound Latency Avg
Beale	(137.956, 138.889)	138.422	137.966	138.306	137.973
Ellsworth	(154.584, 154.918)	154.751	154.595	154.624	154.590
Goodfellow	(117.217, 117.516)	117.366	117.494	117.493	117.513
Hill	(113.020, 113.130)	113.075	113.087	113.059	113.109
Little Rock	(142.733, 142.978)	142.856	142.892	142.923	142.797
Moody	(120.720, 120.986)	120.853	120.727	120.745	120.758
Mt Home	(158.923, 159.142)	159.032	159.091	158.965	159.094
Randolph	(89.102, 89.231)	89.166	89.163	89.231	89.221
Robins	(65.980, 66.078)	66.029	66.045	66.076	66.065
Tinker	(90.988, 91.068)	91.028	91.034	91.003	90.996
Vandenberg	(167.728, 167.969)	167.848	167.890	167.739	167.770
Whiteman	(167.276, 167.633)	167.455	167.569	167.307	167.276

Table 5. Outbound Latency Averages in Milliseconds

Base	Outbound Latency 90% Confidence Interval	Baseline Outbound Latency Avg	12.8 Mbps DDoS Attack Outbound Latency Avg	38.4 Mbps DDoS Attack Outbound Latency Avg	64 Mbps DDoS Attack Outbound Latency Avg
Beale	(102.194, 103.079)	102.636	102.195	102.302	102.258
Ellsworth	(115.987, 115.992)	115.989	115.991	115.992	115.991
Goodfellow	(85.174, 85.179)	85.176	85.176	85.176	85.179
Hill	(98.019, 98.020)	98.019	98.019	98.020	98.020
Little Rock	(108.212, 108.216)	108.214	108.216	108.214	108.216
Moody	(93.287, 93.290)	93.289	93.289	93.290	93.290
Mt Home	(131.809, 131.812)	131.811	131.811	131.810	131.811
Randolph	(72.438, 72.439)	72.439	72.439	72.439	72.439
Robins	(54.571, 54.573)	54.572	54.572	54.573	54.573
Tinker	(80.485, 80.486)	80.485	80.486	80.486	80.486
Vandenberg	(142.123, 142.125)	142.124	142.124	142.123	142.124
Whiteman	(123.809, 123.813)	123.811	123.812	123.811	123.811

As shown by the data in Tables 4 and 5, BGP black hole routing had no adverse effects on the inbound or outbound latency of the data being transmitted on the network. The 90% confidence intervals were derived from the data obtained from the baseline network. It is apparent that the averages of each base in both the inbound and outbound latency averages are within the 90% confidence interval. Therefore, with a 90 percent confidence level this research concludes there is no significant difference in the latency of the network due to BGP black hole routing. The equality of the latency data also proves the network is successfully being defended by the black hole routing.

Effectiveness of BGP black hole routing when not all of the border routers are participating in the black hole routing

It seems obvious that BGP black hole routing wouldn't be as effective when not all of the border routers are dropping attack traffic, but this research explores the possibilities and the results are presented in this section. This section includes an analysis of the network being defended by one, three, and five border routers, in which

the border router directly connected to the base under attack is not black hole routing. In addition, it also includes an analysis of the network defended only by the border router directly connected to the base under attack. This section also studies the effects of the three DDoS attacks against two different bandwidth capabilities, which are 9 Mbps and 40 Mbps. The 9 Mbps communication channel belongs to Mt Home and its border router is Hill, while the 40 Mbps communication channel belongs to Randolph and its border router is Kelly. This research reveals the differences in the inbound bandwidth utilization, queuing delay, and inbound latency.

Inbound bandwidth utilization is the first metric to analyze. Bandwidth utilization should decrease with an increase in the number of routers dropping attack packets. It should also decrease when there are fewer attack packets being delivered into the network. Table 6 below lists the utilization averages of the links.

Table 6. Bandwidth Utilization Averages

Router	Attack Scenario	90% Confidence Interval of Fully Protected Network	Fully Protected Network	Network Defended by Border Router Connected to Base Under Attack	Network Defended by 5 Routers	Network Defended by 3 Routers	Network Defended by 1 Router
Hill	12.8 Mbps	(54.497, 54.721)	54.609	54.739	78.487	99.266	100.000
	38.4 Mbps	(54.501, 54.688)	54.594	54.658	99.287	100.000	100.000
	64 Mbps	(54.521, 54.712)	54.617	54.669	100.000	100.000	100.000
Kelly	12.8 Mbps	(50.863, 51.023)	50.943	51.010	56.923	64.872	77.769
	38.4 Mbps	(50.868, 51.064)	50.966	51.022	68.128	91.682	99.840
	64 Mbps	(50.855, 51.060)	50.958	50.984	78.016	99.847	100.000

Figures 11 through 16 plot the inbound bandwidth utilized for the communication links under the three different attack scenarios. The remainder of the communication links within the network was not affected by the DDoS attacks.




Figure 11 – 9 Mbps Pipe Defended Against a 12.8 Mbps Attack




Figure 12 – 40 Mbps Pipe Defended Against a 12.8 Mbps Attack




Figure 13 – 9 Mbps Pipe Defended Against a 38.4 Mbps Attack




Figure 14 – 40 Mbps Pipe Defended Against a 38.4 Mbps Attack




Figure 15 – 9 Mbps Pipe Defended Against a 64 Mbps Attack




Figure 16 – 40 Mbps Pipe Defended Against a 64 Mbps Attack

As illustrated by Figures 11 through 16, the inbound bandwidth utilization decreases as the number of routers dropping attack packets increases and the amount of attack traffic on the network decreases. The one exception to that is when the border router connected to the base is the only router dropping attack packets, and in that case the border router is dropping all of the attack traffic, therefore the bandwidth utilization will not change. This is only beneficial if it is known what border router is connected to the base under attack. With the assumption the DDoS traffic is evenly distributed among all of the border routers, this research proves that the inbound bandwidth utilization will increase when one or more routers are not configured to drop the attack traffic, with the one exception mentioned. It would be difficult to attempt to defend bandwidth utilization of the communication link under attack without having all of the border routers successfully dropping attack traffic.

The second metric to discuss is queuing delay. As discussed in the previous section of this chapter, the queuing delay of all the routers increases due to the increase in the amount of traffic being processed by the routers. The amount of increase in the border routers not directly connected to the base router under attack is insignificant, as shown for the 64 Mbps attack in Table 7 below.

Table 7. Queuing Delay Averages in Microseconds

Router	All Routers Defending Against 64 Mbps Attack	One Router Defending Against 64 Mbps Attack Against 9 Mbps Pipe	One Router Defending Against 64 Mbps Attack Against 40 Mbps Pipe	Base Border Router Defending Against 64 Mbps Attack Against 9 Mbps Pipe	Base Border Router Defending Against 64 Mbps Attack Against 40 Mbps Pipe	Three Routers Defending Against 64 Mbps Attack Against 9 Mbps Pipe	Three Routers Defending Against 64 Mbps Attack Against 40 Mbps Pipe	Five Routers Defending Against 64 Mbps Attack Against 9 Mbps Pipe	Five Routers Defending Against 64 Mbps Attack Against 40 Mbps Pipe
Beale	10.697	10.913	10.913	10.911	10.913	10.913	10.913	10.723	10.723
Denver	10.665	10.667	10.665	10.887	10.891	10.676	10.672	10.686	10.686
Hill	10.793	18.048	10.974	21.55	10.974	13.897	10.794	11.922	10.813
Kelly	10.896	11.074	18.625	11.074	22.287	10.901	14.108	10.917	12.104
Robins	11.089	11.183	11.181	11.182	11.179	11.183	11.186	11.097	11.098
Tinker	10.884	11.031	11.029	11.029	11.03	10.878	10.884	10.896	10.893

Table 7 illustrates that when a border router is not dropping attack packets, the queuing delay increases at most 216 nanoseconds, as in the case of the Beale router. Beale's queuing delay is 10.697 microseconds when it is configured to drop packets and the highest its queuing delay reaches is 10.913 microseconds when it is not configured to drop packets. Since 216 nanoseconds is the worst case of increase, this section will focus on the border routers directly connected to the bases under attack. Figures 17 and 18 plot the results of the network being attacked by the three different DDoS attack scenarios. As shown, the queuing delay is related to the number of routers protecting the network. Table 8 further demonstrates that even though the queuing delays decrease with fewer packets and more routers defending, the average queuing delays of the routers are worse.




Figure 17 – Queuing Delay of Border Router Connected to 9 Mbps Pipe




Figure 18 – Queuing Delay of Border Router Connected to 40 Mbps Pipe

Table 8. Queuing Delay Averages of Routers in Microseconds

Router	Attack Scenario	Queuing Delay 90% Confidence Interval with Network Fully Protected	Queuing Delay with Network Fully Protected	Queuing Delay with Network Defended by 5 Routers	Queuing Delay with Network Defended by 3 Routers	Queuing Delay with Network Defended by Denver Router	Queuing Delay with Network only Defended by Border Router Connected to Base
Hill	12.8 Mbps	(10.259, 10.263)	10.261	10.276	10.633	10.961	11.014
	38.4 Mbps	(10.508, 10.605)	10.511	10.612	11.990	13.487	14.144
	64 Mbps	(10.790, 10.797)	10.793	10.973	13.897	18.048	21.550
Kelly	12.8 Mbps	(10.339, 10.343)	10.341	10.363	10.726	11.078	11.111
	38.4 Mbps	(10.514, 10.609)	10.607	10.702	12.126	13.724	14.324
	64 Mbps	(10.893, 10.899)	10.896	11.076	14.108	18.625	22.287

An interesting finding to note is that the router connected to the 9 Mbps pipe consistently has a lower queuing delay than the router connected to the 40 Mbps pipe. The reason for

this is the router's queues had to be configured to only hold one second of traffic to allow the simulations to run due to memory constraints of the system on which the simulations were run. Therefore, the router connected to the 9 Mbps pipe dropped more overall traffic due to a smaller queue and thus resulting in a shorter queuing delay for packets it actually processed. As stated in chapter 3, it was not feasible to determine which packets were valid and which were attack traffic, therefore packets dropped were not measured as a metric for this research. As stated in the previous section, the queuing delay of the routers is affected by the amount of attack traffic on the network. This research also demonstrates that queuing delay is also affected by the number of routers dropping attack traffic. Even with the border router connected to the base under attack being the only router dropping attack traffic, the queuing delay is increased due to the fact it has to drop all of the attack traffic. Therefore, this data clarifies the obvious that a reduction in the number of routers dropping attack packets does result in an insignificant increase of the queuing delay at the border router connected to the base under attack. In addition, as stated in the previous section, the queuing delay of the border routers shouldn't have a major impact on the performance of the network due to the robustness of the routers.

The last metric to analyze is the inbound latency of the communication links under attack. The latency of the other communication links within the network shows no significant difference from the baseline system. The inbound delay of the bases under attack should increase with each increase in attack traffic being allowed to pass through the network. Figures 19 and 20 plot the latency data obtained in this research. Latency is determined by propagation, transmit, and queue. Propagation is defined as the distance traveled by the speed of light. In this simulation, the propagation will not change due to

the increase in attack traffic, so propagation was disregarded as a reason for an increase in latency. Transmit is defined by the size of the packet divided by the bandwidth.

Therefore, transmit did not change in this research either.




Figure 19 – Inbound Latency of Base with 9 Mbps Pipe




Figure 20 – Inbound Latency of Base with 40 Mbps Pipe

The only logical explanation for an increase in latency is due to an increase in queuing delay. As shown previously in this section, the queuing delay of the routers was only increased by microseconds. A more thorough look revealed that the increase in latency was caused by the communication link being saturated. Table 9 lists the amount of milliseconds that were added to the latency under each scenario due to the communication link's queue.

Table 9. Average Added Queuing Delay of Communication Link in Milliseconds

Router	Attack Scenario	Hill with Protection by Hill Only	Kelly with Protection by Kelly Only	Hill with Protection by Denver Only	Kelly with Protection by Denver Only	Hill with Protection by Denver, Kelly, and Tinker	Kelly with Protection by Denver, Hill, and Tinker	Hill with Protection by every Router except Hill	Kelly with Protection by every Router except Kelly	Hill with Full Protection	Kelly with Full Protection
Hill	12.8 Mbps	0.008		67.843		0.974		0.063		0.008	
	38.4 Mbps	0.008		78.003		73.463		1.017		0.008	
	64 Mbps	0.008		81.185		78.611		70.395		0.008	
Kelly	12.8 Mbps		0.002		0.019		0.011		0.004		0.002
	38.4 Mbps		0.002		1.061		0.049		0.011		0.002
	64 Mbps		0.002		15.955		1.097		0.020		0.002

As shown by the above figures, the inbound latency is related to the number of routers protecting the network, the amount of attack traffic on the network, and the bandwidth utilization of the base under attack. The inbound latency of the network with only the border router connected to the base under attack is the same as the network when all six routers are dropping attack traffic. This is due to the fact that the border router is dropping all of the attack traffic and thus not overloading the link's queue. Therefore, by knowing which border router the system under attack is connected to it is feasible to drop the attack traffic at that border router only and the inbound latency will not be affected. The data also suggests that if you know the exact amount of bandwidth available on a communication link and the exact amount of attack traffic coming inbound, that you could possibly configure the network to only have a portion of the routers dropping attack

packets and the inbound latency would not be affected. The data more strongly suggests, it is wiser to ensure every border router drops attack packets to ensure inbound latency is minimally affected.

Effectiveness of remotely triggering BGP black hole routing on a network like the NIPRNET while it is under a DDoS attack

To determine the feasibility of remotely triggering border routers after an attack has started, this research looks at the amount of time it takes the border routers to converge with the update sent from the trigger router to start dropping attack traffic. This research also looks at how fast the network recovers once the BGP black hole routing has been triggered, by looking at the bandwidth utilization, router queuing delay, and latency data.

As stated in Chapter 3, these simulations were configured to simulate the attack traffic attacking the network for ten seconds before the trigger router sent the update to start dropping the attack traffic. Figure 21 illustrates the router convergence data obtained. The data illustrates that each of the border routers has been updated to start dropping packets in no more time than 32 milliseconds. The data obtained during this research also supports that each border router is dropping attack traffic by 200 milliseconds after the update is sent by the trigger router. The data also shows that the routers will take longer to apply the updates with an increase in attack traffic. A 38.4 Mbps attack against a 9 Mbps communication link increases the router's convergence by an average of 73 microseconds over a 12.8 Mbps attack, whereas a 38.4 Mbps attack against a 40 Mbps communication link increases the router's convergence by an average of 37 microseconds over a 12.8 Mbps attack.




Figure 21 – Remote Triggered Update Router Convergence

In addition, a 64 Mbps attack against a 9 Mbps communication link increases the router’s convergence by an average of 52 microseconds over a 38.4 Mbps attack, whereas a 64 Mbps attack against a 40 Mbps communication link increases the router’s convergence by an average of 90 microseconds over a 38.4 Mbps attack. The data implies that on average the routers will take an additional 2 microseconds to converge with each additional 1 Mbps of attack traffic introduced into the network.

The next area to look at is bandwidth utilization. Figures 22 and 23 illustrate how the utilization of the two attacked bases reacts to the attack and to the black hole routing. As shown in the graphs, the utilization rates of the two communication links return to normal around the 200 millisecond point, which is consistent with the router convergence data. Therefore, remotely triggering the border routers defended the networks bandwidth in less than one second.




Figure 22 – Inbound Bandwidth Utilization of 9 Mbps Pipe




Figure 23 – Inbound Bandwidth Utilization of 40 Mbps Pipe

The next metric to explore is the queuing delay of the border routers. Figures 24 and 25 illustrate the data obtained from this research. As expected, the router queuing delays return to normal within 200 milliseconds of the black hole routing update. This data demonstrates that remotely triggering BGP black hole routing does successfully protect the border routers' queuing delays.




Figure 24 – Router Queuing Delay of 9 Mbps Pipe




Figure 25 – Router Queuing Delay of 40 Mbps Pipe

The final metric to examine is the latency of the network. Figures 26 and 27 below graph the inbound latency results obtained. Outbound latency was not affected during these simulations. The data shows that the latency of the network is restored to normal within the 200 milliseconds of the remote-triggered update, once again suggesting

that remote-triggered black hole routing is successful in protecting latency within a network in a minimal amount of time.




Figure 26 – Inbound Latency of 9 Mbps Pipe




Figure 27 – Inbound Latency of 40 Mbps Pipe

This data presented in this section suggests that remotely triggering border routers to black hole attack traffic during a DDoS attack would be highly successful. The large Internet communication links are key to ensuring the updates are delivered in a small amount of time and the network is restored in a reasonable time.

Effectiveness of customer-triggered BGP black hole routing as compared to remote-triggered black hole routing in defending a network under attack

The scenario for testing this goal is to set up the network as in the previous section. The attack traffic is allowed to freely flow through the network for ten seconds before the base router attempts to send an update to the border router it is connected to start dropping attack packets. With customer-triggered black hole routing, the updating of the base router would either have to be a manual process or there would have to be an intrusion detection system connected to the router that would need to be able to update the base router policies via some sort of batch script. In the case of a manual update, ten seconds to detect a DDoS attack and to update the router to send the update to the border router would be rather quick. In the case of an automated update, ten seconds would border on the slow side. Since customer-triggered black hole routing could be accomplished in either fashion, ten seconds was chosen to compare the results from this scenario to the scenario from the previous section. This research looks at the router convergence times, as well as how fast the network recovers once the BGP black hole routing has been triggered.

The convergence time of the border routers is determined by the amount of attack traffic, size of the routers queue, the communication link's delay, and the bandwidth of the communication link. Figure 28 below displays the results obtained from this research. As you can see, the 40 Mbps communication link gets the updates to the border routers rather quickly under the 12.8 and 38.4 Mbps attacks.




Figure 28 - Customer Triggered Update Router Convergence

The reason for this is due to the fact that the border router connected to the 40 Mbps communication link has a larger queue size as described in Chapter 3, the bandwidth available is 40 Mbps, and the delay on the link is 65 milliseconds (ms). Taking all of these factors into consideration, the probability that the BGP TCP packets will be dropped and have to be retransmitted is 0 percent for the 12.8 Mbps attack and 7 percent for the 38.4 Mbps attack. In contrast, for the 9 Mbps communication link under attack, the border router connected to it has a much smaller queue size, an available bandwidth of only 9 Mbps, and a delay of 126 ms. Therefore the probability the BGP TCP packets will be dropped is 34 percent for the 12.8 Mbps attack and 54 percent for the 38.4 Mbps attack. As for the 64 Mbps attack, the probability the BGP TCP packets will be dropped is 21 percent for the 40 Mbps communication link and 75 percent for the 9 Mbps communication link. These probabilities were obtained from the queuing data obtained

during this simulation. Therefore, in order for customer-triggered black hole routing to be effective, it would have to be triggered on a link with the necessary bandwidth and minimal delay as well as a router with a large enough queue to handle the amount of attack traffic on the network. It is interesting to point out that in each of the 64 Mbps scenarios, the border routers attached to the communication link under attack received the update from the base but took extra time to send the update to the other five routers. It is speculated that this phenomenon is attributed to the fact that the BGP updates are sent via TCP and the border routers attached to the bases under attack were still dropping packets due to their queue sizes. Therefore, it took longer for them to complete the TCP three-way handshake with the remaining border routers.

As for how the network responded, the first metric explored is bandwidth utilization. Figures 29 and 30 display the inbound bandwidth utilization data obtained.




Figure 29 – Inbound Bandwidth Utilization of 9 Mbps Pipe

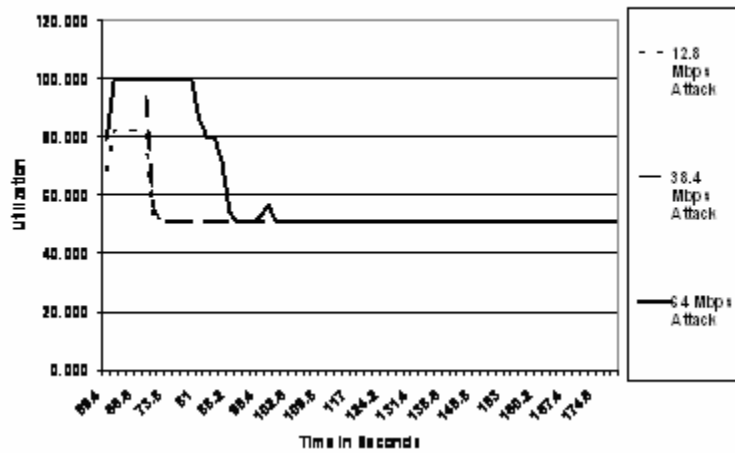


Figure 30 – Inbound Bandwidth Utilization of 40 Mbps Pipe

The data shows that the utilization returns to normal after the successful triggering of the black hole routing. Unfortunately, the bandwidth is consumed for an amount of time that is directly related to the amount of time it takes the border routers to receive the update from the base and to converge.

The next metric to examine is the queuing delay of the border routers. Figures 31 and 32 below graph the results. As demonstrated, the queuing delay of the border routers




Figure 31 – Router Queuing Delay of 9 Mbps Pipe




Figure 32 – Router Queuing Delay of 40 Mbps Pipe

is greatly impacted by the 64 Mbps attack. It elevates the queuing delay from microseconds to seconds. This is a major impact to a network that is constructed to meet certain delay averages. A network like the NIPRNET would definitely be impacted if its router’s queuing delays were elevated into seconds.

The final metric is latency. Figures 33 and 34 below graph the results obtained.




Figure 33 – Inbound Latency of 9 Mbps Pipe




Figure 34 – Inbound Latency of 40 Mbps Pipe

Once again, the network is highly affected by the 64 Mbps attack. The 64 Mbps attack also affects the outbound latency as illustrated in Figures 35 and 36 below.




Figure 35 – Outbound Latency of 9 Mbps Pipe under 64 Mbps Attack




Figure 36 – Outbound Latency of 40 Mbps Pipe under 64 Mbps Attack

Most of the latency data can be attributed to the fact that the router’s queuing delays were elevated into the seconds. There is a direct correlation between the outbound latency data and the router’s queuing delay data. These latency results would definitely be unacceptable on a network like the NIPRNET.

This section suggests that customer-triggered BGP black hole routing isn’t as effective as remotely triggering border routers via more robust communication links. Since BGP uses TCP packets to communicate and pass updates, the links between the bases and the border routers don’t seem to be large enough to successfully automate the triggering of the border routers to start black hole routing attack traffic.

Summary

The evaluation and analysis of the goals proposed in Chapter 3 was accomplished in this chapter. The model output data was analyzed to assist in providing answers to the questions posed by the goals. Some of the questions have distinctive yes or no answers,

while others do not, based on the factors presented in this chapter. Finally, this analysis process demonstrates its merit in aiding decision-makers in determining solutions which best meet a defined set of criteria. With the active participation of the NSA and educational institutions, this process can be iterated, allowing for tradeoff analysis studies.

V. Conclusions

Goal Restatement

The goal of this thesis was to evaluate BGP black hole routing on a network like the NIPRNET. The thesis was to look at whether BGP black hole routing had any adverse effects on the network, whether it was effective in defending the network from a DDoS attack, whether it was effective in defending a network when not all of the routers were black holing the DDoS traffic, how effective was remotely triggering the border routers to black hole when the network was already under attack, and to determine if customer triggered black hole routing was as effective as remote triggered black hole routing. Assumptions were made where network data was not available.

Conclusions

A computer systems analysis approach was used to analyze the abilities of BGP black hole routing under the scenarios presented by the goals. These systems were evaluated in terms of router queuing delays, latency, bandwidth utilization, and router convergence delays. The following five questions were answered:

1) Does BGP black hole routing have any adverse effects on the normal operation of a network like the NIPRNET?

A network like the NIPRNET displayed no adverse effects due to BGP black hole routing. It was pointed out in Chapter 4 that the queuing delays of the border routers did increase when the network was under attack, but this was attributed to the increase in the amount of traffic on the network.

2) Is BGP black hole routing effective in defending a network like the NIPRNET from DDoS attacks?

BGP black hole routing proved to be successful in defending a network like the NIPRNET against DDoS attacks.

3) Is BGP black hole routing effective when not all of the border routers are participating in the black hole routing?

The results were mixed. As for queuing delay, it increases minimally. Therefore, the queuing delay is impacted when not all routers are dropping attack traffic, but unless the DDoS attack is in the multitude of billions of bits per second, the queuing delay will not have an adverse affect on normal operations of a network like the NIPRNET. Latency and bandwidth utilization increase as well, except in the case where the border router connected to the base under attack is the only router dropping packets. It was proven that latency is affected by the queuing delay of the communication link and if the link is 100 percent utilized it will increase the latency of the packets. It was also demonstrated that defending bandwidth utilization will definitely be problematic without using all of the border routers to drop attack traffic. It is feasible to only drop packets by the border router directly connected to the base under attack, but it isn't optimal for the network.

4) Is remotely triggering BGP black hole routing effective on a network like the NIPRNET while it is under a DDoS attack?

Remotely triggering BGP black hole routing is highly effective as long as it is conducted over robust communication links. In each scenario the border

routers were dropping attack traffic in less than 200 milliseconds.

5) Can customer-triggered BGP black hole routing be as effective as remote-triggered black hole routing in defending a network under attack?

Customer-triggered black hole routing is clearly not as effective as remote-triggered black hole routing. This is due to BGP updates being sent via TCP packets and the communication links between the bases and border routers not being of sufficient size to handle the same amount of traffic as the communication links between the border routers. The study revealed that it could take nearly 50 seconds to propagate the update throughout the network and that is definitely not sufficient in terms of defending a network like the NIPRNET.

Contributions

AFIT is now an integral part of the National Security Agencies (NSA) research into BGP black hole routing. This research lays the foundation and framework for all future AFIT work regarding BGP black hole routing. More importantly, this study gives the NSA a baseline from which to work regarding their ongoing efforts to defend the NIPRNET from DDoS attacks. This research proves that BGP black hole routing can successfully be deployed on the NIPRNET to defend against DDoS attacks. It demonstrates that due to the nature of BGP updates, remote-triggered black hole routing is more effective than customer-triggered black hole routing. Finally, this thesis has reinforced AFIT's partnership with the NSA by once again demonstrating AFIT's ability to solve complex operational problems.

Suggestions for Future Work

This research effort only scratched the surface of possibilities with respect to BGP black hole routing. This research only looked at destination-based black hole routing; source-based black hole routing is one area that would expand on this research. In addition, this research simulated the remote triggering and customer triggering of black hole routing. This research would be expanded by configuring an Intrusion Detection System (IDS) in Opnet to actually detect a DDoS attack and send out the BGP update to each of the border routers. In addition, many simplifying assumptions were necessary due to limitations with the simulation software. It would greatly expand this research if an actual scaled down NIPRNET were put in place throughout the United States to obtain some empirical data.

Appendix A: Model Configurations




Figure 37 – Internal Internet Set-up

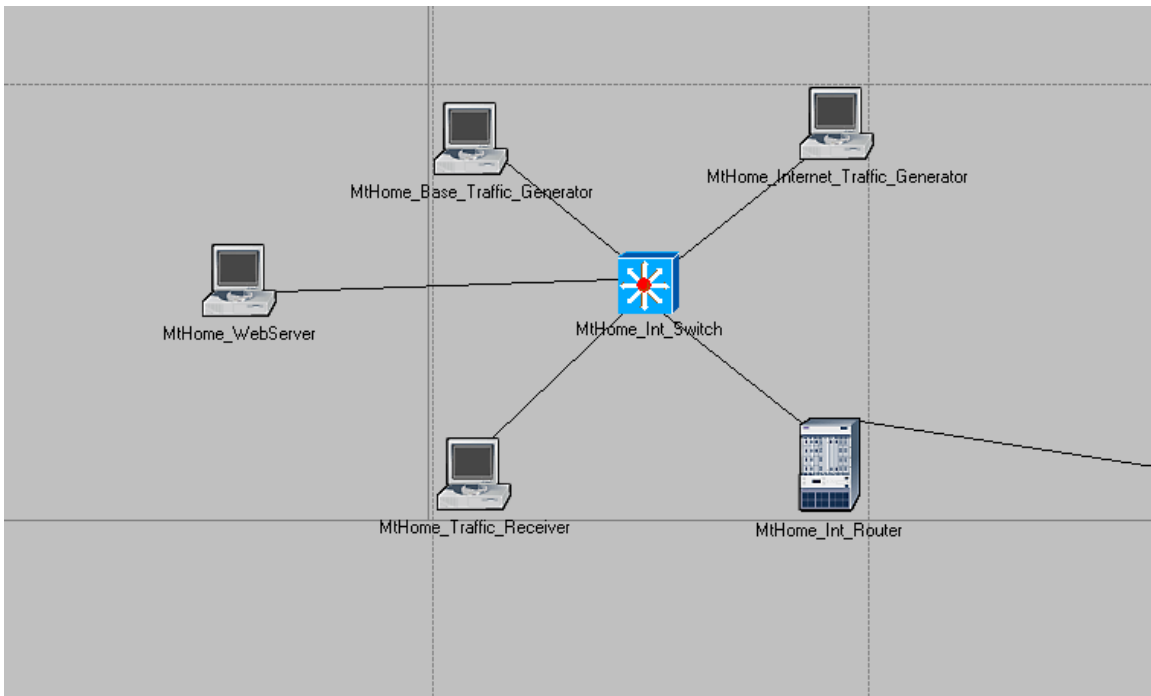


Figure 38 – Internal Base Set-up

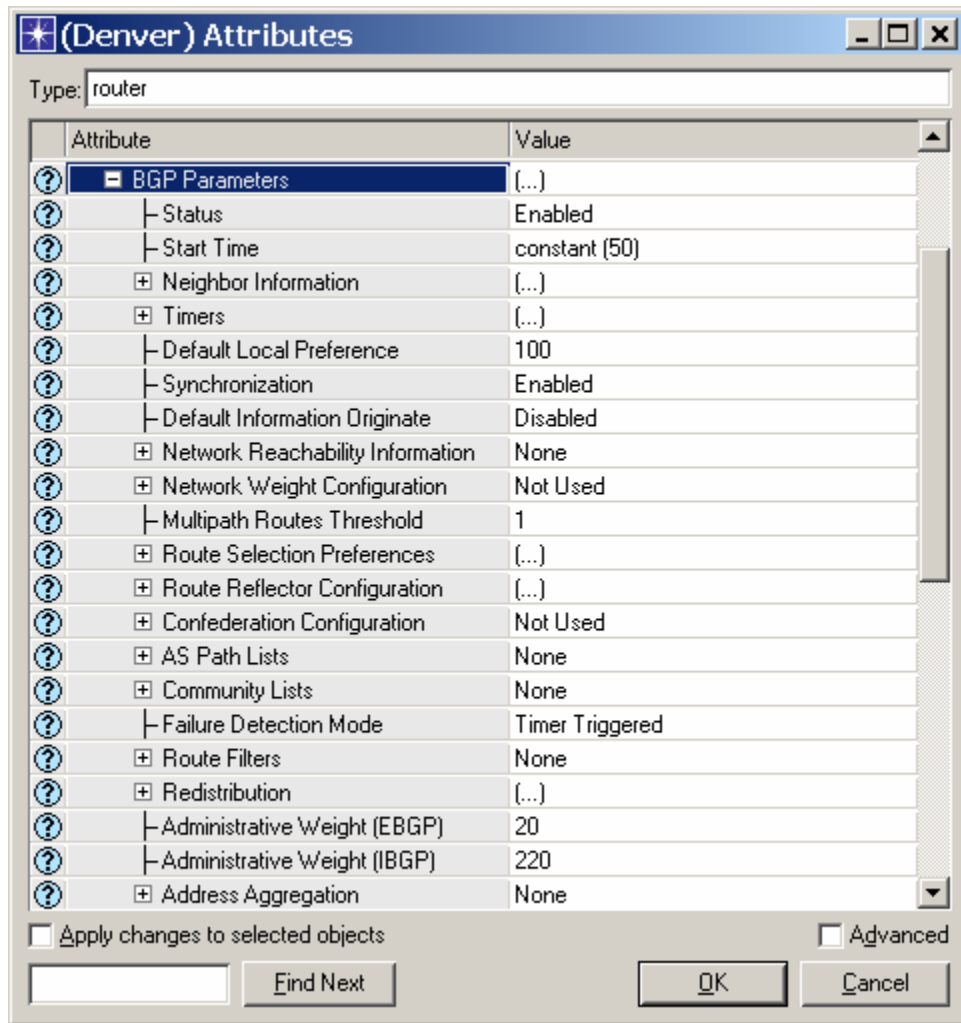


Figure 39 – Border Router BGP Parameters

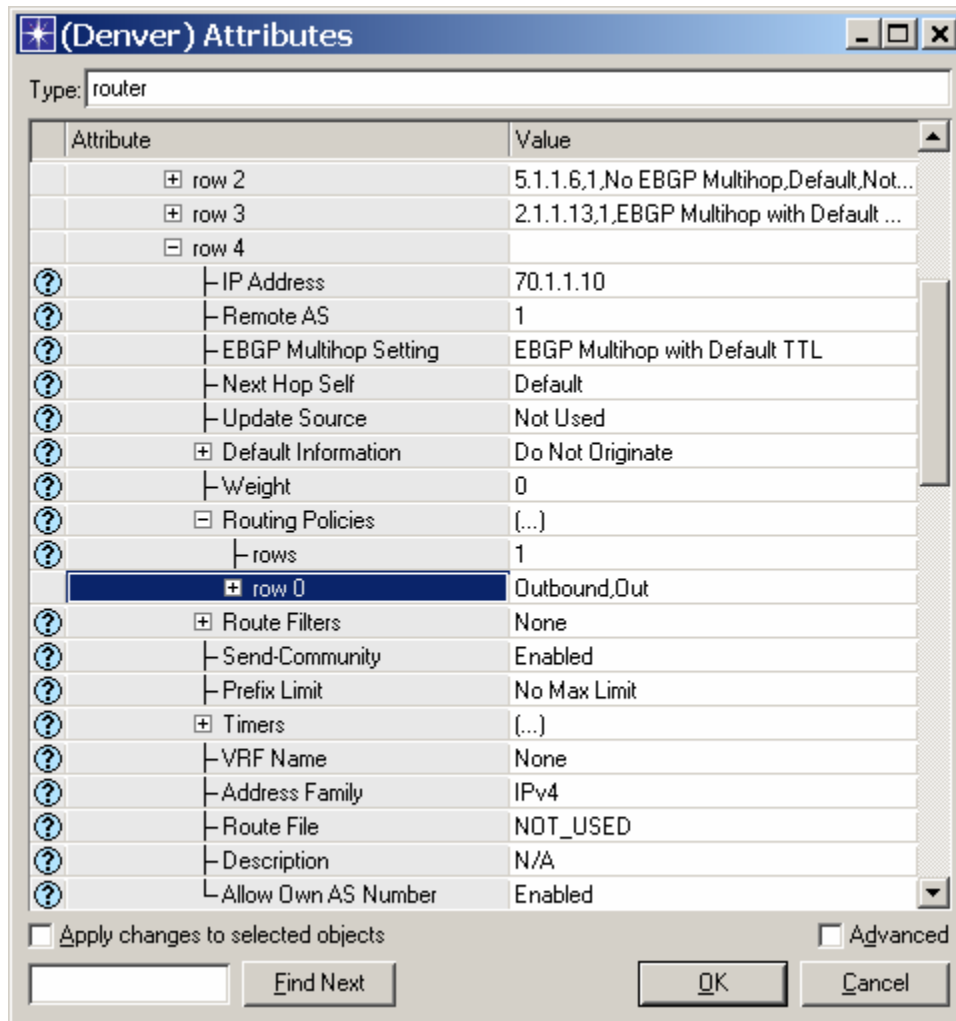


Figure 40 – Border Router BGP Neighbor Information

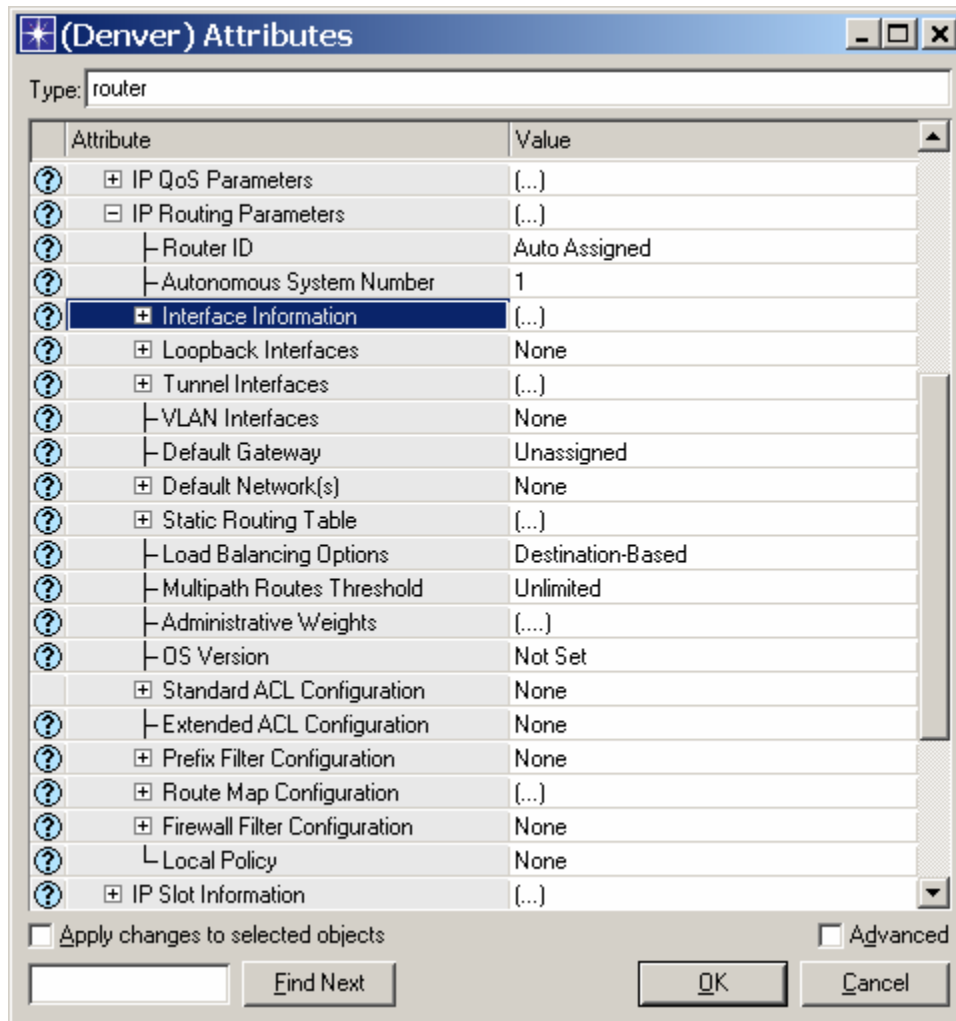


Figure 41 – Border Router IP Routing Parameters

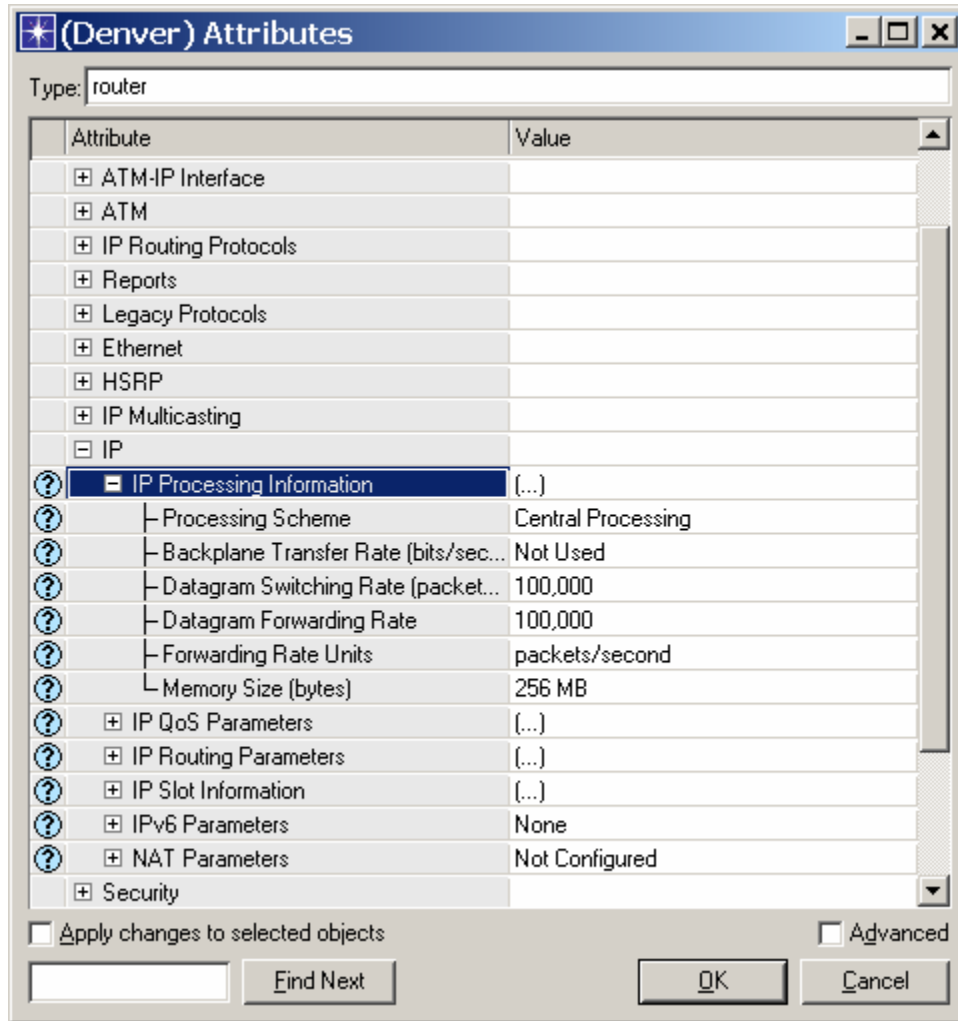


Figure 42 – Border Router IP Processing Information

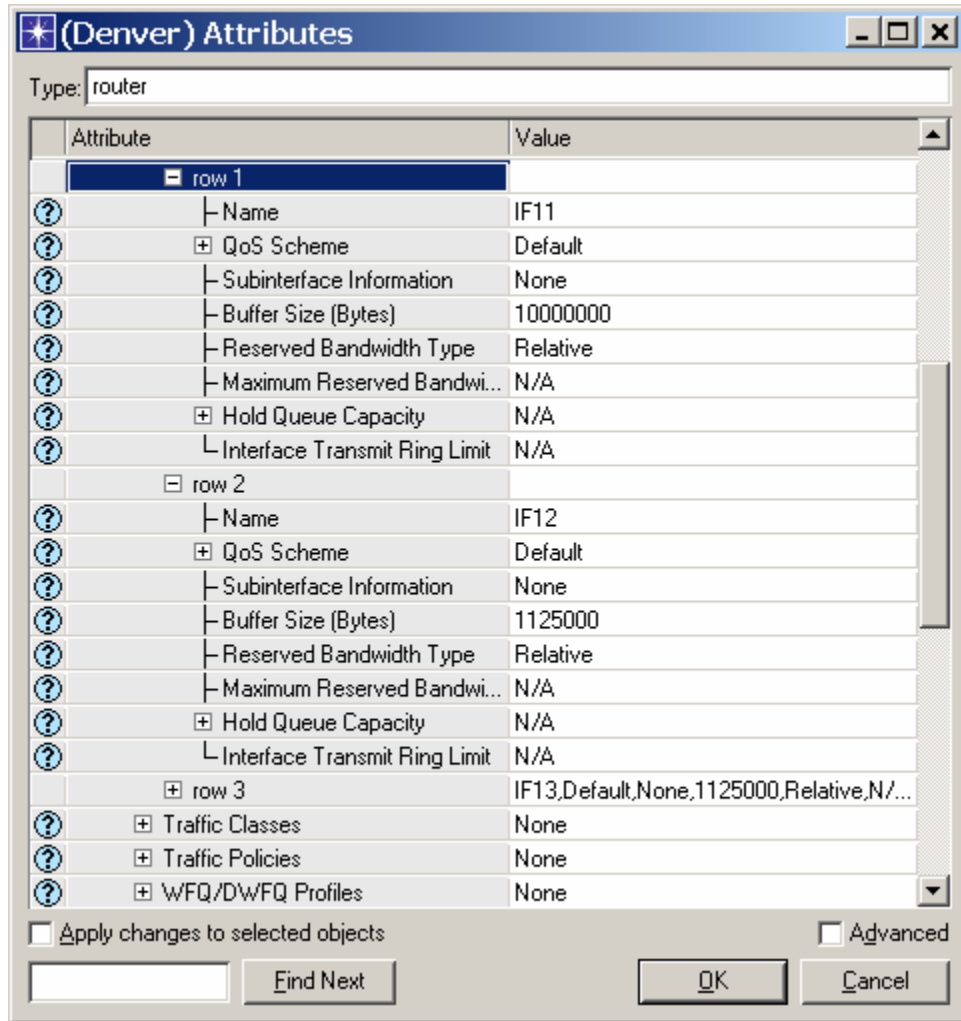


Figure 43 – Border Router IP Quality of Service Configuration of 9 Mbps Pipe

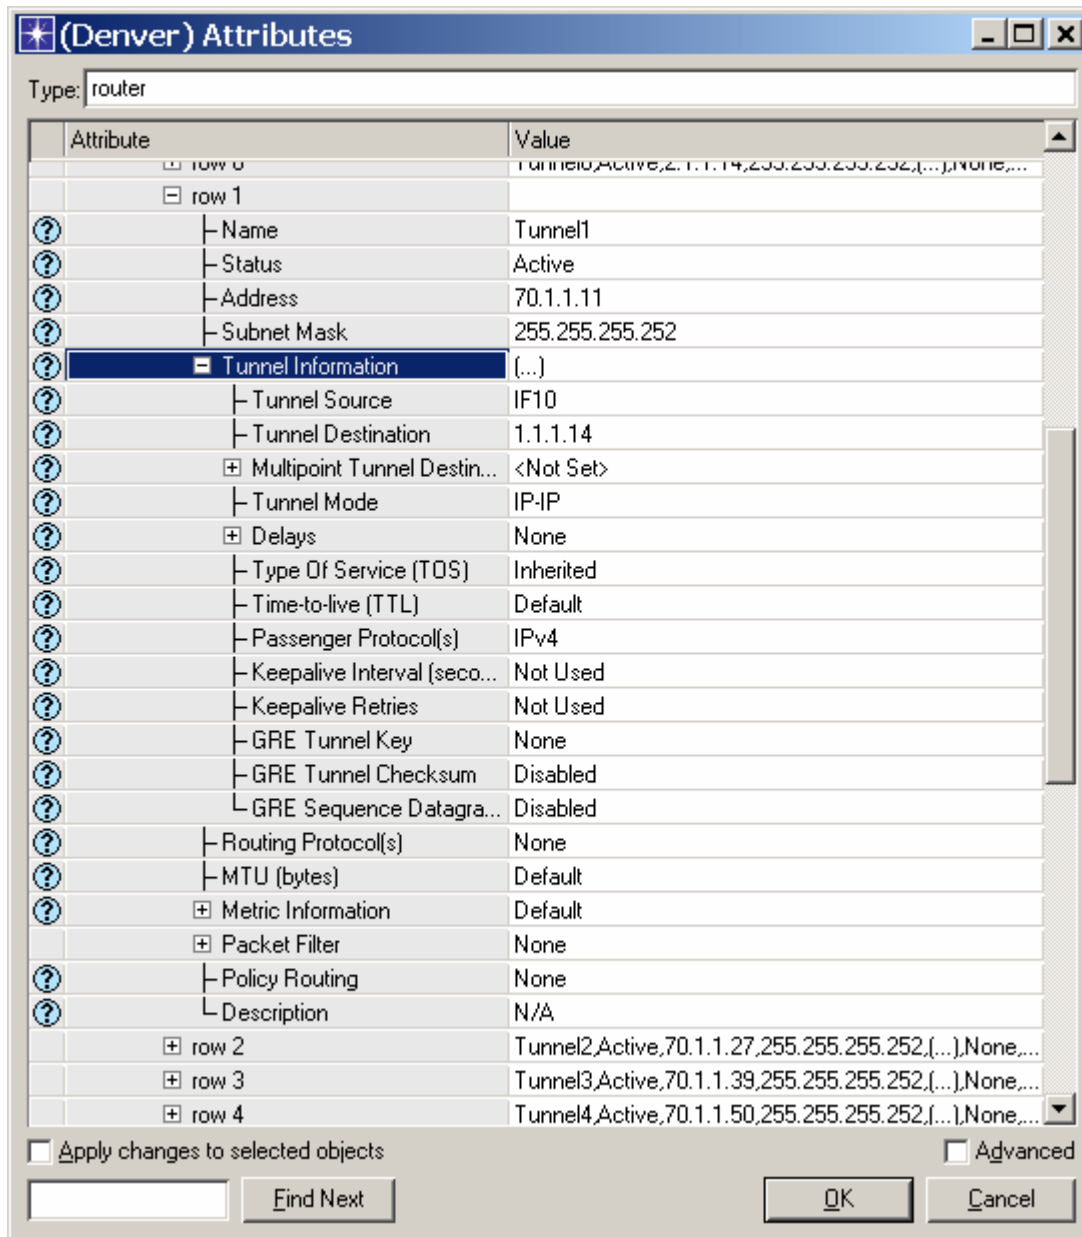


Figure 44 – Border Router IP Tunnel Information

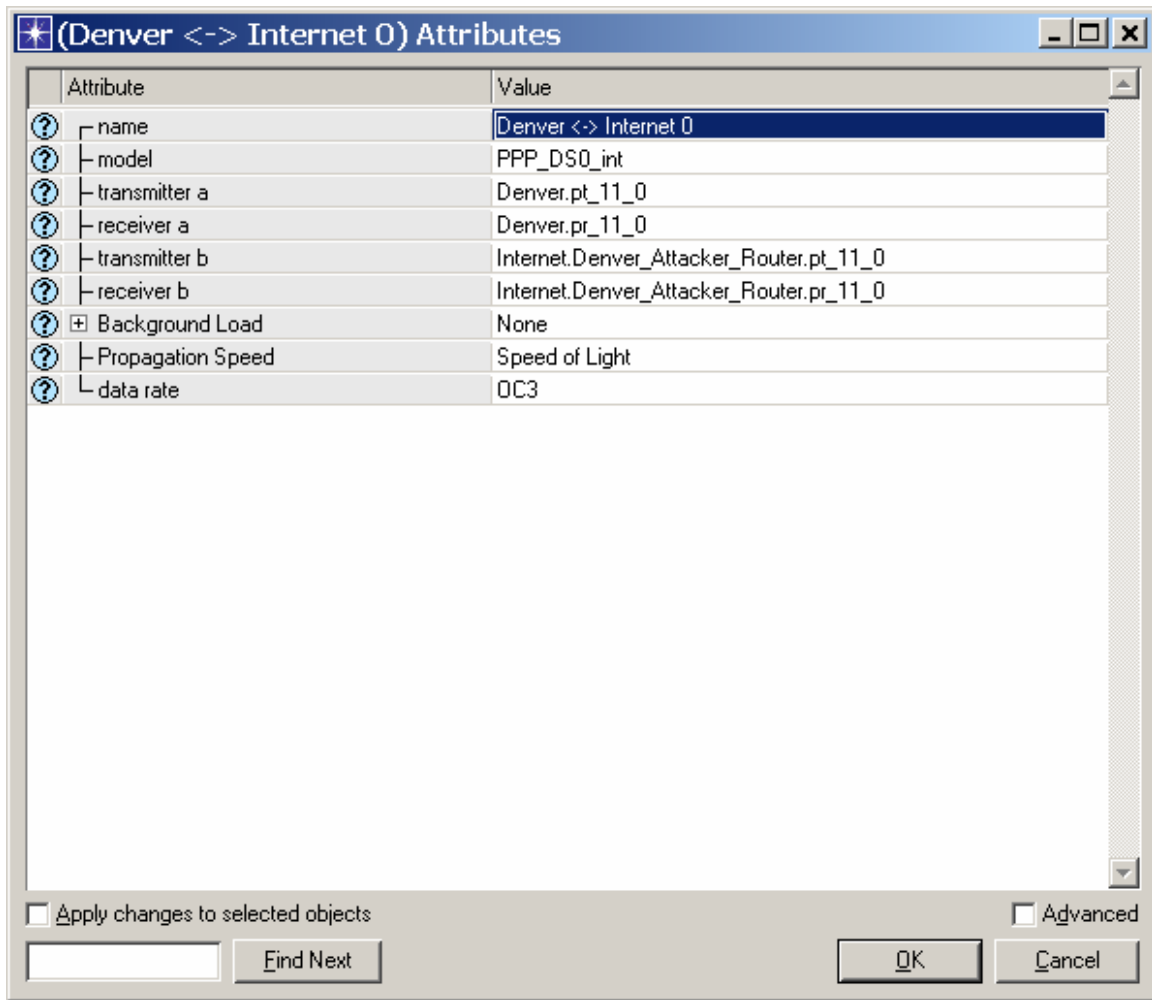


Figure 45 – Border Router to Internet Link Configuration

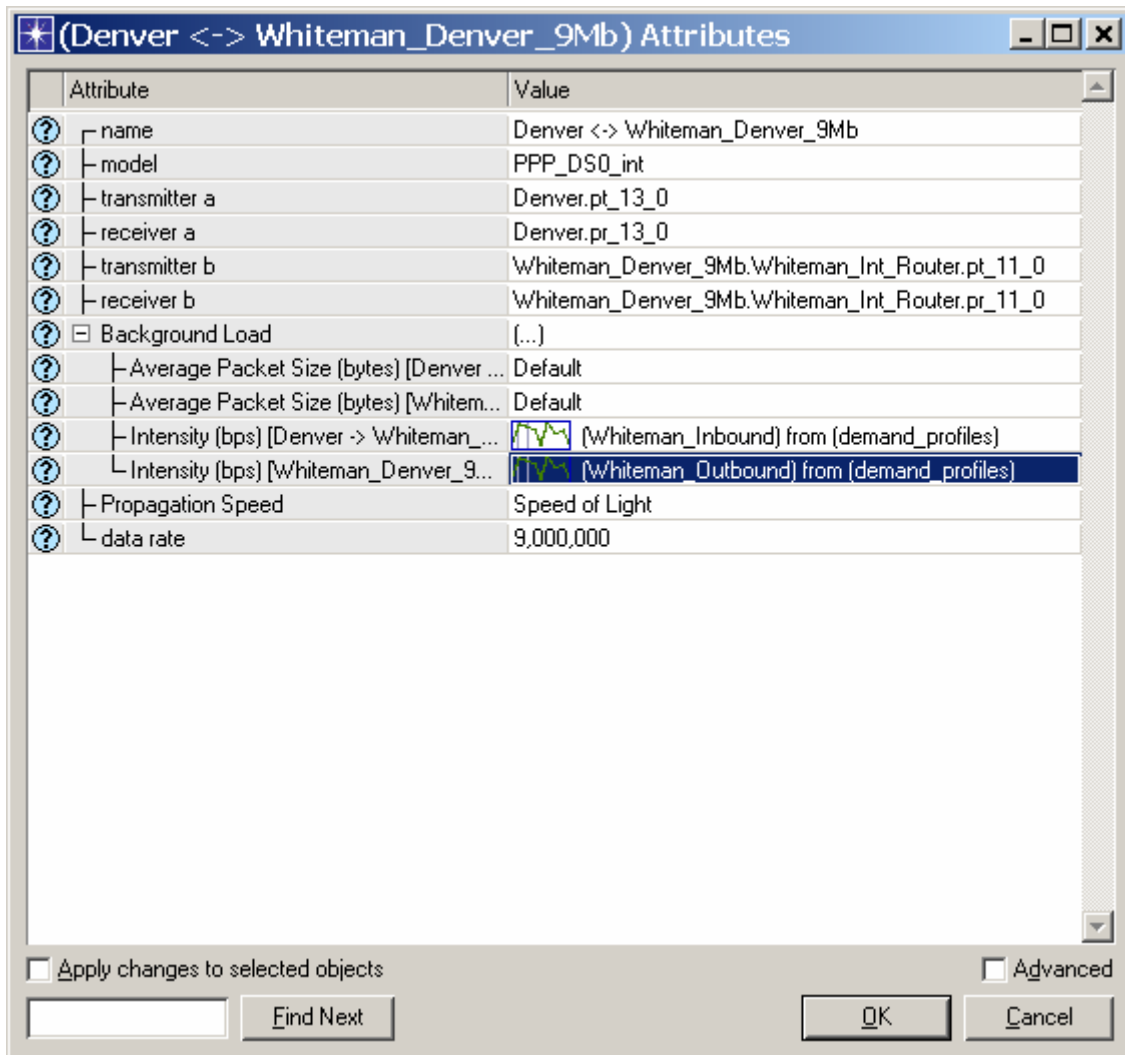


Figure 46 – Border Router to Base Link Configuration

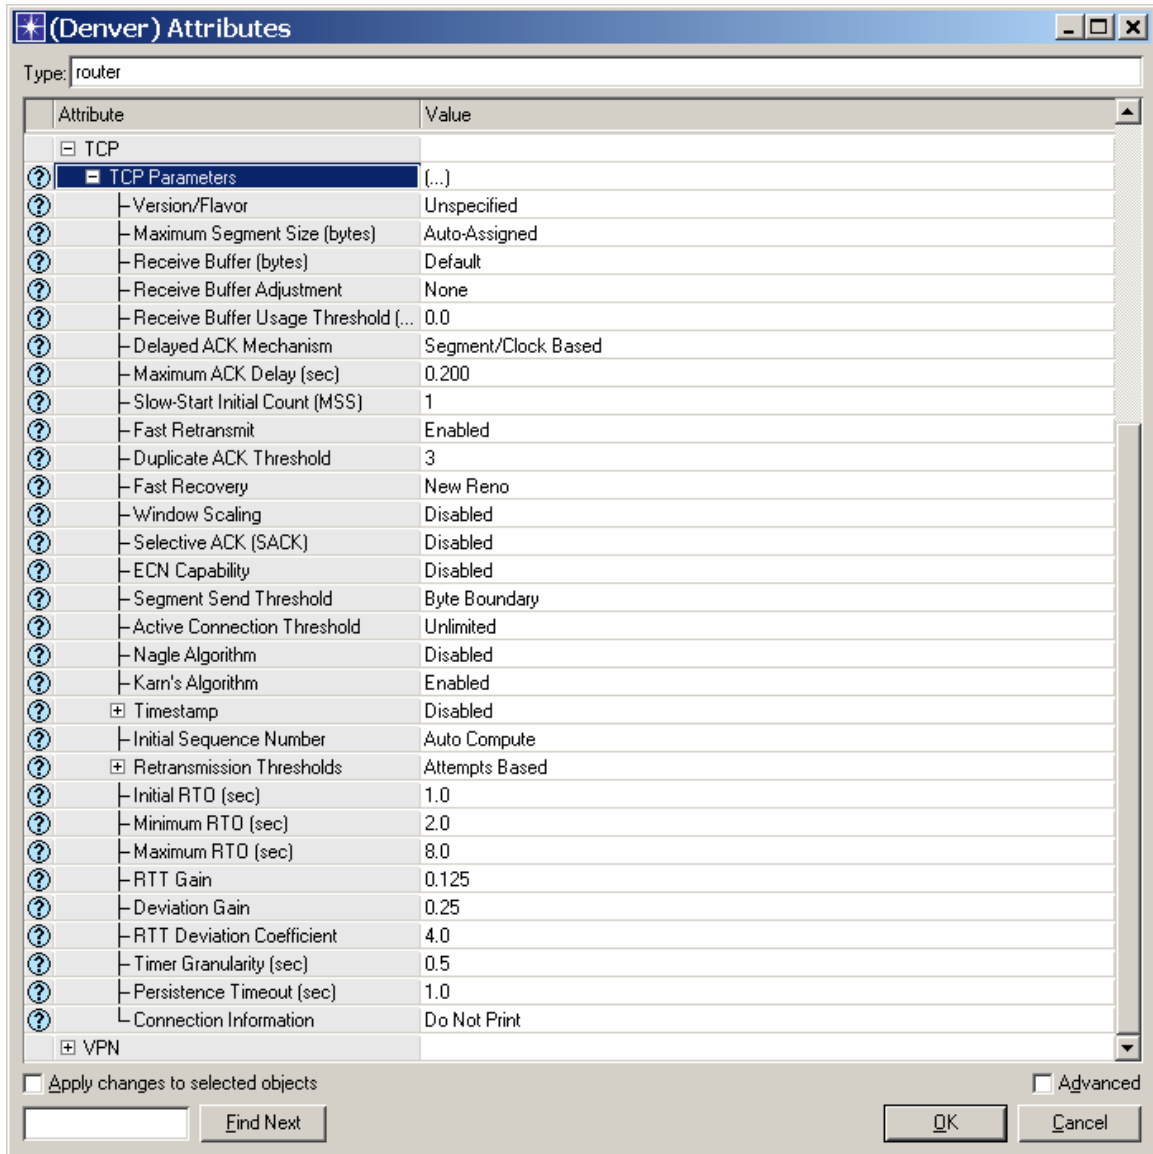


Figure 47 – TCP Settings on All Routers




Figure 48 – Border Router Route Map Configuration




Figure 49 – Border Router Static Routing Table

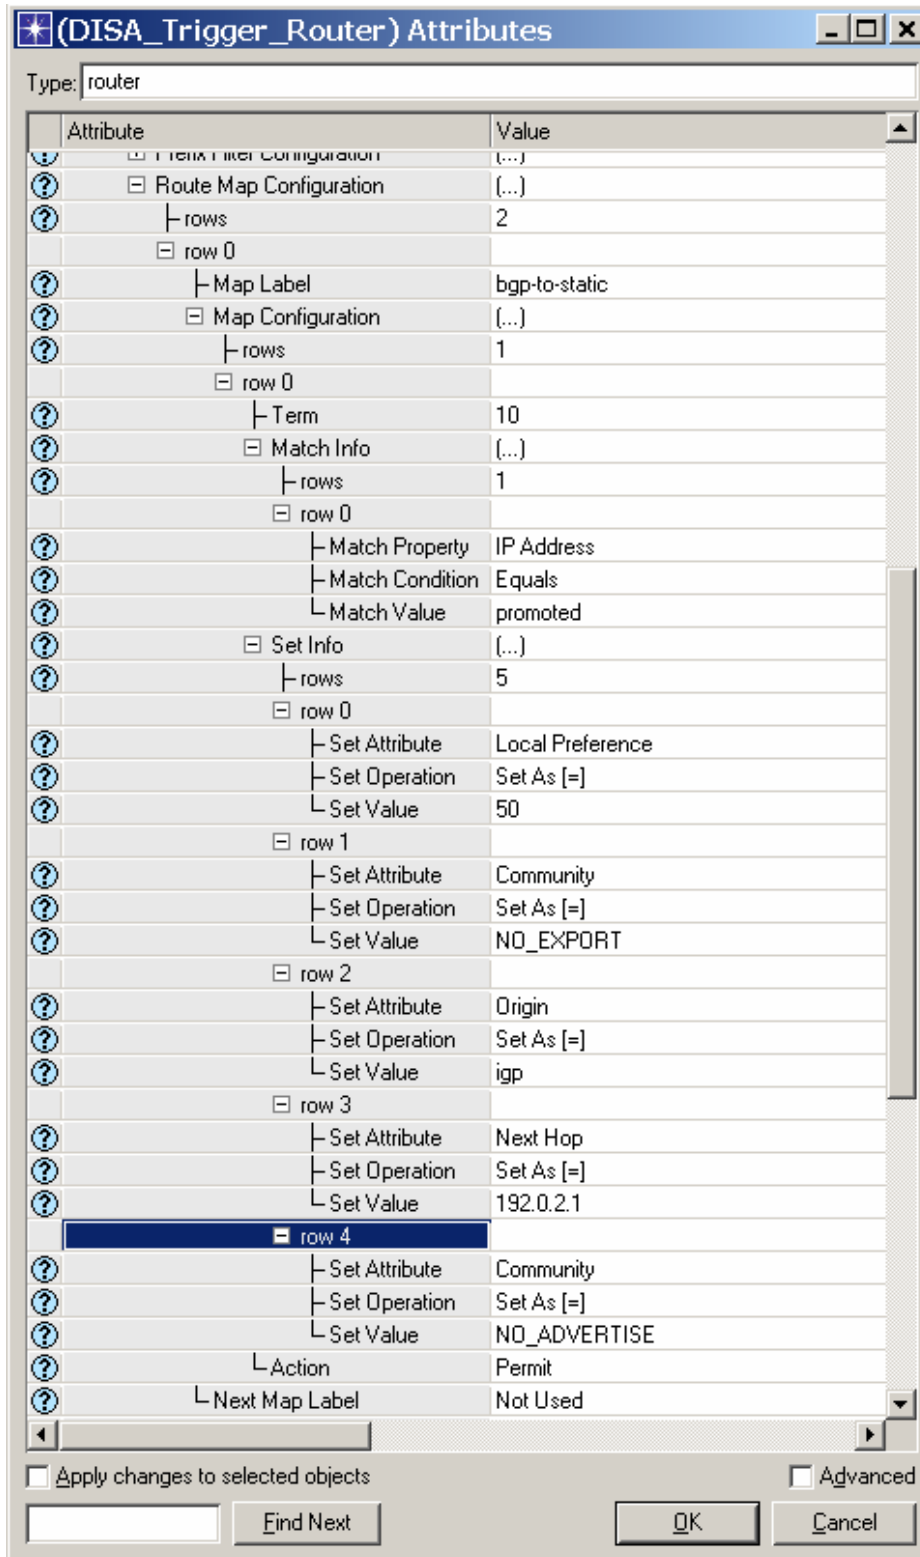


Figure 50 – Trigger-Router Route Map Configuration

(DISA_Trigger_Router.IP Routing Parameters [0].Static Routin... [X]

Destination Address	Subnet Mask	Next Hop	Administrative Wei...	VRF Name
0.0.0.0	0.0.0.0	1.1.1.2	1	None
10.10.5.128	255.255.255.255	Null0	1	None

2 Rows [Delete] [Insert] [Duplicate] [Move Up] [Move Down] [Details] [Promote] [OK] [Cancel]

Figure 51 – Trigger-Router Static Routing Table

Bibliography

- [BEN00] Bennett, T., “Distributed Denial of Service Attacks”, http://www.linuxsecurity.com/resource_files/intrusion_detection/ddos-faq.html, February 2000.
- [CIS96] “OSPF Design Guide”, <http://www.cisco.com/warp/public/104/2.html>, 1996.
- [CIS03] “Cisco Internetworking Technology Handbook, Chapter 39 – Border Gateway Protocol”, http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/bgp.pdf, 1992-2003.
- [CIS04] “Black-Hole Filtering Minimizes Impact of Server Attacks”, http://www.cisco.com/warp/public/779/servpro/promotions/bbip/pdfs/bbip_v5.06.pdf, March 2004.
- [DAT04] Data Connection, “RIP: Routing Information Protocol”, <http://www.dataconnection.com/iprouting/ripprotocol.htm>, 1998-2004
- [GRE02] Greene, B., “Remote Triggered Black Hole Filtering-02”, <ftp://ftp-eng.cisco.com/cons/isp/essentials/>, August 2002.
- [KAL00] Kalyanaraman, S., “Exterior Gateway Protocols: EGP, BGP-4, CIDR”, http://www.ecse.rpi.edu/Homepages/shivkuma/teaching/sp2000/i12_egp/, March 2000.
- [MMR00] Mirkovic, J., J. Martin, and P. Reiher, “A Taxonomy of DDoS Attack and DDoS Defense Mechanisms”, http://lasr.cs.ucla.edu/ddos/ucla_tech_report_020018.pdf, 2000.
- [MOR04] Battle, T., D. McPherson, and C. Morrow, “Customer-Triggered Real-Time Blackholes”, <http://www.nanog.org/mtg-0402/pdf/morrow.pdf>, February 2004.
- [RAJ02] Rajnovic, D., “Black Hole Routers”, <http://www.terena.nl/tech/task-forces/tf-csirt/meeting7/rajanovic-black-hole-routers.pdf>, September 2002.
- [RIV04] Riverhead Networks Whitepaper, “Defeating DDoS Attacks”, http://angell.com/portfolio/Riverhead_WP.pdf, 2004.
- [SAS99] “Dijkstra Algorithm”, http://www.cs.usask.ca/resources/tutorials/csconcepts/1999_8/tutorial/advanced/dijkstra/dijk_descrip.html, 1999
- [SEC03] “Autonomic Systems – Combating DDoS Attacks”, <http://www.securesynergy.com/library/articles/037-2003.php>, March 2003.
- [UCD03] “EEC 189Q: Introduction to Communication Networks”, http://www.ece.ucdavis.edu/~chuah/classes/eec189q/lectures/L7_globalinternet.pdf, October 2003.

[XHN05] “Understanding BGP Session Robustness in Bandwidth Saturation Regime”,
http://cairo.cs.uiuc.edu/~lixiao/application/CV_lixiao.pdf, 2005

[ZAR03] Zaroo, P., “A Survey of DDoS attacks and some DDoS defense mechanisms”, http://www.cs.purdue.edu/homes/zaroo/papers/my_papers/ddos_paper.pdf, June 2003.

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 074-0188*

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 21-03-2005		2. REPORT TYPE Master's Thesis		3. DATES COVERED (From - To) Aug 2003 - Mar 2005	
4. TITLE AND SUBTITLE Analysis of Effects of BGP Black Hole Routing on a Network like the NIPRNET				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Kleffman, Michael D., Captain, USAF				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way, Building 640 WPAFB OH 45433-7765				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GIA/ENG/05-01	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Mr. Neal Ziring NSA/I33 Fort George G. Meade, MD 20755-6000 Phone: 410-854-5762				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The Department of Defense (DoD) relies heavily on the Non-secure Internet Protocol Router Network (NIPRNET) to exchange information freely between departments, services, bases, posts, and ships. The NIPRNET is vulnerable to various attacks, to include physical and cyber attacks. One of the most frequently used cyber attacks by criminally motivated hackers is a Distributed Denial of Service (DDoS) attack. DDoS attacks can be used to exhaust network bandwidth and router processing capabilities, and as a leveraging tool for extortion. Border Gateway Protocol (BGP) black hole routing is a responsive defensive network technique for mitigating DDoS attacks. BGP black hole routing directs traffic destined to an Internet address under attack to a null address, essentially stopping the DDoS attack by dropping all traffic to the targeted system. This research examines the ability of BGP black hole routing to effectively defend a network like the NIPRNET from a DDoS attack, as well as examining two different techniques for triggering BGP black hole routing during a DDoS attack. This thesis presents experiments with three different DDoS attack scenarios to determine the effectiveness of BGP black hole routing. Remote-triggered black hole routing is then compared against customer-triggered black hole routing to examine how well each technique reacts under a DDoS attack. The results from this study show BGP black hole routing to be highly successful. It also shows that remote-triggered black hole routing is much more effective than customer-triggered.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 107	19a. NAME OF RESPONSIBLE PERSON Graham, Robert P. Jr, USAF, AFIT/ENG
REPORT U	ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code) (937) 255-3636, ext 4715; e-mail: Robert.graham@afit.edu

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39-18