

NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

LONGITUDINAL STUDY OF LARGE-SCALE TRACEROUTE RESULTS

by

Dillon Glasser

March 2018

Thesis Advisor: Second Reader: Robert Beverly Justin P. Rohrer

Approved for public release. Distribution is unlimited.

REPORT DO	CUMENTATION PA	GE	Form Approved O	MB No. 0704–0188
Public reporting burden for this collection searching existing data sources, gathering regarding this burden estimate or any or headquarters Services, Directorate for Info to the Office of Management and Budge	on of information is estimated to avera and maintaining the data needed, and ther aspect of this collection of inform ormation Operations and Reports, 1215 t, Paperwork Reduction Project (0704-	ge 1 hour per respo completing and review nation, including sug Jefferson Davis Hig 0188) Washington D	nse, including the time t ving the collection of info gestions for reducing thi hway, Suite 1204, Arlingt C 20503.	for reviewing instruction, rmation. Send comments is burden to Washington ion, VA 22202–4302, and
1. AGENCY USE ONLY (Leave Bland	k) 2. REPORT DATE March 2018	3. REPORT TYP Master's Thesis	E AND DATES COVER 01-05-2016 to 03	RED 3-30-2018
4. TITLE AND SUBTITLE	5. FUNDING NUMB	ERS		
LONGITUDINAL STUDY OF LAR 6. AUTHOR(S) Dillon Glasser	RGE-SCALE TRACEROUTE RES	ULTS		
7. PERFORMING ORGANIZATION N	8. PERFORMING OR	GANIZATION REPORT		
Naval Postgraduate School Monterey, CA 93943	NUMBER			
9. SPONSORING / MONITORING A N/A	ES)	10. SPONSORING / AGENCY REPOR	MONITORING T NUMBER	
11. SUPPLEMENTARY NOTES			I	
The views expressed in this docume Defense or the U.S. Government. IR	ent are those of the author and do B Protocol Number: N/A.	not reflect the offic	cial policy or position	of the Department of
12a. DISTRIBUTION / AVAILABILIT	Y STATEMENT		12b. DISTRIBUTION	CODE
Approved for public release. Distrib	ution is unlimited.			
13. ABSTRACT (maximum 200 words))			
Traceroute is a popular active probin of IP networks. However, traceroute is of traceroute artifacts to find anomal the IPv4 Routed /24 Topology Dat traceroute results, including anomal find attribution to the cause of each possible explanations for unexplaine	ag technique used by researchers, o is susceptible to making inaccurate lies that may be indicative of netwo aset from the Center for Applied ous and unexpected artifacts. We a Finally, we provide a longitudinal d artifacts.	perators, and adver inferences. We per ork errors, misconf Internet Data Ana nalyze the distribu analysis of multi-	rsaries to map the struc- form a large-scale lon figuration, or active de ulysis (CAIDA), we p tion of the observed a protocol label switchi	cture and connectivity gitudinal investigation sception efforts. Using rovide a taxonomy of rtifacts and attempt to ing in order to explore
14. SUBJECT TERMS				15. NUMBER OF
Internet Measurement, Multi-protoc	ol Label Switching, traceroute			16. PRICE CODE
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY OF ABSTI	CLASSIFICATION	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	U	nclassified	UU
NSN 7540-01-280-5500			Standa Prescril	rd Form 298 (Rev. 2–89) bed by ANSI Std. 239–18

Approved for public release. Distribution is unlimited.

LONGITUDINAL STUDY OF LARGE-SCALE TRACEROUTE RESULTS

Dillon Glasser Civilian, Department of the Navy B.S., Stockton University, 2013

Submitted in partial fulfillment of the requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL March 2018

Approved by:

Robert Beverly Thesis Advisor

Justin P. Rohrer Second Reader

Peter J. Denning Chair, Department of Computer Science

ABSTRACT

Traceroute is a popular active probing technique used by researchers, operators, and adversaries to map the structure and connectivity of IP networks. However, traceroute is susceptible to making inaccurate inferences. We perform a large-scale longitudinal investigation of traceroute artifacts to find anomalies that may be indicative of network errors, misconfiguration, or active deception efforts. Using the IPv4 Routed /24 Topology Dataset from the Center for Applied Internet Data Analysis (CAIDA), we provide a taxonomy of traceroute results, including anomalous and unexpected artifacts. We analyze the distribution of the observed artifacts and attempt to find attribution to the cause of each. Finally, we provide a longitudinal analysis of multi-protocol label switching in order to explore possible explanations for unexplained artifacts.

Table of Contents

1]	Introduction	1
1.1	Importance of Network Topology	1
1.2	Scope	4
1.3	Summary of Findings	4
1.4	Thesis Structure.	5
2 1	Background	7
2.1		7
2.2	Historical Traceroute Data.	8
2.3	Large-Scale Active Measurements	9
2.4	Data	10
2.5	Artifacts	13
2.6	Multi-protocol Label Switching	15
3 I	Methodology	19
3.1	Types of Traceroute Artifacts	20
3.2	Data Processing and Classification	25
3.3	Characterization of Discovered Anomalies	27
4 A	Anomalous Traceroute Results	31
4.1	Macro Statistics on Anomalous Traceroutes	32
4.2	DNS Findings of Detected Suspicious Traceroutes	34
4.3	Traceroutes with Both Anomalous Characteristics	43
5 I	Longitudinal Analysis of multi-protocol label switching (MPLS) Results	49
5.1	MPLS Prevalence Across all VPs	49
5.2	Individual Analysis of MPLS Prevalence by VPs	52
5.3	Observed MPLS Change from LEVEL3	69
5.4	MPLS Correlation with Artifacts	70

6	Conclusion and Future Work	71
6.1	Primary Contributions and Takeaways	72
6.2	Future Work	72
Ар	pendix: Vantage Point Boxplots	77
Lis	t of References	85
Ini	tial Distribution List	91

List of Figures

Figure 2.1	MPLS header location. Source: [36]	15
Figure 2.2	MPLS header contents. Source: [37]	16
Figure 3.1	Workflow diagram	20
Figure 4.1	Star Wars Traceroute discovered	36
Figure 4.2	Star Wars Traceroute search for "episode.iv" within the historical domain name system (DNS) in 2017. The command used: find . -name "*.gz" xargs -n 1 -P 10 zgrep -H "episode.iv"	38
Figure 4.3	Star Wars Traceroute search for "may.the.force.be.with.you" within the historical DNS in 2017. The command used: findname "*.gz" xargs -n 1 -P 10 zgrep -H "may.the.force.be.with.you"	39
Figure 4.4	Both IP addresses associated with "frontend.episode-iv.de" and "ridcully.episode-iv.de" point to this website as of 2018-03-11	39
Figure 4.5	Traceroute from 2018-3-29 using the command: traceroute -m 50 bad.horse	40
Figure 4.6	Bad.horse instances found in the suspicious traceroute table's DNS column	41
Figure 4.7	One instance of Bad.horse Traceroute with corresponding DNS names.	42
Figure 4.8	Partial example of one instance of the Comstar anomalous traceroutes	45
Figure 4.9	Demonstration of multiple responses to probe 30 and 31 for another instance of the Comstar anomalous traceroutes	46
Figure 4.10	Reverse order of responses to probes for another instance	46
Figure 4.11	Venn diagram of different traceroute artifact behaviors	48

Figure 5.1	Percentage of traces that traverse a MPLS tunnel across all Team 1 vantage points	50
Figure 5.2	Dublin, Ireland	53
Figure 5.3	Percentage of traceroutes that traverse a MPLS tunnel for the VP in New Orleans, Louisiana, United States	56
Figure 5.4	Downward shift by percentage of MPLS by cycle for the New Orleans, Louisiana (msy-us) VP	60
Figure 5.5	Sao Paulo, Brazil	61
Figure 5.6	Pretoria, South Africa	63
Figure 5.7	Downward shift by percentage of MPLS by cycle for the Pretoria, South Africa (pry-za) VP	64
Figure 5.8	Singapore	65
Figure 5.9	Hamilton, New Zealand	67
Figure A.1	Ames, Iowa, United States	77
Figure A.2	Barcelona, Spain	78
Figure A.3	Stockholm, Sweden	78
Figure A.4	Cheongju, South Korea	79
Figure A.5	Dakar, Senegal	79
Figure A.6	Geneva, Switzerland	80
Figure A.7	Helsinki, Finland	80
Figure A.8	Leipzig, Germany	81
Figure A.9	Monterey, California, United States	81
Figure A.10	Narita, Japan	82
Figure A.11	Oakland, California, United States	82
Figure A.12	San Diego, California, United States	83

Figure A.13	Hong Kong	83
Figure A.14	Sydney, Australia	84
Figure A.15	Toronto, Canada	84

List of Tables

Table 4.1	Summary table of suspicious traceroutes that are closely examined. Seq. IPID and Same /24 columns contain the number of hops dis- playing the observed anomalous characteristics, where a "*" next to the number means it varies between instances observed. Protocols refer to receptiveness to different probing techniques. Breadth is the range of destinations where an artifact will manifest if probed. Multi- VPs indicate whether the traceroute is observed from multiple VP within our dataset. The last column, Occ., represents the number of occurrences we observed from 2013 through 2017.	32
Table 4.2	Top ASN percentages for detected suspicious traceroutes by hop .	33
Table 4.3	Top autonomous system number (ASN) percentages for detected Sus- picious Traceroutes. The "Suspicious Count" column represents the amount of unique traceroute-ASN pairs.	34
Table 4.4	Count of the 89.113.20.0/24 anomalous traces	44
Table 5.1	Median percentages, by year, for all VP traceroutes traversing an MPLS tunnel	50
Table 5.2	msy-us is responsible for a large number of cycles which had total or near 100% MPLS tunnels per trace in 2013	51
Table 5.3	msy-us had many cycles in 2014 that had 100% of traceroutes traverse an MPLS tunnel	52
Table 5.4	Dublin, Ireland, VP traceroutes traversing an MPLS tunnel percent- ages by year	54
Table 5.5	Top 10 ASN appearances for Dublin VP in 2014	54
Table 5.6	Top 10 ASN appearances for Dublin VP in 2015	55
Table 5.7	Dublin, Ireland, VP notable ASN changes	55
Table 5.8	New Orleans VP traceroutes traversing an MPLS tunnel percentages by year	56

Table 5.9	Top 10 ASN appearances for New Orleans vantage point in 2013 .	57
Table 5.10	Top 10 ASN appearances for New Orleans vantage point in 2014 $$.	58
Table 5.11	Top 10 ASN appearances for New Orleans vantage point in 2015 .	58
Table 5.12	New Orleans VP notable ASN changes	59
Table 5.13	Sao Paulo VP traceroutes traversing a MPLS tunnel percentages by year	61
Table 5.14	Sao Paulo VP notable ASN changes	62
Table 5.15	Pretoria VP traceroutes traversing a MPLS tunnel percentages by year	63
Table 5.16	Pretoria, South Africa, VP notable ASN changes	64
Table 5.17	Singapore VP traceroutes traversing a MPLS tunnel percentages by year	66
Table 5.18	Singapore VP notable ASN changes	66
Table 5.19	Hamilton, New Zealand, VP traceroutes traversing a MPLS tunnel percentages by year	68
Table 5.20	Hamilton, New Zealand, VP notable ASN changes	68
Table 5.21	LEVEL3 demonstrated no noticeable change in frequency of appear- ance in traceroutes, but significant change in MPLS	69
Table 5.22	All VP notable MPLS ASN Suspicious traceroute counts	70

List of Acronyms and Abbreviations

AWS	Amazon Web Services
Ark	Archipelago Measurement Infrastructure
AS	autonomous system
ASN	autonomous system number
BGP	border gateway protocol
CAIDA	The Center for Applied Internet Data Analysis
CTD	CAIDA Traceroute Data
CDN	content distribution network
DISA	Defense Information Systems Agency
DNS	domain name system
ICMP	internet control message protocol
ΙΟΤ	Internet-of-Things
IP	internet protocol
ISP	Internet service provider
IPID	IPv4 identification field
MPLS	multi-protocol label switching
msy-us	New Orleans, Louisiana
NPS	Naval Postgraduate School
PPS	packets per second

pry-za	Pretoria, South Africa
--------	------------------------

- **RFC** request for comment
- **RTT** round-trip time
- **TPB** The Pirate Bay
- TTL time-to-live
- TCP transmission control protocol
- **UDP** user datagram protocol
- **VP** vantage point
- **VRF** virtual routing and forwarding

Acknowledgments

I would like to thank my advisor, Robert Beverly, for his his encouraging attitude, patience, and expertise. I have learned an incredible amount from him in a short period of time. I also would like to thank Justin Rohrer for his insight, recommendations, and patience. In addition, I would like to thank CAIDA for the wonderful resources and staff. Young Hyun and Matthew Luckie provided tools and detailed responses that were invaluable. My gratitude also extends to the SFS program and NPS for giving me the opportunity to work on this thesis. Furthermore, I want to thank Aakash Taneja for his guidance in my undergraduate studies and beyond. Finally, I would never have finished if not for the moral support of Leanne, Ruby, and Winston.

CHAPTER 1: Introduction

The Internet is a dynamic, trillion-dollar ecosystem prevalent in every aspect of society. However, this ecosystem does not continue to exist and function without significant effort. To keep the network running, network operators rely on a variety of network management systems and techniques, as well as a suite of tools for debugging the network when problems arise. An important aspect of network management, particularly with respect to routing and defending against attacks, is accurately understanding the logical topology of the network. An accurate logical topology of the network facilitates an understanding of the interconnections between network devices.

Researchers, network operators, and policymakers rely on accurate data in the form of network topologies. Unfortunately, the diagnostic tools available today to determine the network topology are limited due to the lack of underlying support from the network, as well as information hiding. An Internet service provider (ISP) may practice information hiding for competitive or security reasons, especially across administrative boundaries. The state-of-the-art tools currently in use for inferring network topologies are largely based on traceroute techniques.

Therefore, it should be of concern to the community that previous research has demonstrated techniques that can deceive and distort network topologies formed through the use of traceroute. This is due to the fact that there is no integrity in traceroute. Furthermore, misconfigurations and mis-implementation of networking devices can cause anomalous returned results for a given traceroute. These results can cause incorrect inferences of the true network topology. The prevalence and scale of these anomalous traceroute causing networking devices requires investigation.

1.1 Importance of Network Topology

Network Topological maps are interesting and insightful diagrams. However, they may be under-appreciated as powerful and useful tools. This section will focus on a few of the plethora of situations where Network Topology is useful, with traceroute as the regularly sought tool.

1.1.1 Diagnostics

If a destination is unreachable and there is a need to find where in the path a problem occurs, a network operator may utilize traceroute. It can be important to determine which parts of a service are down and in what geographic regions, a recent example being the Amazon Web Services (AWS) outages [1]. Additionally, a problem does not need to arise for traceroute to be used as a diagnostic tool. Many organizations simply want to know where their traffic is flowing on the way to a particular destination. A specific scenario could be an ISP, or some other large organization, may have agreements with many ISPs who charge varying amounts for traffic over their domain. A large entity concerned about cost of their internet traffic could use traceroute to identify which providers are receiving their traffic. This large entity could then change their traffic to flow over a provider that charges less on the way toward a particular destination. While this is just one hypothetical example, there is little doubt that the organization would benefit from the use of a network topological map in this scenario and many like it.

1.1.2 Network Optimization

To keep pace with ever-increasing demand, network operators and providers must optimize their networks. A content distribution network (CDN) is one such entity that commonly considers network topologies in order to deliver content as fast and reliably as possible. An example of a CDN is Akamai, who distributes content for thousands of enterprises such as Adobe, Airbnb, Defense Information Systems Agency (DISA), IBM, and many others [2] [3]. Akamai must manage over 60,000 servers on 1,000 networks in 70 countries in order to facilitate their CDN obligations [3]. Something that can make this challenging is that a CDN must deliver content in a timely manner over an infrastructure that it does not have control of, using protocols that are not designed with speed in mind. Akamai even acknowledges that border gateway protocol (BGP) is inefficient, which is a concern to their mission as a CDN. "[BGP] was never designed for performance: BGP bases its route calculations primarily on AS hop count, knowing nothing about the topologies, latencies, or real-time congestion of the underlying networks" [3]. This quote emphasizes that a large CDN such as Akamai gives much attention to network optimization, a large component of

which network topologies play a role.

1.1.3 Protecting Critical Infrastructure

The Internet extends across the entire globe and through many countries. However, there is still a finite amount of physical connections that connect one entity to another. Even the United States is connected to other countries through a finite amount of network links. For a smaller entity such as a CDN, ISP, or some other provider the links could conceivably be at risk of severing from a motivated adversary, accident, or natural disaster. From a defenders point of view, it would be important to identify those links, enumerate them, and implement a contingency plan in order to mitigate risk.

1.1.4 Cyber Mission Planning

In many ways, cyber operations parallel physical world scenarios. As with any combat or rescue mission, it would be unwise to proceed without attempting to discern details about a territory or terrain. Likewise, a crucial component of any cyber mission is the intelligence-gathering phase, which includes gaining knowledge of the adversary's network. As one researcher states, "intelligence gained from an organization's network topology could allow malicious actors to prioritize their mission objectives and identify especially weak or critical points in the network" [4].

An example of researchers utilizing traceroute would be the team behind Argus, a system for detecting IP prefix hijackings [5]. Shi et al. were able to demonstrate that a previously theoretical cyber attack, BGP prefix hijacking, did in fact exist in the Internet [5]. Also, Shi et al. did not just use normal traceroute data, but rather CAIDA Traceroute Data (CTD), which is what will be used in this research. Shi et al. were able to find 40K routing anomalies within a year period, from which 220 were determined to be prefix hijackings.

In general, "network and security researchers rely on topological maps of the logical Internet to address problems ranging from critical infrastructure protection to policy" [6]. An accurate network topology is often crucial for all of the situations detailed above, which traceroute can assist in creating. Researchers and policymakers would struggle to analyze a trillion-dollar ecosystem if not for traceroute. It is difficult to study the internet because of the "radically distributed ownership of its constituent parts, and an operational climate that generally disincents sharing data with researchers" [7]. An effective method of obtaining this valuable data is through utilizing traceroute.

1.2 Scope

This thesis will focus on traceroute artifacts observed in the CAIDA Archipelago routed IPv4 /24 topology dataset. We will analyze CTD from 2013 through 2017 for the IPv4 Internet. Our analysis will look for traceroutes that appear suspicious based on criteria detailed in Section 3.1. We will also be looking for traceroutes that traverse a multi-protocol label switching (MPLS) tunnel. We then will analyze the suspicious traceroutes in detail in an attempt to discover their root cause. We will also study the trends in the utilization of MPLS by various autonomous system number (ASN) to explore any correlation with the suspicious traceroutes.

Other datasets, such as iPlane, will not be studied. The artifacts observed are anticipated to come from many different provider's networks. Validation of results will be attempted by contacting network providers for specific information about their networks. While we will try to find outside validation, it is out of scope to completely validate all results as we simply do not have access to all networks and are unable to compel network providers to work with us. However, authors have had luck in this research area in the past in both validating their results and identifying false positives [8]. We may look at other measurements in order to confirm suspicions, but we will not be looking in depth at other historical measurements, such as ping data. This research is focused on IPv4, therefore IPv6 is out of scope.

1.3 Summary of Findings

- Average detected MPLS tunnels traversed over all vantage point (VP)s hovers between 33% and 46%, while varying widely between VPs.
- Level 3 identified as a large autonomous system (AS) who stopped utilizing or advertising MPLS in September of 2014 across multiple VPs.
- MPLS as a networking technology in large ISPs is still being used, but there are alternatives that may explain the reduction in overall percentages. Additionally, there may be security concerns incentivizing large ISPs to disregard use of RFC 4950 [9] and instead not advertise that they are even using the technology.

- PROXAD (Free ISP) responsible for the most suspicious traceroute results.
- Discovered a "wild," previously undiscovered, Star Wars traceroute gag within our suspicious traceroute.

1.4 Thesis Structure

- Chapter 1 is the introduction, where we discuss the importance of an accurate network topology, the scope of our research and the summary of our findings.
- Chapter 2 is the background, where we review the basic concepts behind the traceroute tool, the CTD, how the CTD is collected, the traceroute artifacts found by past researchers, and an overview of the MPLS networking technology.
- Chapter 3 focuses on the methodology where we discuss how we will handle the CTD in our attempt to identify the different types of traceroute artifacts. We also detail what defines the different traceroute artifacts. We then discuss how we will go about studying any traceroute artifacts, specifically by looking at their domain name system (DNS) names in the historical DNS dataset and by looking at their associated ASN. We also discuss some other information we may hope to gather such as the geolocation of the artifact, the device it originates from, and the receptiveness to which it responds to various probing techniques. We also discuss in this chapter how will use common traceroute gag strings to search through the historical DNS dataset. Additionally, we discuss our approach of creating a postgresSQL database to facilitate rapid analysis. Finally, we discuss how we will focus on percentages of traceroutes traversing MPLS tunnels on a per VP basis.
- Chapter 4 discusses the results and analysis of our research pertaining to anomalous traceroutes.
- Chapter 5 discusses the results and analysis of our research pertaining to the trend in utilization of MPLS by ASs and their possible influence on anomalous traceroutes.
- Chapter 6 finishes with our conclusion and suggestions for future work.

CHAPTER 2: Background

This chapter introduces the network tools traceroute and Scamper, and their use in creating historical traceroute data. Also covered is the past work concerning traceroute artifacts and the prevalence of the networking technology MPLS.

2.1 Traceroute

Traceroute is a widely used networking diagnostic and measurement tool developed by Van Jacobson [10]. The general concept behind traceroute is to display the forward IP router interface path and round-trip delays of packets while traveling across a network. Traceroute operates by sending packets with increasing time-to-live (TTL) values. The router that receives the packet will decrement the TTL by one and pass it along unless the TTL value is zero. If the TTL value is zero, the router will send an ICMP Time Exceeded message back to the source [11]. The source will then increment the TTL and proceed until the destination replies (e.g., with an ICMP message or TCP FIN), until the sent packets are lost a set number of times in a row (referred to as the gap limit) or a router at a given TTL does not respond with the TTL exceeded messages, depending on implementation. If we do not receive a reply from the destination is unresponsive. The source will continue to compile its list of received router interface IP addresses.

It is important to note that information about the rest of the interfaces of the routers along the path is not exposed, i.e., a traceroute exposes only a single interface of each router along the path. Even though we may have a topology that includes these routers, we will not have all of the available interfaces enumerated and the complete topology could remain unknown. In addition, these many interfaces could lead naïve versions of traceroute to fall prey to false inferences due to load-balancing.

Traceroute can use any transport-layer protocol (TCP, UDP, ICMP) for the TTL-limited probes. UDP can often work well as a probing method due to firewalls having their own default setting of permitting UDP probes [12]. Similarly, TCP SYNs can use a well-known

port, such as port 80, to pass firewalls. ICMP echo requests are advantageous due to being relatively lightweight but are often blocked. Finally, naïve versions of traceroute are susceptible to false inferences, such as false and missing links, due to the prevalence of load-balancing on the Internet. internet protocol (IP) was not instrumented for topology discovery. Traceroute has these limitations because it is effectively a hack that relies on internet control message protocol (ICMP) error messages intended for other purposes. Paris traceroute is an implementation that makes alterations to ICMP and UDP, which are called ICMP-Paris and UDP-Paris, to ensure that all of the probe packets toward a given destination follow the same load-balanced path [13].

2.2 Historical Traceroute Data

Historical network data of all kinds is extremely valuable to researchers. Using historical network data, we observe trends or more deeply investigate events that otherwise would have been missed if not for the ability to compare historic and current topologies, e.g., the development of new connections between networks or changes in networks. For example, Giotsas et al. used historical BGP streams in order to study outages [14]. Their findings showed that they were more successful at detecting outages than the traditional organizations, who only detected 24% of the total outages. In the same vein, it is important and useful to look at historical traceroute data. Luckie et al. utilized macroscopic traceroute data correlated with global BGP data to discover that over their 2.5 year time period, 40% of studied IPv6 routers experienced outages [15].

We can also perform accurate evaluations of optimizations, implementations of request for comments (RFCs), or infrastructure changes. Often it would be expensive or impossible to compare the before and after, but historical traceroute data gives us the ability. A specific example is Lone et al. using historical traceroute to infer the ability to spoof on a given network [8]. RFC 2827 (updated by RFC 3704) describes using ingress filtering (or source address validation) to prevent users within an AS from spoofing an address in order to possibly perform a denial of service attack [16], [17]. Lone et al. were able to leverage one month of historical traceroute data in order to identify 703 provider ASs that lacked source address validation [8]. Historical traceroute data also allows one to compare the current abnormal result against what a path may typically look like. This combination of probing and looking at historical probes can be a powerful tool. For example, Shi et al. created the

Argus tool by leveraging historical traceroutes from The Center for Applied Internet Data Analysis (CAIDA) combined with live BGP data in their efforts to identify potential prefix hijackings [5].

2.3 Large-Scale Active Measurements

When performing large-scale active topology measurements, researchers face a few challenges. These measurements must be performed, as much as possible, in an unbiased and uniform manner. Utilizing multiple VPs is a powerful method to minimize sampling biases. In addition, the program Scamper enables researchers to use a uniform traceroute tool, designed specifically for large-scale measurements, on any operating system.

2.3.1 Utility of Multiple Vantage Points

When conducting internet measurements from a single geographic location, the results might not be reliably representative of the entire internet. That single geographic location is referred to as a Vantage Point (also referred to as a monitor). A VP for the purposes of this research is a single geographic location from which a researcher is studying the internet, mainly by probing. Conducting the same measurements from another VP could produce drastically different results. Shavitt et al. performed research on the importance of distributing VPs when performing Internet topology measurements [18]. They discovered that "although increasing the number of VPs can help in reducing sampling bias, it still does not guarantee unbiased results" [18]. They went on to explain that other characteristics of the VPs needed to be taken into account, such as which ASs the VPs are located within, to avoid biased measurements. Barford et al. delves into the utility of network tomography, the concept of gaining insight into the conditions and configurations in the core of the internet by using a number of VPs [19]. Barford et al. finds that the marginal utility of multiple VPs declines rapidly after two VP. However, they also noted that their research only pertains to "the most basic topographical information about the network" [19]. In our research, we find that using many VPs across different ASs and geolocations helps avoid sampling bias and provides additional insights unique to a VP, as demonstrated in Section 5.2.

Therefore, using multiple VPs can help avoid bias from single locations, but is not sufficient by itself. Another example of such research includes Luckie et al. who studied deployed

TCP implementations and their resilience to blind attacks [20]. Middleboxes could have potentially manipulated their test packets before they reached their destination. Luckie et al. used three vantage points to minimize the risk of an undetected middlebox in the hosting network from impacting their measurements. This demonstrates another utility of using multiple VPs to avoid biased results.

2.3.2 Scamper

For large-scale active measurements, using traceroute's implementation on each operating system is not ideal. In order to gather the data, specialized Internet-wide traceroute tools are necessary. A heavily used program for large-scale active measurement is Scamper [21]. Scamper was created in order to allow researchers to conduct large-scale measurements and archive the data more efficiently. This highly useful tool also enables researchers to solicit volunteers from around the world, no matter their operating system. The problem with traceroute on various operating systems is that they are all implemented differently, making it less than ideal for uniform measurements. In addition, they lack new features like the Paris traceroute implementations [21]. Scamper also exposes ICMP extensions, thereby letting us observe MPLS tunnels, a feature not found in all traceroute variations.

Another useful feature is Scampers ability to set a fixed packets per second (PPS) rate. Setting PPS allows the monitoring node to generate a predictable, and limited, rate of traffic based on the sites configuration policy. Scamper provides researchers with a tool designed for research on any platform, with newer features and better output for archiving and analysis purposes. The two output formats for Scamper are standard ASCII text and a binary file format named warts. There also exists a warts to JSON conversion utility.

2.4 Data

Publicly available Internet measurement datasets, including traceroute data, are available from CAIDA [22]. However, there exists other platforms besides CAIDA that provide traceroute and internet measurement data.

2.4.1 Production Mapping Systems

One example of a large-scale mapping system that assists in constructing a topology of the internet is "iPlane, a scalable service providing accurate predictions of Internet path performance for emerging overlay services" [23]. iPlane utilizes PlanetLab servers at over 300 sites around the world and utilizes traceroute as its primary tool for determining Internet topology. iPlane leverages BGP snapshots collected by RouteViews, which provides a list of all globally routable prefixes [24]. iPlane looks within the BGP snapshot for a ".1 address that responds to either ICMP or UDP probes" [23]. iPlane then performs periodic probing of the targets. Note that the frequency of the topology mapping using traceroute is about once a day [23]. Traceroute only produces the network interfaces from the path to the destination. However, a router along the way can have many network interfaces. iPlane identifies interfaces that belong to the same router by looking for possible aliases and narrowing them down to being likely aliases. iPlane also looks for DNS names assigned to interfaces in pursuit of identifying their geographic location.

Another example of an Internet measurement infrastructure is the DIMES platform [25]. Dimes was launched in 2004 and relies on volunteers to install a software agent which helps to measure the structure and evolution of the internet. The strength of the DIMES platform is in the many thousands of installed agents, in many countries, and many ASs. DIMES is an attempt at creating "mid-level granularity maps" [25], rather than AS level topologies like RouteViews [24] or router level maps such as The IPv4 Routed /24 Topology Dataset from CAIDA [26].

Finally, the RIPE Atlas platform is another measurement platform, consisting of probes and anchors distributed around the globe. RIPE Atlas was designed to measure connectivity and reachability of the Internet [27]. As of 2015, there were 12,800 probes and 109 anchors [27]. The RIPE Atlas platform uses the various "probes against anchors to measure region-based connectivity and reachability [27]." Researchers and other participants are free to install probes in their personal networks or organizational network, with permission. The data collected by RIPE Atlas is publicly available.

2.4.2 CAIDA Infrastructure and Available Data

CAIDA deploys and maintains an internet measurement infrastructure, the Archipelago Measurement Infrastructure (Ark) [26]. Ark utilizes Scamper across an ever-increasing number of VPs, with newly deployed VPs generally being deployed on a Raspberry Pi computer [28]. The traceroute data is collected using Scamper in the warts file format. Ark collects this data by splitting up the probing into three teams. Each team has VPs spread out around the world and each team separately looks at the over 10 million /24 prefixes list. The individual team will split the /24 prefixes amongst the vantage points. If a VP is faster than another, it may be responsible for more work. This is important to note as the changing workload of the vantage points can change the paths traversed by each vantage point for each cycle.

In 2007, there were 20 total VPs, with 12 in team 1. At the time of this research in 2018, there are 208 total VPs, with 71 belonging to team 1 [29]. Ark provides an extensive amount of traceroute data from 2007 to 2016 for routed IPv4/24 prefixes. CAIDA has over 10 million /24 prefixes in their list. Instead of probing a static list of IP addresses, Ark randomly picks an address in each /24 prefix for each cycle. A cycle is the round of probing done by all monitors for a given team, identified by an id number. The cycles are numbered sequentially and the next cycle starts only after the previous one finishes.

Also, since Ark runs one cycle after another one finishes, there are times where a specific date may not be available and may be contained within a cycle. For example, if a researcher were looking for traces specifically on May 4th, the researcher may need to look at the cycle that begins on May 3rd since these cycles can span multiple days. The destinations probed are all the routed /24 networks, 10 million, in the IPv4 address space. A static list of IP addresses is not used. Instead, the destination probed is randomly selected from each of the routed IPv4 /24 prefixes for each cycle.

CAIDA also deploys VPs that probe two addresses within each announced IPv6 prefix that is /48 or shorter. We restrict our analysis to IPv4 in this thesis; future work should perform a similar longitudinal evaluation of IPv6 traceroute data.

The Ark from CAIDA was selected as our dataset to perform our research with due to its extensive historical results, dating back to 2008, and for it being a complete platform, with the inclusion of a historical DNS dataset.

2.5 Artifacts

Intuitively, we might expect traceroute to work in two ways. It could work as we expect it to function, by using the TTL field and ICMP Time Exceeded messages to progressively build a list of returned router interface IP addresses leading toward a destination. Another possibility is that we may expect traceroute to progress partially to the destination, but then be blocked before the packet reaches the target destination. Also, this possibility may be due to routers in the middle of the path not responding, while those at the end do if they can be reached. In real-world traces, only 15 percent of random traces complete [30]. But there is a third category that exists of traceroute artifacts with unexpected behavior. Some notable characteristics are apparent when observing a large quantity of traceroute data. Some traceroute results report the same IP address at multiple hops, indicating, for instance, a routing loop. Also, some traceroute results display unexpected header information such as sequential values in the 16-bit identification field. In addition, it is possible that data packets could take different paths than the traceroute probes. These cases could be due to a number of possibilities such as middle boxes, unintentional misconfigurations, or intentional deception.

The packets in traceroute are not authenticated and do not have any integrity. This is both ideal for those that do not want their networks probed by adversaries and troubling for researchers trying to construct a valid network topology. As one researcher noted,

In the case of traceroute probing, the defender may manipulate the return traffic with deception outcomes such as hiding legitimate nodes, seeding virtual nodes, and masquerading as other nodes. Critical servers and routers are common examples of legitimate nodes that benefit from obscurity. [4]

As another researcher states, "An adversary using the traceroute program may be deceived as to the physical topology of a computer network through the careful manipulation of traffic leaving the network in response to the traceroute probe, thus masking the importance of key nodes along the actual route" [30]. Traceroute is also susceptible to middleboxes and intelligent routers giving deceptive replies. This can assist in hiding the true topology of a network or "shunting" traffic into an enclave of VMs to give a false topology [30].

Such deceptive practices are not simply theoretical, but have been deployed in the wild on the Internet. There are documented cases of deployed decoy hardware, such as the famous Star Wars Traceroute [31]. The Star Wars Traceroute was implemented by deploying and configuring two Cisco 1841 Integrated Services Routers. The creator was able to bounce packets between these two routers using virtual routing and forwarding (VRF), with the two routers likely setup with a static route. Then, the DNS PTR records were set so that the IPs in the traceroute would resolve into the iconic scrolling theme as the reverse DNS lookups were performed. In other words, as the IP addresses were displayed for each hop of the traceroute, the DNS name associated with the IP address was returned in order to show the traceroute gag. The addresses of the routers were all owned by a single entity, as were the corresponding DNS PTR records.

Alternatively to the Star Wars Traceroute approach, a user on Reddit [32] utilized iptables and Perl extensions [4]. Another notable example is when The Pirate Bay caused traceroutes to its web server to appear to lead to a North Korean IP address with a convincing routed path to a naïve observer. One of the methods used to confirm this deception was by comparing the round-trip time (RTT) of a transmission control protocol (TCP) three-way handshake to that of the traceroute [33]. As of this writing, another example of an operational purposely modified traceroute is the famous "bad.horse". Performing a traceroute to a destination of "bad.horse" provides a message that gives the lyrics to a song [34]. While these playful games utilizing traceroute are relatively innocuous, they highlight what is possible should a determined actor attempt to deceive an adversary.

MPLS is an example of a networking technology capable of producing unexpected traceroute results and obscuring the true topology obtained from performing active probing. In the 2012 paper "Revealing MPLS Tunnels Obscured from Traceroute", Donnet et al. concluded that in 2011, at least 30% of traceroute paths traveled through MPLS tunnels [35]. This widespread use of MPLS tunneling obscures router paths, potentially leading to inaccurate or deceptive presentations of Internet topology.

Given that these traceroute artifacts exist, it is important to study the prevalence of these artifacts, when they first appeared, and how long they were active. Also, it would be important to discover what destination and sources are responsible. Without understanding these artifacts and their impact on the network topology, how can any researcher have confidence in the network topologies on which they are basing their work? In addition, many assumptions and past studies have been performed with possible overconfidence in

traceroute's accuracy. Fortunately, we have access to CAIDA historical traceroutes. There are many questions that can be answered by looking at this valuable archive. In addition, we can perform our own probing when we find destinations displaying some of the suspicious behavior.

2.6 Multi-protocol Label Switching

Since we perform extensive research into MPLS behavior and utilization within the CTD in an attempt to explain a subset of the traceroute artifacts, we present here relevant details of MPLS.

In conventional IP forwarding, routing decisions are made at each individual router between a source and destination. Conventional routers use IP address prefixes to determine a packet's Forward Equivalency Class (FEC), which is used to make forwarding decisions. Each router along a path reexamines the packet's prefix to determine a new FEC. MPLS deviates from this architecture by using fixed length labels, inserted into packets between the layer 2 and layer 3 headers by MPLS routers. These labels determine the path packets will take on their subsequent hops until exiting the MPLS tunnel. The MPLS header, which contains the label, also has its own embedded time-to-live (TTL), utilized while inside a Label Switched Path (LSP) [36].



Figure 2.1. MPLS header location. Source: [36].

0 1	2	3 4	5	6	7	8	9	1	1	2	3	4	5	6	7	8	9	20	1	2	3	4	5	6	7	8	9	3	1
						L	al	pel										E	×	,	S			1	TT	Ľ			

- Label Label Value (Unstructured), 20 bits
- Exp Experimental Use, 3 bits; currently used as a Class of Service (CoS) field
- S Bottom of Stack, 1 bit
- TTL Time to Live, 8 bits

Figure 2.2. MPLS header contents. Source: [37].

MPLS routers make the FEC assignment once, when the packet enters the LSP. The FEC is encoded in a label, which is inserted into the IP header and travels along with the packet while in the LSP. MPLS can reduce the time required for forwarding decisions, provide VPN services, and enable the implementation of specific routing paths and policies [36].

In order to classify the different MPLS tunnels, Donnet et al. separated MPLS traffic into four methods of tunneling. The methods are based on the inclusion of TTL-propagate and RFC 4950. In TTL-propagate, when a packet enters a Label-Switched-Path (LSP), the MPLS TTL mirrors the IP TTL of the incoming packet. While traversing MPLS routers in the LSP, the MPLS TTL is decremented at each hop. The IP TTL does not change while in the LSP. When the MPLS TTL expires, the router will send an ICMP time-exceeded reply. The IP packet of the traceroute probe has an additional TTL called the quoted TTL or q-TTL, which represents the difference in the IP-TTL from the MPLS-TTL. The q-TTL increments at each successive MPLS router visited. By examining the ICMP time-expired message, Donnet et al. inferred that reply packets with q-TTL's greater than 1 indicates the presence of MPLS tunneling within a traceroute [35].

The other tool used to identify MPLS tunneling is the observation of whether a Label Switched Router (LSR) enables RFC 4950. If an LSR receives an undeliverable MPLS-encapsulated datagram, it strips off the label stack and typically sends an ICMP message to the sender indicating why the datagram was not delivered to its destination. With RFC 4950 enabled, the LSR sending the TTL response can append MPLS information to the
ICMP message using an ICMP extension [9].

Explicit tunnels implement both TTL-propagate and RFC 4950. With explicit tunnels, traceroutes will reveal each MPLS router traversed and identify the MPLS routers inside the Label Edge Routers (LERs), which start and end MPLS tunnels. Implicit tunnels utilize TTL-propagate, but do not use RFC 4950. This will yield a traceroute return, which identifies the MPLS routers in the tunnel, but does not identify them as MPLS routers. Opaque tunnels do not implement TTL-propagate, but do implement RFC 4950. This may result in some MPLS routers being obscured, but will return indications of the last MPLS router in the tunnel prior to the egress LER. Finally, Invisible tunnels do not enable TTL-propagate or RFC4950.

Donnet et al. used their measurement of explicit tunnels as their primary metric for MPLS presence in traceroutes. They concluded that "at least 30% of traceroutes from most vantage points reveal explicit MPLS tunnels and more than 5% of collected IP interfaces explicitly exhibit MPLS capability" [35]. The data collection and analysis performed by Donnet et al. forms the groundwork for our research regarding MPLS.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 3: Methodology

Our methodology is to analyze the CAIDA IPv4 routed /24 topology dataset [26], which we refer to as CAIDA Traceroute Data. This data comes from the Ark [38] measurement infrastructure. Across monitors and time, we attempt to classify each trace as either "normal" or "anomalous." The CAIDA traceroute data we analyze is from the time period starting in 2013 through 2017. Although Ark has hundreds of monitors (also referred to as VPs), we focus on 21 stable monitors. These 21 monitors were functional throughout the entire time period, whereas some of the non-included monitors may have had long running maintenance or were added during our selected time period. These monitors created 81,512 warts files, with the average traceroute count per file being 145,659. There were 11,872,948,709 total traceroutes in the data we analyzed.

We then further break down the anomalous traces into one of three categories based on available artifacts: 1) MPLS; 2) Misconfiguration; or 3) Deception. We then characterize the traces within each of these three categories to better understand their origins. This chapter describes our methodology for performing this classification; Figure 3.1 provides a high-level overview.



3.1 Types of Traceroute Artifacts

In this research, we plan to provide a taxonomy of traceroute results, including a categorization of traceroute artifacts. We also intend to look at what fraction of these artifacts can be attributed to misconfiguration, misimplementation, or potentially malicious behavior. The set of artifacts that we will consider include:

- Sequential IPv4 identification fields (IPIDs) from different routers along the path
- Many IP addresses on the path allocated from the same /24 network prefix
- Incorrect Destination IP addresses or false AS path [33]

We detail each of these artifacts next.

3.1.1 Sequential IPIDs

The IPID field with the IPv4 header enables fragmentation and reassembly. The IPID is required to be unique within the maximum lifetime for all datagrams with a given source and destination address pair [39]. Routers have two general IPID behaviors; they return a random IPID or they return a sequential IPID. In the case of sequential IPIDs, each router has an independent counter and each counter advances independently (depending on the rate of control plane traffic or "velocity") [40]. Therefore, we would expect that the sequence of IPIDs in the ICMP TTL exceeded responses received from each router along a traceroute path to be random. However, if the along a traceroute path the ICMP TTL exceeded IPIDs are sequential, and are responses to different TTLs probes, this suggests that the same counter is being used to respond to probes with different TTLs. This would imply that the same router is responsible for the different TTLs. This is not a normal behavior as packets are not supposed to loop, which provides us with an anomaly. We detect this anomaly by looking at each hop within our CAIDA traceroute data for IPIDs that demonstrate this behavior. We have set a threshold of 5 consecutive IPIDs with a sensitivity of 5 plus or minus the previous hops IPID value. Both numbers are picked arbitrarily to limit our focus to the most obvious examples of this artifact. At 5 consecutive IPIDs, our confidence is

quite high that the artifact has occurred.

Data: A Given Traceroute **Result:** True or False value initialization; threshold and sensitivity are set at 5; while hops remaining do read current; if the difference between IPID and the last IPID is less than the sensitivity then if the difference between probettl and the last probettl is less than 10 then if IPID is not (0 OR less than 40 OR equal to probettl OR equal to the last IP address) then consecutive+=1; end end else consecutive=0; end if consecutive is greater than threshold then return True; end end return False;

Algorithm 1: Finding Sequential IPID

3.1.2 Many IP Address of a Traceroute in the Same /24

In addition to sequential IPIDs we are looking for many IP addresses on a path that are within the same /24 address space. In a typical traceroute, packets follow the forwarding path to traverse efficiently between networks, progressively getting closer to the final destination. Thus, it would be unusual in normal operation to observe responses from many router interfaces within the same /24 network prefix.

For example, in our database one instance has what appears to be a normal traceroute to a destination until it comes upon an IP address of 78.254.248.30, followed by 78.254.248.134, 78.254.248.138, 78.254.248.142, and so on in the same /24 until timing out. Similar to the methodology for detecting consecutive IPIDs in a traceroute, we look at each hop within a traceroute to determine how many IP addresses are within the same /24, with a threshold of 4 determining if we classify it as anomalous. The threshold of 4 was determined to empirically strike a good balance between false positives and interesting results. As we will show later in the results, many of these artifacts are associated with an MPLS tunnel and some are found to be related to traceroute gags.

```
Data: A Given Traceroute
Result: True or False value
initialization;
threshold is set to 4;
while hops remaining do
   read current;
   last difference set to current difference;
   if the current hop IP is equal to the last IP then
       if the difference between the last IP and current hop IP is greater than 0 then
           if the last hop difference is equal to the current hop difference then
               consecutive+=1;
           else
               if the last difference is equal to the current difference subtracted from 256
                then
                   difference is set to difference minus 256;
                   consecutive+=1;
               end
           end
       end
   else
       consecutive=0;
   end
   if consecutive is greater than threshold then
       return True;
   end
end
return False;
```

Algorithm 2: Finding many IP addresses in the same /24

3.1.3 Incorrect destination IP addresses

Incorrect destination IP addresses and false AS paths refer to incidents such as The Pirate Bay (TPB) case study [33], where the destination address being advertised appeared to be in North Korea for the purpose of evading regulations. TPB servers were proven to not be

located in North Korea first by using TCP traceroute. In the typical traceroutes to TPB, user datagram protocol (UDP) or traditional ICMP traceroutes were used. Will, a blogger, used TCP traceroute to demonstrate that the response time indicated that the destination could not possibly be in North Korea [33]. Because of the limitations in how fast packets can travel due to the speed of light, the response times indicated that the destination was limited to somewhere in Europe. In addition to performing a TCP traceroute, by looking at the AS path Will was able to determine that the AS path was spoofed. The AS path appeared to be going through Cambodia and then onwards to North Korea, which appeared very impractical considering the main provider of North Korean internet is AS 4837 (China Unicom). TPB actually had the real route end in Germany [33].

3.1.4 Multiprotocol Label Switching Observations

A large component of this research is dedicated to continuing the work carried out by Donnet et al., who discovered that at least 30% of traceroute paths traveled through MPLS tunnels. As mentioned earlier this widespread use of MPLS tunneling obscures router paths, potentially leading to inaccurate or deception presentations of Internet topology. With this in mind, we will be utilizing our database populated with Ark data to determine the percentage of traces that traverse an MPLS tunnel on a per monitor (VP) basis. The time period we will be examining ranges from 2013 until June of 2017. Further analysis will be done regarding any changes in percentages between years. This analysis will focus heavily on which ASNs are most responsible for increases and decreases in percentage, or if ASNs remain relatively consistent in years that show little or no change.

3.2 Data Processing and Classification

To carry out our analysis, this research will rely heavily upon Scamper and CAIDA's Archipelago dataset. We will also use a Python library written to replicate the functionality of scamper for parsing [41]. The Python library will be used to search for these artifacts within the historical traceroute data. We search for the signatures of each artifact detailed previously. We extract information about the trace for cases where we detect artifacts and set them aside for later analysis. We only analyze team one, although we look at multiple monitors within team one since the different vantage points used may give different results and expose artifacts.

While parsing the CAIDA traceroute data warts files, we deposit results and accompanying traceroute data in a database created in order to query for specific commonalities. More details about the database are given in section 3.2.1. While the database contains the expected information necessary to perform a thorough analysis, the database also points back to the original scamper warts files if further analysis is required beyond what was selected for inclusion in the database. We use the ark dataset from 2013 through June of 2017.

For both the MPLS tunnels and traceroute artifacts, a count is kept of the number of occurrences of each per Ark monitor. There is a counter for MPLS tunnels per VP in a given cycle which can be compared with the total traceroute count of the VP for that cycle to calculate the percentage. This count will be incremented by one for every traceroute that has any MPLS. The traceroute count is incremented for every traceroute from the VP in that cycle. In our database, these VPs in a given cycle are called files. The same applies to our traceroute artifacts, which are the sequential IPID values and the similar IP addresses in the same /24. A count is kept of each occurrence per file which can be compared with the total trace count per file to calculate percentages. There is an independent count for the sequential IPIDs artifact and the many IP addresses in the same /24 artifact, on a per traceroute basis.

Also, for the MPLS tunnels, a count is kept of the total occurrences of explicit, implicit, and opaque tunnels for each VP. This count includes multiple hops and can span different tunnels within a single trace. It cannot be compared with the total trace count to calculate percentages. We also make note of multiple MPLS tunnels within a traceroute with a binary flag which is set if more than one is observed.

3.2.1 MPLS and Suspicious Traceroute Database

In this research, we will be handling large volumes of historical traceroute data. Searching through the entire Ark dataset is time consuming and resource intensive. Instead of sequentially running through the unstructured data with every question, it would be helpful to store the information in an intermediate format. We will take the historical traceroute data from the binary files, called warts files, and place it in a postgresSQL database. We anticipate this task to take some time upfront, however we anticipate it being extremely helpful in our

analysis later.

Once we have our artifacts gathered, we will begin looking for patterns. If we can detect a pattern, such as an AS that displays a lot of these artifacts, we may consider doing some of our own probing on the AS. We will want to look to see which destinations are displaying multiple artifacts, which destinations within a shared network exhibit the same behavior, and if multiple networks exhibit the observed behavior. We can also compare our probes from various vantage points against those from the historical traceroute dataset to see if the artifacts still occur. In addition, we will want to determine when artifacts began appearing, if it ended and when, and if it is still active. We can utilize our historical traceroute data to answer these questions.

3.3 Characterization of Discovered Anomalies

After finding our artifacts, we will use the following techniques to characterize the anomalous traceroutes we discover:

- Historical DNS
- Current DNS
- ASN
- Geolocation
- Additional Probing

3.3.1 Historical DNS information

CAIDA provides a IPv4 Routed /24 DNS Names Dataset, which is invaluable for comparing against relevant addresses found within the CAIDA traceroute data [42]. "DNS names are useful for obtaining additional information about routers and hosts making up the Internet topology. For example, DNS names of routers often encode the link type (backbone vs. access), link capacity, Point of Presence (PoP), and geographic location" [42]. While these historical DNS records are collected soon after the topology traces are collected, it should be noted they are not collected immediately. Also, CAIDA attempts to keep the burden on external DNS nameservers and their own infrastructure low by not performing lookups on previous obtained results dating back 7 days.

We will utilize this historical DNS information in conjunction with our database by matching IP addresses associated with hops that appear anomalous to their historical DNS host names.

3.3.2 Current DNS Availability

If we find an artifact in the historical record, it would be worth looking to see whether the artifact is still active. We will detect if the address relating to the artifact resolves using DNS. If there is an A record for the address of the artifact, we can use a Python script that will use the socket Python module [43] to see if it retrieves a hostname:

Example: socket.gethostbyaddr("205.155.65.20") Returns: www.nps.edu

An IP that does not resolve will return 'Unknown host'.

One problem with this approach is that it can take a long time to do it sequentially on thousands, possibly millions, of IPs. Therefore, we will use the Python interface for an asynchronous DNS resolver library, adns-python [44].

3.3.3 Grepping Historical DNS for Well-known Traceroute Gags

In addition to the Ark dataset, there is an accompanying historical DNS dataset, mentioned earlier in section 3.3.1. By taking the suspicious IP addresses we have identified and searching for their associated DNS names, we plan to manually identify those gags that resemble traceroute gags from previous work. In addition, we plan to search through the DNS dataset by using strings of previously known traceroute gags. Because these traceroute gags are sometimes short lived, we are not guaranteed to find all previously known traceroute gags. However, we expect to find some traceroute gags by searching through the DNS dataset, which will allow us to perform some active probing with the returned IP address.

We do not require any specialized software to search through this Historical DNS dataset. The following example command will suffice for most searches.

• find . -name "*.gz" | xargs -n 1 -P 25 zgrep -H 'Star Wars'

The string can also be replaced with an IP address in the command if desired. If that is the case, it should be noted that a '\s' should be used to make the search more specific.

• find . -name "*.gz" | xargs -n 1 -P 25 zgrep -H '\s72.30.208.85\s'

Without '\s', strings that contain the searched string as a substring would be returned.

3.3.4 Associated AS Number

In order to better understand the artifacts we discover, it is helpful to know which AS the artifact's address belongs to. As one researcher noted, "BGP allows each AS to choose its own administrative policy in selecting routes and propagating reachability information to others [45]." These individually administered policies are often necessary due to ASs having complex relationships between other ASs and providers. Therefore, "such policies imply that AS relationships are an important aspect of Internet structure [45]."

Note that accurately mapping router interface addresses to their correct owners can be problematic; for purpose of this research we assume that the AS that advertises the interface's prefix is the owner.

It can be helpful to know who owns the AS in order to determine if a particular organization is responsible for a consistent artifact or perhaps free of certain artifacts. Also, as was the case with TPB example, looking at the AS along a path can allow a researcher to determine that a path is unlikely, if not impossible [33].

We can determine the AS number of any address that we obtain. We utilize a script that can take any number of IPs and determine the corresponding AS numbers by comparing against a RIB file from routeviews [24]. This script performs longest prefix matching against the route views RIB file, represented as a radix tree, to provide the corresponding AS number. Building the radix tree from the RIB files can take some time. In order to facilitate fast lookups, we preserve the radix tree data structures using the Pickle Python library [46].

Any AS that has a large percentage of the total artifacts found will be identified and examined closely.

3.3.5 Geolocation of the Artifact's IP Address

Once we have identified our set of artifacts, it would be useful to know their geographic location in order to discover any commonalities or patterns. If they are predominately from a certain geographic region, we can narrow our research to that region. One way in which we can do this is by looking at the AS number to make a general statement about where an artifact is occurring.

We can also use a number of services to take an IP address and produce a corresponding location. One service we may use is Maxmind GeoIP to determine the location by country and city [47].

3.3.6 Artifacts Receptiveness to Different Probing Techniques

While all of these artifacts should be detectable in the Ark dataset, some artifacts may require additional probing. Certain probes may be able to get through and display artifacts. ICMP can often be blocked, while UDP and TCP are still able to get through. Further, different probes may elicit different behavior, and different probes may be used to the way in which deception is implemented. In the case of an MPLS router, we may be able to determine that it is there if it is implementing RFC 4950 and ICMP extension is in the header. In addition, different probing techniques can be used to discover the likelihood of an advertised IP address being spoofed by looking at the TTL values [33].

CHAPTER 4: Anomalous Traceroute Results

The primary objective of this research was to examine a large dataset of historic IPv4 traceroute data in order to identify and classify artifacts. Using the methods discussed in Chapter 3, we flagged hundreds of thousands of anomalous traceroutes. Across 5 years of CAIDA traceroute data, 561,178 traceroutes displayed many hops in the same /24 IPv4 address space, while 340,382 exhibited sequential IPIDs. These quantities are far too great to individually analyze each instance. However, since the traceroutes and their respective hops are within our database, we are able to rapidly query for results as well as compare against other datasets such as the historical DNS dataset from CAIDA.

We also wanted to enumerate, and perform a deeper analysis, of traceroutes that display both artifact characteristics. In our selected CAIDA traceroute data over the 5 years, we discovered 73 traceroutes which display both consecutive IPIDs and have many IPv4 addresses in the same /24. We will perform a more in-depth analysis on these 73 traceroutes. The rest of this chapter focuses on finding trends within this large number of artifacts and exploring explanations for said artifacts.

Table 4.1. Summary table of suspicious traceroutes that are closely examined. Seq. IPID and Same /24 columns contain the number of hops displaying the observed anomalous characteristics, where a "*" next to the number means it varies between instances observed. Protocols refer to receptiveness to different probing techniques. Breadth is the range of destinations where an artifact will manifest if probed. Multi-VPs indicate whether the traceroute is observed from multiple VP within our dataset. The last column, Occ., represents the number of occurrences we observed from 2013 through 2017.

Traceroute	Seq.	Same	Protocols	Duration	Breadth	Multi-	Occ.
	IPID	/24				VPs	
StarWars	17	18	ICMP-	2016-	1 IP	Yes	3
			Paris	2017			
Badhorse	0	27	All	2014-	1 IP	Yes	4
				Present			
Comstar	545*	134*	ICMP-	2013-	/24	Yes	3
			Paris	2016			
Befree	17*	17*	ICMP-	2013-	/24	Yes	32
			Paris	2016			

Table 4.1 represents the suspicious traceroutes observed in our CAIDA traceroute data that we closely examined. This table serves as a summary of the findings that will be discussed below, as each traceroute will be discussed in detail in Section 4.2. Section 4.1 will display macroscopic statistics and analysis of all of our identified anomalous traceroutes.

4.1 Macro Statistics on Anomalous Traceroutes

As discussed in section 3.1, we are interested in finding sequential IPIDs and many IP addresses on a path in the same /24. To facilitate that, we can utilize our suspicious traceroute table within our database to query.

With our interest in ASs, we looked into what is the most commonly occurring ASN by hops and by traceroute count. Table 4.2 represents the top 5 ASNs that were flagged as having traceroute artifacts of 5 or more unique hops within the same /24 or 5 or more sequential IPIDs.

The percentage is based on the sum of all suspicious hops for a given ASN divided by the sum of all suspicious hops in the entire table. The table also includes ASN 0, which indicates that there was no associated ASN within the queried RIB file. These hops could be part of a traceroute artifact example that includes many hops, perhaps dozens, until the traceroute times out. These artifacts can sometimes extend many hops, and therefore could contribute to higher counts in the suspicious count column, skewing the results. However, we believe looking at the artifact hop count is still useful for understanding the ASN which contribute most toward the artifacts.

Table 4.3 accounts for the multiple hops potentially skewing results by only considering one ASN instance of an artifact per traceroute. In other words, even if an ASN were responsible for multiple consecutive IPIDs within a traceroute, it would only be counted once.

In both Tables 4.2 and 4.3, PROXAD is the leading contributor to the suspicious hop and traceroutes. PROXAD is a French telecommunications company, known as "Free" [48].

Suspicious Hop Count and Percentage by ASN							
ASN	AS Name	Suspicious Count	Percentage				
All	-	20,390,676	100%				
12322	PROXAD	6,743,134	33%				
0	-	5,177,520	25%				
174	Cogent Com.	1,379,537	7%				
786	JANET	679,482	3%				
3292	TDC	508,122	2%				
3356	LEVEL3	309,193	1.5%				

Table 4.2. Top ASN percentages for detected suspicious traceroutes by hop

Count of Suspicious Unique Trace-ASN Pairs and Percentage of Total by ASN							
ASN	AS Name	Suspicious Count	Percentage				
All	-	2,806,606	100%				
12322	PROXAD	399,570	14%				
0	-	329,651	12%				
174	Cogent Com.	172,975	6%				
9902	Neocomisp	95,928	3%				
701	Verizon	74,887	3%				
3292	TDC	72,835	3%				
16086	DNA Oyj	71,689	3%				
3356	LEVEL3	71,335	3%				

Table 4.3. Top ASN percentages for detected Suspicious Traceroutes. The "Suspicious Count" column represents the amount of unique traceroute-ASN pairs.

4.2 DNS Findings of Detected Suspicious Traceroutes

This DNS Findings section refers to searching for interesting DNS strings in the set of traceroutes already identified as having anomalous characteristics. As part of our research into suspicious traceroutes, we continue to look at the DNS names associated with the hops along these traceroute paths. While many of the hops of the suspicious traceroutes have no matching entry in the historical DNS database, the ones that do provide us with insights into the entities being crossed, such as large ISP names associated with their routers. Even more interesting are the traceroute gags. With known traceroute gags, we have the advantage of searching for strings within the gags to identify through manual examination of the DNS names paired with the suspicious traceroutes.

4.2.1 Star Wars Traceroute Gag Findings

Initially, we were unable to find the gag through a string search because it does not explicitly announce itself as a "Star Wars" themed traceroute, as can be seen in Figure 4.1. Searching

for "episode.iv" across the entire DNS dataset would have likely led to the traceroute gag given enough time. However, without knowing the approximate month, these grep searches can take a substantial amount of time for each string in the entire historical DNS dataset. Narrowing down searches to the month by looking for the traceroute gag characteristics first is far more efficient.

That being said, we searched for the string "star" in an attempt to be more thorough. The search returned more than just the Star Wars Traceroute gag, which we identified before the search completed by its anomalous characteristics. It also returns a large amount of suspicious traceroutes most of which are making use of MPLS from ASN 8359 (MTS PJSC), the largest mobile operator in Russia. These appear in the database as "alfabank.moscow.access.comstar.ru". Also, the Singapore Government Network, ASN 4657, has the name "Starhub-Internet", which appeared many times.

Without a specific string to search for and an approximate time, historical DNS searches can be unwieldy. However, when examining instances of anomalous traceroutes that demonstrated both of the anomalous characteristics (consecutive IPIDs and many IPv4 addresses in the same /24), we discovered a previously unknown Star Wars Traceroute Gag. This provides an example of our techniques being able to discover such traceroute gags in the "wild" Internet.

Our technique was to first identify our subset of suspicious traceroutes using the algorithms discussed in Section 3.1. We matched each response IP address with the corresponding historical DNS name. This allowed us to identify the iconic scrolling "Star Wars" title. Figure 4.1 demonstrates one instance in time of the discovered Star Wars Traceroute, as well as some of the corresponding data in our database. Figure 4.1 demonstrates the sequential IPIDs values as well as the IP addresses in the same /24.

hop_ip		ipid	dns
139.18.1.254		1	augate-if.rz.uni-leipzig.de
139.18.122.87	I	2	clusaug-if.rz.uni-leipzig.de
188.1.239.49	I	63591	cr-lap1-be3.x-win.dfn.de
188.1.144.26	I	58414	cr-tub2-hundredgige0-9-0-5.x-win.dfn.de
188.1.144.57	Ι	54013	cr-ham1-hundredgige0-1-0-0.x-win.dfn.de
188.1.244.178	I	0	kr-ros51.x-win.dfn.de
139.30.0.20	I	7	nx1-gate.uni-rostock.de
139.30.253.117	I	8050	FAIL.NON-AUTHORITATIVE.in-addr.arpa
139.30.208.195	Ι	9971	episode.ivl
139.30.208.135	Ι	9972	a.new.hopel
139.30.208.137	Ι	9973	a.long.time.ago.in.a.galaxy.far.far.away1
139.30.208.139	Ι	9974	<pre>it.is.a.period.of.civil.war.rebell</pre>
139.30.208.141	Ι	9975	<pre>spaceships.striking.from.a.hiddenl</pre>
139.30.208.143	Ι	9976	<pre>base.have.won.their.first.victoryl</pre>
139.30.208.145	Ι	9977	against.the.evil.galactic.empirel
139.30.208.147	Ι	9978	<pre>during.the.battle.rebel.spies.managed1</pre>
139.30.208.149	Ι	9979	<pre>to.steal.secret.plans.to.the.empiresl</pre>
139.30.208.151	Ι	9980	ultimate.weapon.the.death.star.anl
139.30.208.153	Ι	9981	armored.space.station.with.enoughl
139.30.208.155	Ι	9982	<pre>power.to.destroy.an.entire.planetl</pre>
139.30.208.157	Ι	9983	<pre>pursued.by.the.empires.sinister.agents1</pre>
139.30.208.159	Ι	9984	<pre>princess.leia.races.home.aboard.her1</pre>
139.30.208.161	I	9985	<pre>starship.custodian.of.the.stolen.plansl</pre>
139.30.208.163	I	9986	<pre>that.can.save.her.people.and.restore1</pre>
139.30.208.165	I	9987	<pre>freedom.to.the.galaxyl</pre>
139.30.208.132	I	26	<pre>may.the.force.be.with.youl</pre>
	F	igure 4.1.	Star Wars Traceroute discovered

Now that we had discovered this Star Wars Traceroute, we wanted to determine if the traceroute gag was still active. Performing multiple traceroutes from the original VP utilizing ICMP-Paris showed that the traceroute gag was not active. UDP-Paris and TCP traceroutes were attempted as well with no success. We also tried probing addresses within the same /24 as the original destination. The IP address 139.30.208.195 provides a response

to an ICMP-Paris or TCP probe, but not UDP-Paris. None of the other IP addresses in the same /24 respond. However, when searching for the individual DNS names by their corresponding hop IP addresses, they all still resolve to their original strings.

Since the traceroute was unable to be actively probed, we decided to attempt to identify when the traceroute gag stopped being operational. The instance demonstrated in Figure 4.1 is from 2016-01-02. One way to approximate when the traceroute gag ended is to look at the historical DNS. From viewing Figure 4.3, it appears that the traceroute gag was active into 2017 but stopped around late August. There are no DNS results for this artifact as of 2018-03-18.

./05/dns-names.17.20170525.tx	t.gz:1495672227	139.30.208.195	episode.ivl
./05/dns-names.17.20170514.tx	t.gz:1494786875	139.30.208.195	episode.ivl
./05/dns-names.17.20170531.tx	t.gz:1496229908	139.30.208.195	episode.ivl
./05/dns-names.17.20170527.tx	t.gz:1495866513	139.30.208.195	episode.ivl
./05/dns-names.17.20170510.tx	t.gz:1494437364	139.30.208.195	episode.ivl
./05/dns-names.17.20170506.tx	t.gz:1494080162	139.30.208.195	episode.ivl
./05/dns-names.17.20170530.tx	t.gz:1496132353	188.40.99.69	ridcully.episode-iv.de
$\sqrt{05/dns-names}$, 17, 20170518, tx	rt.gz:1495103735	139.30.208.195	enisode ivl
/05/dns-names.17.20170501.tx	rt.az:1493661085	139.30.208.195	episode.ivl
/11/dns-names 17 20171123 tx	rt az 1511479147	188 40 99 69	ridcully enisode-iv de
/02/dns-names 17 20170222 tx	rt az:1487779471	139 30 208 195	enisode ivl
$\sqrt{02}/dns-names 17 20170208 tx$	rt az:1486513411	139 30 208 195	enisode ivl
$\sqrt{02}/dns_names 17 20170200.tx$	rt az:1488106598	139 30 208 195	enisode ivl
$\sqrt{02}/dns_names 17.20170220.tx}$	rt az:1486167205	139 30 208 195	enisode ivl
(02)/dnc names 17.20170204.tx	rt. gz. 1480107205	120 20 200 105	episode iv
$(02/d12 - 11d1125 \cdot 17 \cdot 20170214 \cdot 1x)$.t.yz.1407103172	120 20 200 105	episode iv
(0°) (dnc names 17.20170805.1X	.t.gz.1501945600	120 20 200 105	episode iv
(02/drs-names.17.20170808.tx	11.yz:1502195958	159.50.208.195	episode.ivi
./02/dns-names.1/.201/0202.tx	t.gz:1486004348	188.40.99.69	ridcully.episode-iv.de
./08/dns-names.1/.201/0802.tx	t.gz:1501694131	139.30.208.195	episode.ivi
./08/dns-names.1/.201/0825.tx	t.gz:1503682250	139.30.208.195	episode.ivl
./06/dns-names.17.20170618.tx	t.gz:1497783001	139.30.208.195	episode.ivl
./06/dns-names.17.20170611.tx	t.gz:1497174028	139.30.208.195	episode.ivl
./06/dns-names.17.20170630.tx	t.gz:1498846140	188.40.99.69	ridcully.episode-iv.de
./06/dns-names.17.20170626.tx	t.gz:1498503489	139.30.208.195	episode.ivl
./06/dns-names.17.20170629.tx	t.gz:1498758366	139.30.208.195	episode.ivl
./06/dns-names.17.20170621.tx	t.gz:1498043500	139.30.208.195	episode.ivl
./12/dns-names.17.20171207.tx	t.gz:1512643910	139.30.208.195	episode.ivl
./01/dns-names.17.20170106.tx	t.gz:1483703370	139.30.208.195	episode.ivl
./01/dns-names.17.20170113.tx	t.gz:1484321095	139.30.208.195	episode.ivl
./01/dns-names.17.20170130.tx	t.gz:1485754461	139.30.208.195	episode.ivl
./09/dns-names.17.20170922.tx	t.gz:1506069853	188.40.99.90	<pre>frontend.episode-iv.de</pre>
./09/dns-names.17.20170907.tx	t.gz:1504795556	139.30.208.195	episode.ivl
./03/dns-names.17.20170323.tx	t.gz:1490292614	139.30.208.195	episode.ivl
./03/dns-names.17.20170309.tx	t.gz:1489058381	139.30.208.195	episode.ivl
./03/dns-names.17.20170331.tx	t.gz:1490989903	139.30.208.195	episode.ivl
./03/dns-names.17.20170324.tx	t.gz:1490345584	188.40.99.90	<pre>frontend.episode-iv.de</pre>
./03/dns-names.17.20170320.tx	t.gz:1490040021	139.30.208.195	episode.ivl
./03/dns-names.17.20170313.tx	t.gz:1489393532	139.30.208.195	episode.ivl
./03/dns-names.17.20170317.tx	t.gz:1489770287	139.30.208.195	episode.ivl
./04/dns-names.17.20170416.tx	t.gz:1492360292	139.30.208.195	episode.ivl
./04/dns-names.17.20170409.tx	t.gz:1491769459	188.40.99.90	<pre>frontend.episode-iv.de</pre>
./04/dns-names.17.20170423.tx	t.gz:1492942358	139.30.208.195	episode.ivl
./04/dns-names.17.20170428.tx	t.gz:1493373646	139.30.208.195	episode.ivl
./04/dns-names.17.20170406.tx	t.gz:1491500736	139.30.208.195	episode.ivl
./04/dns-names.17.20170411.tx	t.gz:1491923191	139.30.208.195	episode.ivl
./04/dns-names.17.20170420.tx	t.gz:1492703124	139.30.208.195	episode.ivl
./07/dns-names.17.20170708.tx	t.gz:1499521812	139.30.208.195	episode.ivl
./07/dns-names.17.20170705.tx	t.gz:1499255449	139.30.208.195	episode.ivl
./07/dns-names.17.20170711.tx	t.gz:1499798070	139.30.208.195	episode.ivl
./07/dns-names.17.20170724 tx	t.gz:1500903443	207.154.246.96	algo.episode-iv.de
./07/dns-names.17.20170730 +x	t.gz:1501421743	139.30.208.195	episode.iv1
./07/dns-names.17.20170706.tx	t.gz:1499364368	207.154.246.96	algo.episode-iv.de
./07/dns-names 17 20170714 +x	t.gz:1500030879	139.30.208 195	episode.iv
/07/dns-names 17 20170702 +v	rt.gz:1498991657	139.30.208 195	enisode iv1
.,,			-r 1

Figure 4.2. Star Wars Traceroute search for "episode.iv" within the historical DNS in 2017. The command used: find . -name "*.gz" | xargs -n 1 -P 10 zgrep -H "episode.iv"

Figure 4.2 shows the 2017 results for a search of "episode.iv". Within this search, we also noticed a "frontend.episode-iv.de" and "ridcully.episode-iv.de", which are present in the previous years as well. Seeing that one of these DNS names is "frontend", it would seem logical to attempt to access the IP address associated with it. Both IP addresses point to the simple website shown in Figure 4.4. Traceroutes to these addresses do not produce the traceroute gag.



Figure 4.4. Both IP addresses associated with "frontend.episode-iv.de" and "ridcully.episode-iv.de" point to this website as of 2018-03-11

4.2.2 Bad.horse Traceroute Gag Findings

The Bad.horse traceroute gag is known to us publicly through news and enthusiast forms [49], [50]. We were able to find 5 instances of the Bad.horse Traceroute gag in our CAIDA traceroute data, which is still active at the time of this writing [34]. See Figure 4.5, which shows a single traceroute to the "bad.horse" destination using the default traceroute client with default settings in linux. The returning DNS names are part of a song from the

referenced video. Visiting "bad.horse" in a web browser will provide you with a video associated with the traceroute gag.

traceroute to bad.horse (162.252.205.157), 50 hops max, 60 byte packets 1 thunder.caida.org (192.172.226.253) 0.359 ms 1.037 ms 1.009 ms 2 bwctl9.sdsc.edu (192.12.207.9) 0.325 ms 0.353 ms 0.338 ms 3 dc-sdg-agg4--sdsc-1.cenic.net (137.164.23.129) 0.782 ms 0.813 ms 0.894 ms 4 137.164.11.10 (137.164.11.10) 2.450 ms 2.481 ms dc-tus-agg3--sdg-agg4-100ge.cenic.net (137.164.11.8) 2.517 ms 5 137.164.11.22 (137.164.11.22) 3.269 ms 3.278 ms 3.265 ms 6 * * * 7 * * * 8 las-b21-link.telia.net (80.239.195.60) 3.061 ms 3.575 ms 3.537 ms 9 las-b24-link.telia.net (62.115.136.47) 12.868 ms 11.703 ms 11.672 ms 10 atlanticmetro-ic-306052-las-b3.c.telia.net (62.115.42.54) 3.210 ms 3.122 ms 3.249 ms 11 ve1450.cr2.lga3.atlanticmetro.net (108.60.128.73) 75.686 ms 70.486 ms 71.880 ms 12 e6-3.cr1.lga12.atlanticmetro.net (108.60.148.78) 73.800 ms 73.631 ms 70.503 ms 13 e2-20.cr2.lga11.atlanticmetro.net (69.9.32.221) 73.366 ms 72.215 ms 74.325 ms 14 t03.nycmc1.ny.us.sn11.net (162.252.205.4) 101.313 ms 94.933 ms 93.523 ms 15 bad.horse (162.252.205.130) 93.975 ms 92.837 ms 93.328 ms 16 bad.horse (162.252.205.131) 104.648 ms 107.939 ms 106.458 ms 17 bad.horse (162.252.205.132) 108.636 ms 109.131 ms 103.074 ms 18 bad.horse (162.252.205.133) 108.648 ms 107.530 ms 112.525 ms 19 he.rides.across.the.nation (162.252.205.134) 116.155 ms 119.027 ms 115.260 ms 20 the.thoroughbred.of.sin (162.252.205.135) 129.049 ms 124.693 ms 128.101 ms 21 he.got.the.application (162.252.205.136) 123.583 ms 126.720 ms 126.717 ms 22 that.you.just.sent.in (162.252.205.137) 130.120 ms 126.396 ms 128.390 ms 23 it.needs.evaluation (162.252.205.138) 150.720 ms 144.407 ms 144.404 ms 24 so.let.the.games.begin (162.252.205.139) 143.431 ms 144.525 ms 143.636 ms 25 a.heinous.crime (162.252.205.140) 151.507 ms 151.276 ms 150.888 ms 26 a.show.of.force (162.252.205.141) 151.534 ms 151.292 ms 153.595 ms 27 a.murder.would.be.nice.of.course (162.252.205.142) 154.684 ms 160.003 ms 158.527 ms 28 bad.horse (162.252.205.143) 163.447 ms 165.010 ms 164.151 ms 29 bad.horse (162.252.205.144) 176.849 ms 174.031 ms 175.667 ms 30 bad.horse (162.252.205.145) 171.615 ms 187.756 ms 175.138 ms 31 he-s.bad (162.252.205.146) 177.826 ms 181.542 ms 179.771 ms 32 the.evil.league.of.evil (162.252.205.147) 183.206 ms 186.902 ms 183.323 ms 33 is.watching.so.beware (162.252.205.148) 184.541 ms 184.739 ms 182.833 ms 34 the.grade.that.you.receive (162.252.205.149) 193.063 ms 190.923 ms 189.881 ms 35 will.be.your.last.we.swear (162.252.205.150) 205.969 ms 202.697 ms 203.023 ms 36 so.make.the.bad.horse.gleeful (162.252.205.151) 208.786 ms 200.288 ms 209.764 ms 37 or.he-ll.make.you.his.mare (162.252.205.152) 205.680 ms 205.338 ms 206.433 ms 38 o_o (162.252.205.153) 213.477 ms 213.171 ms 213.788 ms 39 you-re.saddled.up (162.252.205.154) 220.033 ms 218.396 ms 217.518 ms 40 there-s.no.recourse (162.252.205.155) 228.576 ms 231.465 ms 224.031 ms 41 it-s.hi-ho.silver (162.252.205.156) 225.696 ms 227.416 ms 237.066 ms 42 signed.bad.horse (162.252.205.157) 235.920 ms 228.957 ms 230.351 ms Figure 4.5. Traceroute from 2018-3-29 using the command: traceroute -m

50 bad.horse

We discovered this specific, known, gag by simply searching for the string "horse" in our suspicious traceroute table. Figure 4.6 demonstrates what was returned when searching for the string in the table.

<pre>traceanalysis=> select hop_ip, trace_id, dns from suspicious where dns LIKE '%horse%';</pre>	
hop_ip trace_id dns	
162.252.205.145 3250435396 bad.horse	
162.252.205.131 3250435396 bad.horse	
162.252.205.130 2284610512 bad.horse	
162.252.205.131 2284610512 bad.horse	
162.252.205.132 2284610512 bad.horse	
162.252.205.133 2284610512 bad.horse	
162.252.205.143 2284610512 bad.horse	
162.252.205.144 2284610512 bad.horse	
162.252.205.145 2284610512 bad.horse	
162.252.205.157 2284610512 signed.bad.horse	
162.252.205.151 2284610512 so.make.the.bad.horse.gleeful	
162.252.205.194 1153505021 bad.horse	
162.252.205.133 3250435396 bad.horse	
162.252.205.132 3250435396 bad.horse	
162.252.205.143 3250435396 bad.horse	
162.252.205.144 3250435396 bad.horse	
162.252.205.151 3034559376 so.make.the.bad.horse.gleeful	
162.252.205.157 3034559376 signed.bad.horse	
162.252.205.151 3250435396 so.make.the.bad.horse.gleeful	
162.252.205.157 3250435396 signed.bad.horse	
(20 rows)	
Figure 4.6. Bad.horse instances found in the suspicious traceroute table'	s
DNS column	

Figure 4.7 displays a complete instance of a Bad.horse Traceroute found within our dataset. It clearly demonstrates the artifact characteristic of having many IP addresses within the same /24, but it does not demonstrate the consecutive IPIDs characteristic. This particular example of finding a traceroute gag acts as validation that we are able to find active traceroute gags within our database.

hop_ip | ipid | dns ------84.88.81.121 | 0 | FAIL.NON-AUTHORITATIVE.in-addr.arpa 84.88.19.149 | 56446 | anella-ccaba-upc.cesca.cat 130.206.211.69 | 0 | anella-val1.ae2-454.uv.rt1.val.red.rediris.es 130.206.245.90 | 0 | uv.ae5.telmad.rt4.mad.red.rediris.es 62.40.124.192 | 0 | rediris.mx1.gen.ch.geant.net 62.40.98.113 | 0 | ae1.mx1.mil2.it.geant.net 217.29.66.125 | 39158 | he.mix-it.net 184.105.222.129 | 41703 | 10ge3-3.core1.zrh1.he.net 184.105.222.49 | 46401 | 10ge7-17.core1.par2.he.net 184.105.81.77 | 63769 | 100ge7-1.core1.nyc4.he.net 184.105.223.161 | 58246 | 100ge7-2.core1.chi1.he.net 184.105.223.178 | 34737 | 100ge10-1.core1.msp1.he.net 216.66.78.110 | 0 | ip-house.gigabitethernet3-6.core1.msp1.he.net 216.250.189.170 | 56428 | c4500-1.mpls.iphouse.net 209.240.64.149 | 4091 | egw-iphouse.mplsc1.mn.us.sn11.net 162.252.204.66 | 28385 | t00.nycmc1.ny.us.sn11.net 162.252.205.130 | 63695 | bad.horse 162.252.205.131 | 54557 | bad.horse 162.252.205.132 | 34837 | bad.horse 162.252.205.133 | 20304 | bad.horse 162.252.205.134 | 3166 | he.rides.across.the.nation 162.252.205.135 | 47726 | the.thoroughbred.of.sin 162.252.205.136 | 49814 | he.got.the.application 162.252.205.137 | 6048 | that.you.just.sent.in 162.252.205.138 | 21211 | it.needs.evaluation 162.252.205.139 | 30008 | so.let.the.games.begin 162.252.205.140 | 32623 | a.heinous.crime 162.252.205.141 | 25898 | a.show.of.force 162.252.205.142 | 22754 | a.murder.would.be.nice.of.course 162.252.205.143 | 24325 | bad.horse 162.252.205.144 | 9582 | bad.horse 162.252.205.145 | 32850 | bad.horse 162.252.205.146 | 27480 | he-s.bad 162.252.205.147 | 6554 | the.evil.league.of.evil 162.252.205.148 | 43612 | is.watching.so.beware 162.252.205.149 | 36701 | the.grade.that.you.receive 162.252.205.150 | 47416 | will.be.your.last.we.swear 162.252.205.151 | 38223 | so.make.the.bad.horse.gleeful 162.252.205.152 | 41114 | or.he-ll.make.you.his.mare 162.252.205.153 | 55506 | o\x5fo 162.252.205.154 | 65279 | you-re.saddled.up 162.252.205.155 | 58301 | there-s.no.recourse 162.252.205.156 | 44206 | it-s.hi-ho.silver 162.252.205.157 | 31904 | signed.bad.horse (44 rows)

traceanalysis=> select hop_ip, ipid, dns from suspicious where trace_id = 2284610512 order by probe_ttl;

Figure 4.7. One instance of Bad.horse Traceroute with corresponding DNS names.

4.3 Traceroutes with Both Anomalous Characteristics

In this section, we analyze the 73 traceroutes that display consecutive IPIDs and Many IP Addresses in same /24. These 73 traceroutes involve repetitions of the same or similar destinations in the same /24 over the 5 years of this CAIDA traceroute data, indicating an approximate life span of the artifact. For example, 89.133.20.0/24 appears multiple times from 2013 until 2016. It does not make any appearance in 2017. Some of these traceroutes include a highly unusual amount of responses ranging from hundreds to over 2,000 total responses to less than 20 probes. We discuss these specific anomalous traceroutes in this section in greater detail and describe them as a new classification of artifact.

4.3.1 Befree Destinations

Nearly half of our 73 identified anomalous traceroutes are the same artifact repeated over 4 years. 32 of our 73 traceroutes that display both anomalous characteristics are in the 89.113.20.0/24 address space. Table 4.4 shows the distribution over our CAIDA traceroute data timeline. Many of these traces traverse through an IP address with a DNS name of "gw-befree.retn.net" before continuing on to repeated responses of 89.113.16.69 with incrementing IPIDs. The IP address associated with "gw-befree.retn.net" belongs to ASN 9002 (RETN), an international ISP. This IP address appears to be a gateway, which offers the numerous incrementing IPIDs responses all of which belong to ASN 39749 (Befree) and are of the same IP address.

Other instances of these 32 that do not traverse the "befree gateway", but still have Befree destinations, instead traverse Russian domain names, such as "spb-81-211-88-182.sovintel.spb.ru", associated with ASN 3216 (PJSC VimpelCom). One instance of this traceroute, which traverses the Russian domain, originated from a VP in Canada. Another instance run just prior to this was from a VP in the Netherlands and resulted in the "befree gateway".

Running a traceroute in 2018 with UDP-Paris, from the same Canadian VP, to 89.113.20.25 stops at the same IP address with the domain name "spb-81-211-88-182.sovintel.spb.ru". Running another traceroute with the same settings but from a VP in the Netherlands results in only 5 hops and then stops before proceeding as far as from the Canadian VP. Probing with ICMP-Paris and TCP produces the same result.

89.113.20.0/24 appearance by year					
Year	Count				
All	32				
2013	5				
2014	17				
2015	9				
2016	1				
2017	0				

Table 4.4. Count of the 89.113.20.0/24 anomalous traces

4.3.2 Russian ISP

Within our group of 73 suspicious traceroutes that display both characteristics over 5 years, 3 of the 73 destinations belong to the Russian ISP, MTS. MTS is also known as Comstar, which is the name that appears in the historical DNS associated with the responses. These 3 traces have destinations in the 89.175.85.0/24 space and were active in 2013, with the last known one occurring on 2016-04-26. Nearly all of the responses contain ICMP extension, indicating MPLS, although it is not clear that the trace traverses a tunnel. Many of the responses including ICMP extension are multiple responses to a single probe, perhaps indicating that the trace does not actually traverse the tunnel but instead reveals multiple different MPLS addresses. These addresses could be different MPLS ingress routers or interfaces. Figure 4.8 demonstrates some of the over 600 responses of one of the instances.

hop_ip	Ι	ipid	I	asn	I	probe	I	dns
195 206 249 17	+-	 13111	+-	49770	+	 2	+-+	FAIL NON_AUTHORITATIVE in_addr arna
05 1/3 207 173	1	11803	1	40770	÷	2	1	FAIL NON AUTHORITATIVE in addr arpa
95.145.207.175	1	44005	1	49770	÷	ر ۸	1	TALL.NON-AUTHORITATIVE.III-audi.aipa
95.145.207.229		21200	1	49770	-	4	-	mx-corel.internetport.se
213.132.98.113		31269	!	12552	1	5		FAIL.NON-AUTHORITATIVE.in-addr.arpa
83.145.21.205	1	0	1	24862	1	6		FAIL.NON-AUTHORITATIVE.in-addr.arpa
194.68.128.161		13897	I	9902	I	7	I	netnod-ix-ge-b-sth-1500.mts.ru
195.34.50.166	I	58460	I	8359	I	8	I	oct-cr03-be5.278.spb.stream-internet.net
212.188.2.38	Ι	786	I	8359	I	9	I	<pre>mag9-cr01-be1.78.msk.stream-internet.net</pre>
195.34.50.150	Ι	42526	I	8359	I	10	I	<pre>ss-cr04-be2.77.msk.stream-internet.net</pre>
195.34.59.106	Ι	23379	I	8359	I	11	I	a197-cr04-be5.77.msk.stream-internet.net
195.34.53.5	Ι	27287	I	8359	I	12	I	<pre>ss-cr02-po2.221.msk.stream-internet.net</pre>
195.210.128.105	Ι	0	I	8359	I	13	I	asr-sm-0.moscow.core.comstar.ru
212.248.3.221	Ι	9194	I	8359	I	14	I	bb-sm-1-ge-4-22.moscow.core.comstar.ru
82.204.255.246	Ι	48226	I	8359	I	15	I	bb-tg-0-te5-3.moscow.core.comstar.ru
212.248.3.146	Ι	38232	I	8359	I	16	I	ngn-7606-te3-0-1.moscow.core.comstar.ru
89.175.143.33	Ι	27233	I	8359	I	19	I	<pre>rusdel4.moscow.access.comstar.ru</pre>
89.175.143.190	Ι	40936	I	8359	I	19	I	<pre>interprombank2.moscow.access.comstar.ru</pre>
89.175.226.45	Ι	23196	I	8359	I	19	I	FAIL.SERVER-FAILURE.in-addr.arpa
89.175.143.179	Ι	12564	Ι	8359	I	19	I	interprombank.moscow.access.comstar.ru
89.175.143.19	I	34784	I	8359	I	19	I	newstep.moscow.access.comstar.ru
89.175.143.18	I	58451	I	8359	I	19	I	monarossi.moscow.access.comstar.ru
89.175.143.73	Ι	63660	I	8359	I	19	I	<pre>mroformitel.moscow.access.comstar.ru</pre>
89.175.143.31	I	15563	I	8359	I	19	I	<pre>rusdel2.moscow.access.comstar.ru</pre>
89.175.143.152	Ι	43174	I	8359	I	19	I	bershka.moscow.access.comstar.ru

Figure 4.8. Partial example of one instance of the Comstar anomalous traceroutes

Another instance of this Comstar anomaly is demonstrated in Figure 4.9, with over 1000 responses. Hundreds of responses are generated in response to probe 30 alone, with corresponding DNS names that refer to different businesses (interprombank, monex, zara). Probe 31 then generates hundreds more responses with 89.175.143.1, "cr-145-vl2620.access.comstar.ru", as the only response. This name may potentially refer to the model of the router, perhaps a Cisco 2620, while "vl" could refer to "Virtual Lan". Figure 4.10 shows the third instance, which has over 700 responses. Figure 4.10 demonstrates the same anomaly observed in Figure 4.9, however the order of the probe responses are reversed.

89.175.143.1	22150 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22151 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22128 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22183 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22152 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22153 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22154 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22155 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22156 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.1	22157 8359	30 cr-145-vl2620.access.comstar.ru
89.175.143.249	47233 8359	<pre>31 diva.moscow.access.comstar.ru</pre>
89.175.143.111	839 8359	<pre>31 goldstreet.moscow.access.comstar.ru</pre>
89.175.143.107	46023 8359	<pre>31 fashiongroup.moscow.access.comstar.ru</pre>
89.175.143.34	60779 8359	31 skyter.moscow.access.comstar.ru
89.175.143.107	46029 8359	<pre>31 fashiongroup.moscow.access.comstar.ru</pre>
89.175.143.34	60785 8359	31 skyter.moscow.access.comstar.ru
89.175.143.18	33367 8359	<pre>31 monarossi.moscow.access.comstar.ru</pre>
89.175.143.145	15503 8359	31 lpk.moscow.access.comstar.ru
89.175.143.107	46034 8359	<pre>31 fashiongroup.moscow.access.comstar.ru</pre>
89.175.143.236	52386 8359	<pre>31 italikam.moscow.access.comstar.ru</pre>
89.175.143.89	43502 8359	31 edinayevropasb.moscow.access.comstar.ru
89.175.143.249	47234 8359	<pre>31 diva.moscow.access.comstar.ru</pre>
89.175.143.20	29008 8359	<pre>31 system-fd.moscow.access.comstar.ru</pre>
89.175.143.111	840 8359	<pre>31 goldstreet.moscow.access.comstar.ru</pre>



89.175.226.23	Ι	42866	Τ	8359	1	30		adidas.moscow.access.comstar.ru
89.175.226.56	Ι	37857	Τ	8359	1	30		edan.moscow.access.comstar.ru
89.175.143.42	Ι	53581	Τ	8359	1	30		<pre>sanfeshn.moscow.access.comstar.ru</pre>
89.175.143.7	Ι	24476	Ι	8359	1	30		<pre>veter2.moscow.access.comstar.ru</pre>
89.175.143.100	Ι	59429	Τ	8359	1	30		restorants.moscow.access.comstar.ru
89.175.226.32	Ι	55588	Τ	8359	1	30		adidas.moscow.access.comstar.ru
89.175.143.221	Ι	55086	Ι	8359	1	30		ravelo.moscow.access.comstar.ru
89.175.226.5	Ι	54609	T	8359	1	30		<pre>interjazz.moscow.access.comstar.ru</pre>
89.175.226.56	Ι	37858	Τ	8359	1	30		edan.moscow.access.comstar.ru
89.175.226.28	Ι	53634	Ι	8359	1	30		rusmoda.moscow.access.comstar.ru
89.175.226.33	Ι	7099	T	8359	1	30		<pre>nrgroup.moscow.access.comstar.ru</pre>
89.175.143.42	Ι	53582	Τ	8359	1	30		<pre>sanfeshn.moscow.access.comstar.ru</pre>
89.175.143.7	Ι	24477	Ι	8359	1	30		<pre>veter2.moscow.access.comstar.ru</pre>
89.175.143.100	Ι	59430	Ι	8359	1	30		${\tt restorants.moscow.access.comstar.ru}$
89.175.143.1	Ι	48906	T	8359	1	31		cr-145-vl2620.access.comstar.ru
89.175.143.1	Ι	48592	Ι	8359	1	31		cr-145-vl2620.access.comstar.ru
89.175.143.1	Ι	48594	T	8359	1	31		cr-145-vl2620.access.comstar.ru
89.175.143.1	Ι	48595	Τ	8359	1	31		cr-145-vl2620.access.comstar.ru
89.175.143.1	Ι	48596	Ι	8359	1	31	I	cr-145-vl2620.access.comstar.ru
89.175.143.1	Ι	48597	T	8359	1	31	I	cr-145-vl2620.access.comstar.ru

Figure 4.10. Reverse order of responses to probes for another instance

Active probes to the original destinations do not replicate this anomaly and appear to stop

at "last-hop2-ngn.comstar.ru" before the gap limit is exceeded. This could possibly indicate that this was unintended information leakage for this ISP and they corrected the problem sometime between 2016 and 2017.

4.3.3 Other Traceroutes with Many Responses to a Single Probe

This newly observed characteristic of many responses to a single probe seems to be common across a number of the 73 traceroutes exhibiting both characteristics, with some ranging from a few responses while others contain thousands. Section 4.3.1 is an example of this characteristic of multiple responses to a single probe. The traces actually do not demonstrate many IP addresses in the same /24, since it does not display multiple hops. Instead, multiple identical responses to a single probe are provided.

We observe one instance of a traceroute with over 2000 responses. Similar to the artifact discussed in 4.3.2, there are many identical responses to one probe followed by thousands of responses in the 200.32.179.0/24 address space to another probe. With this particular artifact, there are no corresponding DNS results. Active probes sent to the original destination do not demonstrate the artifact, which first appeared in 2013.

While we correctly identified these traceroutes as being anomalous, they do not fit well into our original artifact classifications. Therefore, we propose a new characteristic to look for when researching traceroute artifacts. In addition to traceroutes that contain many hops with IP addresses in the same /24, traceroutes with many responses to a single probe should be considered as anomalous behavior. However, these artifacts are not mutually exclusive. Just as in this section, we identified 73 instances of traceroutes displaying sequential IPIDs and many IP addresses in the same /24, it is possible for a subset of these to also contain the many responses characteristic. Figure 4.11 demonstrates the possible overlap of anomalous behaviors. For example, most traceroutes observed with many responses to a single probe likely had those responses come from a single interface or router due to the incrementing IPIDs that were also present.



Figure 4.11. Venn diagram of different traceroute artifact behaviors

CHAPTER 5: Longitudinal Analysis of MPLS Results

As stated in Chapter 4, the goal of this thesis is to examine a large dataset of historic IPv4 traceroute data in order to identify and classify artifacts. In pursuit of this goal, we formed a hypothesis that MPLS could be responsible for a subset of these artifacts. Using the methodology described in Chapter 3, we hope to discover trends in the prevalence of MPLS tunnels that may influence the frequency in appearance of these artifacts. We also hope to discover deceptive practices in the "wild" Internet by studying these historical datasets. We will focus our research on the two artifact types discussed in subsections 3.1.1 (Sequential IPIDs) and 3.1.2 (Many IP addresses in the same /24).

5.1 MPLS Prevalence Across all VPs

Figure 5.1 is a boxplot representation of all VPs in our dataset. For each year, every instance of a VP is represented by the percentage of traceroutes that traversed a MPLS tunnel. The x-axis represents each year of our dataset, from 2013 until 2017. The y-axis represents the percentage of traceroutes that traverse an MPLS tunnel. While the median percentage in 2013 is 46%, many outliers appear near 100%. This could indicate that these monitors are only a few hops away from an ISP utilizing MPLS, guaranteeing most destinations must go through a tunnel. There also appears to be outliers on the bottom tails that may indicate that these VPs are passing through a high number of ASs that are not utilizing MPLS or not advertising their use of MPLS.



Figure 5.1. Percentage of traces that traverse a MPLS tunnel across all Team 1 vantage points

Table 5.1. Median percentages,	by year, for all VF	^O traceroutes traversing a	n
MPLS tunnel			

Change in the median percent of MPLS tunnels for all VPs by year				
Year	Percentage			
2013	46%			
2014	41%			
2015	33%			
2016	34%			
2017	36%			

MPLS prevalence across all monitors in the entire dataset for team 1 from 2013 through 2017 remained relatively stable, although the median percentage did drop to a low of 33% in 2015, a trend shared by many of the VPs surveyed later in this chapter. The median percentage of traceroutes that traversed an MPLS tunnel range from approximately 33% in 2015 to

46% in 2013, which is noticeably higher than the 30% in 2011, discovered by previous research [35]. Despite being higher, the percentage is still within a reasonable difference when considering MPLS utilization has likely increased as large ISP continually optimize their networks for increased traffic since 2011. Donnet et al. discovered approximately 30% of paths tested traversed an MPLS tunnel from data collected in 2011 [35]. They utilized 25 PlanetLab VPs and collected all data on 2011-08-24 [35]. Donnet et al. used scamper, the same tool that CAIDA uses to collect their CTD [21].

However, when looking at the data more closely it is apparent that there are large variations between certain monitors. Figure 5.1 demonstrates this quite clearly by showing many outliers in the boxplot, specifically the many outliers near 100% of traceroutes that traverse a MPLS router in years 2013 and 2015. By querying our database for the VPs with the largest MPLS percentages per cycle in yearly time periods, we are able to identify one monitor that is contributing to these outliers. Tables 5.2 and 5.3 demonstrates that New Orleans, Louisiana (msy-us) is one such VP contributing to the outliers in 2013 and 2014, likely because of being a few hops away from a large ISP heavily utilizing MPLS. We will delve into the msy-us VP later in Section 5.2.2.

Top 10 VPs by percentage of traceroutes with MPLS 2013		
Monitor	Date	Percentage
msy-us	2013-11-08	100%
msy-us	2013-11-26	100%
msy-us	2013-12-12	100%
msy-us	2013-11-14	100%
msy-us	2013-10-26	99%
msy-us	2013-10-18	99%
msy-us	2013-12-23	99%
msy-us	2013-10-21	99%
msy-us	2013-10-18	99%
msy-us	2013-12-12	99%

Table 5.2. msy-us is responsible for a large number of cycles which had total or near 100% MPLS tunnels per trace in 2013.

Top 10 VPs by percentage of traceroutes with MPLS 2014			
Monitor	Date	Percentage	
msy-us	2014-09-22	100%	
msy-us	2014-03-13	100%	
msy-us	2014-08-09	100%	
msy-us	2014-01-24	100%	
msy-us	2014-07-10	100%	
msy-us	2014-05-17	100%	
msy-us	2014-07-10	100%	
msy-us	2014-01-27	100%	
msy-us	2014-03-31	100%	
msy-us	2014-04-05	100%	

Table 5.3. msy-us had many cycles in 2014 that had 100% of traceroutes traverse an MPLS tunnel

Therefore, due to these large variations between different VPs, it is necessary to look at percentages of traces traversing an MPLS router on a per-vantage point basis to better understand these trends. The following vantage points come from team 1 in the dataset. They are spread across the globe with at least one monitor positioned on the American, South American, African, Asian, Australian, and European continents. Many of the vantage points displayed dramatic increases or decreases in the percentage of traces that traversed an MPLS tunnel from year to year. One potential cause could be a large ISP moving away from utilizing MPLS for optimization of traffic and moving toward another technology. Another possibility could be that peering arrangements and agreements have changed between the large ISPs. And yet another possibility is that the large ISPs have not reduced their usage of MPLS at all, but instead are no longer advertising publicly the MPLS within their infrastructure.

5.2 Individual Analysis of MPLS Prevalence by VPs

If a large ISP is no longer utilizing MPLS, we would expect to see that particular ISP AS number appearance to drop from one year to the next. The first step is to find all traceroute
hop IP addresses associated with MPLS, meaning we have determined them to be a hop as part of an MPLS tunnel, and match them with their corresponding AS number. In our database, we stored hops of a traceroute that appeared to have any form of MPLS (explicit, implicit, or opaque). By looking at these IP addresses and matching them with an AS number utilizing a route views RIB file from that month and year, we can then gather a count of the AS numbers for that vantage point for that year. If we notice a trend across multiple VPs for a specific ASN, our confidence that the ASN is the responsible party to the change in frequency of MPLS will be higher.

5.2.1 Dublin, Ireland, VP

Figure 5.2 represents Dublin, Ireland's VP over a time period of 5 years as a boxplot. This represents every single cycle that the Dublin VP participated in during this time period. The x-axis represents each year of our dataset, between 2013 until 2017. The y-axis represents the percentage of traceroutes that traverse an MPLS tunnel. There is a noticeable drop in the percentage of traces that traverse an MPLS tunnel between 2014 and 2015. Table 5.4 displays the percentage change between years, 2015 being the lowest.



Figure 5.2. Dublin, Ireland

Change in percent of MPLS tunnels for dub-ie VP by year		
Year	Percentage	
2013	56%	
2014	47%	
2015	27%	
2016	30%	
2017	31%	

Table 5.4. Dublin, Ireland, VP traceroutes traversing an MPLS tunnel percentages by year

Table 5.5. Top 10 ASN appearances for Dublin VP in 2014

Change in percent of MPLS tunnels for dub-ie VP by year		
ASN	Count	
3356	134,002,617	
7018	18,387,985	
2914	9,054,786	
6453	5,630,394	
2828	1,902,673	
6830	1,541,080	
7843	1,450,913	
8928	1,242,488	
4538	1,070,583	
15557	1,034,256	

By examining the count of each hop's ASN in Table 5.5, we see that the top result in 2014 was ASN 3356, which according to as-rank.caida.org [48] is the AS number associated with Level 3 Communications. Level 3 accounted for 67% of all MPLS tunnels traversed in 2014 from the Dublin, Ireland vantage point. The number two most traversed ASN with MPLS tunnels was ASN 7018 with 9% of all MPLS tunnels traversed in 2014 from the Dublin, Ireland so the set of all MPLS tunnels traversed in 2014 from the Dublin, Ireland set of all MPLS tunnels traversed in 2014 from the Dublin, Ireland set of all MPLS tunnels traversed in 2014 from the Dublin, Ireland set of all MPLS tunnels traversed in 2014 from the Dublin, Ireland monitor. ASN 7018 belongs to AT&T Services.

Change in percent of MPLS tunnels for dub-ie VP by year		
ASN	Count	
7018	16,715,880	
2914	12,154,127	
6453	6,478,810	
2828	3,670,696	
7843	2,999,685	
6461	2,141,185	
8928	1,617,884	
6830	1,505,907	
15557	1,430,427	
2603	1,294,895	

Table 5.6. Top 10 ASN appearances for Dublin VP in 2015

ASN 3356 went from a count of more than 134 million in 2014 to not even being represented in the top 10 in 2015. ASN 3356 count for 2015 was 48,291, representing a mere .06% of all MPLS tunnels. ASN 7018 moved up in position to number one, accounting for 21% of all MPLS tunnels.

Change in percent of MPLS tunnels by notable AS				
ASN	AS Name	2014 Percentage	2015 Percentage	
3356	LEVEL3	67%	<1%	
7018	AT&T	9%	21%	
2914	NTT America	4.5%	16%	
6453 TATA COM. 2.7% 8%			8%	
2828	Verizon Business	1%	5%	

Table 5.7. Dublin, Ireland, VP notable ASN changes

Figure 5.7 demonstrates the changes of notable ASNs in percentage of hops detected to be part of an MPLS tunnel between 2014 and 2015.

5.2.2 New Orleans, Louisiana, United States, VP

Figure 5.3 representing New Orleans displays years 2013 and 2014 as having near 100% of traces traversing an MPLS tunnel, then suddenly dropping down to near 30% in 2015. Figure 5.3 represents information for the New Orleans VP in the same way that it was represented in subsection 5.2.1.



Figure 5.3. Percentage of traceroutes that traverse a MPLS tunnel for the VP in New Orleans, Louisiana, United States

Change in percent of MPLS tunnels for msy-us VP by year		
Year	Percentage	
2013	99.9%	
2014	83%	
2015	28%	
2016	30%	
2017	31%	

Table 5.8. New Orleans VP traceroutes traversing an MPLS tunnel percentages by year

The top two ASNs appear to be 22773 and 3356. ASN 22773 is associated with Cox Communications while ASN 3356 is Level 3. Cox Communications appears to account for 39% of all MPLS tunnels while Level 3 accounts for 31%.

Change in percent of MPLS tunnels for msy-us VP by year		
ASN	Count	
22773	12,852,353	
3356	10,133,079	
7018	1,913,998	
2914	1237,725	
6461	788,145	
1273	788,036	
2828	497,621	
6453	473,663	
286	353,295	
6830	289,370	

Table 5.9. Top 10 ASN appearances for New Orleans vantage point in 2013

In 2014, both Cox and Level 3 remain the top two, although they both experience a slight decrease. Cox accounts for 36% while Level 3 accounts for 25%. The other ASNs in the top 10 experience an increase.

Change in percent of MPLS tunnels for msy-us VP by year		
ASN	Count	
22773	52,676,772	
3356	36,907,034	
7018	10,236,839	
2914	6,989,558	
1273	4,988,280	
6461	4,108,392	
2828	3,931,801	
6453	2,840,653	
286	2,559,341	
7843	1,635,882	

Table 5.10. Top 10 ASN appearances for New Orleans vantage point in 2014

In 2015, Cox and Level 3 disappear entirely from the top 10. Cox and Level 3 decrease to .02% and .005%. ASN 2914, NTT America, takes over as the top ASN accounting for 16% and showing an increase in utilization of MPLS from previous years.

Change in percent of MPLS tunnels for msy-us VP by year		
ASN	Count	
2914	9,608,661	
7018	8,667,452	
6461	6,327,404	
1273	3,256,418	
6453	3,209,127	
286	3,071,563	
7843	1,805,772	
6830	1,456,503	
2764	1,125,284	
4538	941,831	

Table 5.11. Top 10 ASN appearances for New Orleans vantage point in 2015

Change in percent of MPLS tunnels by notable AS				
ASN	AS Name	2013	2014	2015
22773	Cox Com.	39%	36%	<1%
3356	LEVEL3	31%	25%	<1%
7018	AT&T	6%	7%	15%
2914	NTT	4%	5%	16%
6461	Zayo	2%	3%	11%

Table 5.12. New Orleans VP notable ASN changes

Table 5.12 demonstrates the change in percent of hops that traverse an MPLS tunnel by ASN. Cox and LEVEL3 clearly have made some change in their infrastructure which results in the almost complete elimination of traceroute hops traversing an MPLS tunnel in 2015. The other ASN do increase in percentage, however this is mostly due to the disappearance of the previous top 2 as their MPLS hop count does not increase significantly, and in the case of ASN 7018 the count actually decreases.

Based on Figure 5.3, it is obvious the median percentage of traces that traverse a MPLS tunnel are lower in 2015. However, the bottom outliers in 2014 appear to indicate that the trend towards the lower percentage may have started sometime within 2014. Therefore, we look to our database in an attempt to pinpoint a particular date range when the downward shift in traces that traverse a MPLS tunnel occurred.

```
traceanalysis=> select x.monitor, x.cycle_START, X.PERC from
(select distinct on (monitor, cycle_start) monitor, cycle_start,
(mpls_cnt/trace_cnt::float) as perc from files where
cycle_start >= '20140920' and cycle_start < '20141101'
and trace_cnt !=0 and monitor = 'msy-us')x order by x.cycle_start asc;
monitor | cycle_start | perc
msy-us | 2014-09-20 19:31:58 | 0.999985729779097
msy-us | 2014-09-22 07:14:58 |
                                             1
msy-us | 2014-09-23 19:26:54 | 0.589769278150069
msy-us | 2014-09-25 06:59:27 | 0.338230006677282
msy-us | 2014-09-26 18:20:17 | 0.336733579773213
msy-us | 2014-09-28 05:50:16 | 0.335882511045037
msy-us | 2014-09-29 17:58:03 | 0.337114149117613
msy-us | 2014-10-01 08:16:05 | 0.332174828857293
msy-us | 2014-10-02 21:52:42 | 0.337209302325581
msy-us | 2014-10-04 10:05:37 | 0.33030270994115
msy-us | 2014-10-05 22:36:09 | 0.332686069860324
msy-us | 2014-10-07 11:04:28 | 0.331644277564114
msy-us | 2014-10-08 22:25:36 | 0.318520914068966
msy-us | 2014-10-10 11:13:01 | 0.306986300216401
msy-us | 2014-10-12 00:48:48 | 0.304582460749204
msy-us | 2014-10-13 14:59:14 | 0.308271311513473
msy-us | 2014-10-15 04:12:43 | 0.306724660268287
msy-us | 2014-10-16 16:02:21 | 0.302791001426435
msy-us | 2014-10-18 04:42:57 | 0.303879837565439
msy-us | 2014-10-19 19:02:08 | 0.30209726443769
msy-us | 2014-10-22 22:48:41 | 0.268631849085131
msy-us | 2014-10-24 12:05:33 | 0.267448723624022
msy-us | 2014-10-26 00:38:14 | 0.269964464735368
msy-us | 2014-10-27 13:16:24 | 0.268901967351691
msy-us | 2014-10-29 02:31:41 | 0.270672258773399
msy-us | 2014-10-30 15:47:55 | 0.268869366875201
(26 rows)
```

Figure 5.4. Downward shift by percentage of MPLS by cycle for the msy-us VP

Figure 5.4 displays an approximate date and time for the start of the downward trend. From 2014-09-20 through 2014-09-22 we can observe the normal, expected, percentage from all previous data in our dataset for the msy-us VP. However, from 2014-09-23 onwards we observe a steady decrease in percentage. Based on our previous observations in Table 5.12, we believe that September 23rd, 2014 is the approximate inflection point for when ASN 22773 (Cox) and 3356 (LEVEL3) made significant changes to their infrastructure, affecting MPLS routers within.

5.2.3 Sao Paulo, Brazil, VP

Figure 5.5 demonstrates a noticeable upward trend in the percentage of traceroutes that traverse a MPLS tunnel over each year. This is a different trend from what was observed in figures 5.3 and 5.2. Table 5.13 also indicates that there is a clear percentage increase in traceroutes traversing MPLS tunnels, starting from 45% up to 93%. Inspecting which ASN is responsible for the MPLS tunnels is required to attempt an explanation.



Figure 5.5. Sao Paulo, Brazil

Table 5.13.	Sao Paulo VI	^D traceroutes	traversing a	MPLS tunne	el percentages
by year					

Change in percent of MPLS tunnels for Sao-br VP by year		
Year	Percentage	
2013	45%	
2014	52%	
2015	69%	
2016	89%	
2017	93%	

Table 5.14 demonstrates the continued trend of LEVEL3 (ASN 3356) accounting for less of the traceroutes that traverse a MPLS tunnel, starting in 2015. However, what is different about the Sao Paulo VP is that another telecom, ALGAR, increases in percentage to offset the decline of LEVEL3. The other ASN remain around the same hops that traverse a MPLS tunnel count and percentage.

Change in percent of MPLS tunnels by notable AS				
ASN	AS Name	2014 Percentage	2015 Percentage	
16735	ALGAR	27%	54%	
3356	LEVEL3	20%	<1%	
7018	AT&T	16%	10%	
2914	NTT	9%	15%	
6453	TATA Com.	4%	2%	

Table 5.14. Sao Paulo VP notable ASN changes

5.2.4 Pretoria, South Africa, VP

Based on Figure 5.6, it is obvious the median percentage of traces that traverse a MPLS tunnel are lower in 2015 compared to 2013 and 2014. However, similar to Section 5.2.2, the bottom quartile in 2014 indicates that the trend towards the lower percentage may have started sometime within 2014. Therefore, we look to our database in an attempt to pinpoint a particular date range when the downward shift in traces that traverse a MPLS tunnel occurred for Pretoria, South Africa (pry-za).



Figure 5.6. Pretoria, South Africa

Table 5.15. Pretoria VP traceroutes traversing a MPLS tunnel percentages by year _____

Change in percent of MPLS tunnels for pry-za VP by year			
Year	Percentage		
2013	55%		
2014	50%		
2015	33%		
2016	30%		
2017	30%		
2014 2015 2016 2017	50% 33% 30% 30%		

Change in percent of MPLS tunnels by notable AS				
ASN	AS Name 2014 Percentage 2015 Percentage			
3356	LEVEL3	56%	<1%	
7018	AT&T	11%	18%	
2914	NTT	10%	20%	
6461	Zayo	3%	9%	
174	Cogent Comm.	2%	11%	

Table 5.16. Pretoria, South Africa, VP notable ASN changes

traceanalysis=> select x.monitor, x.cycle_START, X.PERC from (select distinct on (monitor, cycle_start) monitor, cycle_start, (mpls_cnt/trace_cnt::float) as perc from files where cycle_start >= '20140801' and cycle_start < '20141001' and trace_cnt !=0 and monitor = 'pry-za')x order by x.cycle_start asc; monitor | cycle_start perc pry-za | 2014-08-01 16:07:26 | 0.581443063946199 pry-za | 2014-08-03 05:10:10 | 0.594077115900588 pry-za | 2014-08-04 18:18:06 | 0.60913290915197 pry-za | 2014-08-06 19:40:48 | 0.583545418950666 pry-za | 2014-08-08 10:39:16 | 0.584125526215078 pry-za | 2014-08-09 23:46:59 | 0.584190582304175 pry-za | 2014-08-11 12:32:49 | 0.584744276350406 pry-za | 2014-08-13 00:33:15 | 0.560803569320838 pry-za | 2014-08-14 15:03:05 | 0.551814683401296 pry-za | 2014-08-16 05:59:16 | 0.563573282747604 pry-za | 2014-09-11 19:47:07 | 0.330208079528288 pry-za | 2014-09-13 11:12:13 | 0.367767934658195 pry-za | 2014-09-15 01:20:51 | 0.341819312827572 pry-za | 2014-09-16 15:37:53 | 0.329869996929061 pry-za | 2014-09-17 18:53:42 | 0.32658971409793 pry-za | 2014-09-19 07:50:19 | 0.332321152856408 pry-za | 2014-09-20 19:31:58 | 0.335214773648257 pry-za | 2014-09-22 07:14:58 | 0.332586260258914 pry-za | 2014-09-23 19:26:54 | 0.325013320021312 pry-za | 2014-09-25 06:59:27 | 0.324354741896759 pry-za | 2014-09-26 18:20:17 | 0.345762575128071 pry-za | 2014-09-28 05:50:16 | 0.33413245430027 pry-za | 2014-09-29 17:58:03 | 0.329333871076158 (23 rows)

Figure 5.7. Downward shift by percentage of MPLS by cycle for the pry-za VP

Figure 5.7 displays an approximate date and time for the start of the downward trend. From

2014-08-01 through 2014-08-16 we can observe the normal, expected, percentage from all previous data in our dataset for the pry-za VP. However, from 2014-09-11 onwards we observe a decrease in percentage that remains constant from that point onwards in our dataset. Based on our previous observations in Table 5.16, we believe that 2014-09-11 is the approximate inflection point for when ASN 3356 (LEVEL3) made significant changes to their infrastructure, affecting MPLS routers within. This trend for ASN 3356 appears common, even among VP that are experiencing an increase in the percentage of traceroutes traversing a MPLS tunnel.

5.2.5 Singapore VP

Figure 5.8 and Table 5.17 demonstrate a steady increase in percentage of traceroutes that traverse a MPLS tunnel over each year. As a comparison, Figure 5.5 shows a dramatic increase from a starting point of 45% to 93%, an almost 50% total increase over 5 years. The Singapore VP starts at 61%, displaying a more modest 5 year total increase of 25%. As with the other VPs, inspecting the different ASN contributions towards the total MPLS tunnels provides insight into these changes.



Figure 5.8. Singapore

Table 5.17. Singapore VP traceroutes traversing a MPLS tunnel percentages by year

Change in percent of MPLS tunnels for sin-sg VP by year			
Year	Percentage		
2013	61%		
2014	67%		
2015	73%		
2016	84%		
2017	86%		

Table 5.18. Singapore VP notable ASN changes

				-
Change in percent of MPLS tunnels by notable AS				
ASN	AS Name	2013	2015	
		Percentage	Percentage	Percentage
2914	NTT	42%	48%	57%
3356	LEVEL3	21%	11%	9%
7018	AT&T	9%	9%	7%
6453	TATA	7%	6%	4%
6461	Zayo	3%	1%	2%

Similar to the other VPs, we can observe two trends across the various ASNs in Table 5.18. The first trend is that LEVEL3 declines in the percentage of MPLS tunnels seen by the Singapore VP, but not in as extreme of a way as other studied VPs. The second trend is that NTT increases in its share of the percentage over the 5 years. Sao Paulo, Pretoria, New Orleans, and Dublin ASN comparison tables all demonstrate the same trend of an increase in NTT MPLS tunnels. However NTT is not the top contributor of MPLS tunnels for these VP, like it is in Singapore. Therefore, NTT appears to be the ASN mainly responsible for the 25% increase in MPLS tunnels over the past 5 years for the Singapore VP.

5.2.6 Hamilton, New Zealand, VP

Figure 5.9 demonstrates an upward trend in the percentage of traceroutes that traverse a MPLS tunnel over each year, with an odd spike and many outliers near 50% in 2014. This is a slightly different trend from what was observed in figures 5.8, where the increase was consistent and continuous through the 5 years. Table 5.20 indicates that the total increase over 5 years, excluding the anomalous spike in 2014, is 10%. Inspecting which ASNs are most responsible for the MPLS tunnels is required to attempt an explanation.



Figure 5.9. Hamilton, New Zealand

Table 5.19. Hamilton, New Zealand, VP traceroutes traversing a MPLS tunnel percentages by year

Change in percent of MPLS tunnels for hlz-nz VP by year			
Year	Percentage		
2013	62%		
2014	78%		
2015	70%		
2016	71%		
2017	72%		

Table 5.20. Hamilton, New Zealand, VP notable ASN changes

Change in percent of MPLS tunnels by notable AS					
ASN	AS Name	2013	2015		
		Percentage	Percentage	Percentage	
4826	Vocus	63%	54%	<1%	
3356	LEVEL3	10%	12%	<1%	
7018	AT&T	6%	6%	9%	
2914	NTT	1%	1%	20%	
7575	AARNet	<1%	<1%	17%	
6461	Zayo	<1%	<1%	30%	

Table 5.20 demonstrates large changes between 2013 and 2015 for individual ASN. Hamilton, New Zealand's VP continues the trend of a sudden decrease in percentage of MPLS tunnels traversed for LEVEL3 in 2015. Hamilton also has a regional ASN, Vocus, who accounts for more than half of the MPLS seen in 2013 and 2014, but then suddenly drops to below 1% in 2015. Meanwhile, NTT makes a large increase in 2015. With the exception of Vocus, these trends seem consistent with what has been observed from other VPs, however the Hamilton VP seems to have especially exaggerated changes on a per ASN basis.

5.3 Observed MPLS Change from LEVEL3

Of all the ASNs observed during the analysis by VPs, LEVEL3 demonstrated the most consistent and dramatic change. LEVEL3 across all studied VPs had their percentage of MPLS tunnels drop near 2015. The cause of this one specific ASN changing within the data could perhaps be explained by LEVEL3 implementing changes within their infrastructure. LEVEL3 could be moving away from MPLS as a networking technology. LEVEL3 also could be no longer advertising the MPLS routers and instead using them in an opaque or invisible mode. The motivation for doing so could be due to security or competitive advantage reasons.

We performed analysis of LEVEL3 before and after the observed LEVEL3 (ASN 3356) event in an attempt to determine other alternative explanations for why there was a sudden drop in MPLS. For example, an explanation could be that LEVEL3 is no longer a major transit network or that there was a significant routing change. Based on the shifts observed in Figures 5.4 and 5.7, we picked a cycle before the event, 2014-08-01, and a cycle after the event had manifested, 2014-10-30. We looked for the total amount of traceroutes in each cycle, the total amount of traceroutes with a hop that traversed ASN 3356, and the count for those hops belonging to ASN 3356 that contained MPLS. Table 5.21 summarizes our findings that the percentage of traceroutes through ASN 3356 did not significantly change, while the percentage of traceroutes with hops that have IP addresses belonging to ASN 3356 with MPLS. This signifies that it is more likely some change in LEVEL3's infrastructure.

LEVEL3 Cycle Comparison					
Cycle	cle Traceroutes ASN ASN 3356 & MPLS ASN MPLS				
		3356		3356 %	%
2014-08-01	10,254,907	2,111,141	1,593,134	21%	75%
2014-10-30	10,588,199	2,041,599	981	19%	.04%

Table 5.21. LEVEL3 demonstrated no noticeable change in frequency of appearance in traceroutes, but significant change in MPLS

We reached out to LEVEL3 on 2018-2-27 for possible validation of our observations and any other information that they would be willing to share. We have not received a response at the time of this writing.

5.4 MPLS Correlation with Artifacts

A question worth attempting to answer is if there is any correlation between MPLS and these suspicious traceroutes. To attempt to answer this question, we looked in our artifact database for the ASN that appeared most frequently near the top for percentage of traceroutes that traversed a MPLS tunnel. Table 5.22 shows that the percentage of the total suspicious traceroutes appears to be low for some of the most common ASN in Section 5.1.

Suspicious Count and hop Percentage by Notable MPLS ASN				
ASN	AS Name	Suspicious Count	Percentage	
All	-	20,390,676	100%	
3356	LEVEL3	309,193	1.5%	
2914	NTT	60,530	.3%	
16735	ALGAR	44,745	.2%	
7018	AT&T	18,666	.1%	

Table 5.22. All VP notable MPLS ASN Suspicious traceroute counts

Another way to look at this data is to query the suspicious traceroutes and identify what percentage indicated the presence of MPLS. 253,153 of the 901,478 total suspicious traceroutes identified in our database contained at least one hop with ICMP extension, indicating the presence of a MPLS tunnel. This indicates that about 28% of our suspicious traceroutes contained detectable MPLS. Of the 20,390,676 hops and responses in the suspicious traceroute table, only 4,906,474 contained MPLS. This is an even lower percentage of 24%. This does not provide an indication of a correlation between MPLS and these suspicious traceroutes.

CHAPTER 6: Conclusion and Future Work

If nothing else, this work demonstrates that MPLS is still prevalent in today's Internet, although its utilization or advertisement of utilization can vary on an AS over time. LEVEL3 has clearly changed how their internal networks utilize MPLS, which drastically changed in approximately 2015. One theory to explain this change could be that LEVEL3 and other large ASs are moving away from the use of MPLS towards a new, better networking technology. These backend networking technology could also be a shift towards an entire IPv6 infrastructure within the AS. Another theory could be that the ASs that show a decrease in the percentage of traceroutes that traverse a MPLS tunnel are simply not advertising the MPLS enabled routers anymore. This could be explained by the security concerns that a motivated adversary may create a denial of service on an ingress MPLS router, causing degradation of the large AS's network. These large ASs can use internal tools designed specifically for diagnostics of MPLS router. Of course, this would harm researchers abilities to study MPLS prevalence on the Internet, but to a network operator this would be a lower priority.

We also identified a number of traceroute artifacts using our original classifications and algorithms. We were able to discover some well known traceroute gags within our 5 years of CAIDA traceroute data, such as the "Star Wars Traceroute" and "Bad.Horse". We also were able to discover some previously unknown anomalous traceroutes, such as our Comstar and Befree examples. When inspecting suspicious traceroutes that demonstrated both sequential IPIDs and many hop IP addresses in the same /24, we discovered that some of these traceroutes also exhibited the characteristic of having multiple responses to single probe. Finally, although many of the suspicious traceroutes that we studied contained MPLS within them, we can not claim that MPLS utilization is correlated with the artifacts that we observe. As one datapoint, LEVEL3, the largest contributor to MPLS tunnels, only accounted for 3% of suspicious traceroutes while the largest contributor at 14%, PROXAD, did not appear as a top contributor of MPLS. Also, only 28% of our suspicious traceroutes traversed an MPLS tunnel, which is under the average percentage of MPLS tunnels traversed

for all traceroutes.

A motivation for studying these anomalous traceroutes was their potential impact on topology graphs. Clearly, some of the anomalous traceroutes discovered are not real topologies, specifically the traceroute gags. Some of the other examples, such as the Befree instance, are likely some kind of misconfiguration. These could potentially lead to small false inferences, leading to a non-perfect network topology. However, the anomalous traceroutes do not make up a significant fraction of the entire traceroute dataset. On a macroscopic level, it does not appear that researchers would need to be significantly concerned about these artifacts in relation to the accuracy of topology graphs.

6.1 **Primary Contributions and Takeaways**

- Discovered from 2013-2017, the average detected MPLS tunnels traversed over all VPs hovers between 33% and 46%, while varying widely between VPs.
- Identified Level 3 as a large AS who stopped utilizing or advertising MPLS in September of 2014 across multiple VPs.
- MPLS as a networking technology in large ISPs is still being used, but there are alternatives that may explain the reduction in overall percentages. Additionally, there may be security concerns incentivizing large ISPs to disregard use of RFC 4950 [9] and instead not advertise that they are even using the technology.
- Demonstrated that MPLS does not appear to be the likely cause of the majority of these anomalous traceroutes.
- PROXAD ("Free" ISP) responsible for the most suspicious traceroute results.
- Discovered a "wild", previously undiscovered, Star Wars traceroute gag within our suspicious traceroute.
- Identified traceroutes with many responses to a single traceroute probe, another anomalous traceroute characteristic for researchers to look for in the future.

6.2 Future Work

We use this section to describe research that we have left for the future due to lack of time. Future researchers should consider working on the following areas to expand upon work done in this thesis.

6.2.1 MPLS Research

New methods for discovering MPLS tunnels have been published since starting this research, which should certainly be used to study the utilization of MPLS as a technology [51]. The importance of continuing to study MPLSs prevalence in todays internet is evident and summarized nicely by this quote,

For years, Internet topology research has been conducted through active measurement. For instance, Caida builds router level topologies on top of IP level traces obtained with traceroute. The resulting graphs contain a significant amount of nodes with a very large degree, often exceeding the actual number of interfaces of a router. Although this property may result from inaccurate alias resolution, we believe that opaque MPLS clouds made of invisible tunnels are the main cause. Using Layer-2 technologies such as MPLS, routers can be configured to hide internal IP hops from traceroute. Consequently, an entry point of an MPLS network appears as the neighbor of all exit points and the whole Layer-3 network turns into a dense mesh of high degree nodes. [51]

These researchers have developed methods to assist in finding these opaque MPLS tunnels. These new methods should be used to study LEVEL3 and other ASNs to discover if the methods we used for detecting MPLS missed parts of their MPLS infrastructure. One explanation for our observations is that the ASN changed how they configure their MPLS routers to respond to traceroute probes.

Future researchers should also reach out to the large ASN network operators for possible insights into the change in MPLS. We reached out to LEVEL3 ourselves, but have not heard back as of this writing. Contacting other ASNs, such as NTT who demonstrated an increase in percentage of MPLS, could also be useful in providing insight into what caused these changes.

Finally, we were only able to look at 5 vantage points in-depth regarding the change in their observed MPLS tunnels. In the Appendix are boxplots of other VPs that demonstrate varied changes in the percentage of MPLS over the 5 years. Future researchers may want to consider these VPs if searching for other regional ASN that may have made changes to their infrastructure regarding MPLS.

6.2.2 Suspicious Traceroutes

In Section 4, we analyze various suspicious traceroutes that we detected. These anomalous traceroutes were detected by looking for two main characteristics. Figure 4.11 demonstrates a third characteristic, multiple responses to a single probe, that we believe future researchers should look for when studying suspicious traceroutes in large-scale traceroute data. The majority of our 73 traceroutes that demonstrated the original characteristics we were looking for involving multiple hops in the same /24 and sequential IPIDs also demonstrated the new characteristic. It would be worth going through our 5 years of CAIDA traceroute data to look for traceroutes that only demonstrate the multiple responses to a single probe, or to look through future years worth of data.

In addition, the historical DNS dataset is extremely useful for searching for well-known traceroute gags or patterns. Future researchers interested in specific traceroute gags should compile DNS strings to search for within this dataset. This is particularly useful for gags that are no longer active or for identifying the approximate time that a gag was active. The DNS dataset is especially useful for traceroute gags that manifest to probes to multiple IP addresses within a /24 space, rather than being responsive to only to a single IP address.

6.2.3 Future Rapid Analysis

CAIDA is continuously collecting IPv4 traceroute data that can be analyzed by future researchers. We believe that the techniques and methods used in this research could be streamlined to perhaps identify both trends in MPLS and suspicious traceroutes, perhaps on a monthly basis. By automating this analysis, the quick turn around could allow a researcher to study these trends and artifacts in the present time. This can be especially important for suspicious traceroutes, as they often cease to exist in a short amount of time. For example, the Star Wars Traceroute ceased to exist after approximately 14 months. Being able to identify these anomalous traceroutes quickly could be very useful. Observing MPLS trends quickly can also be useful, as contacting network operators of large ISP may be more successful if they more recently made a change. In addition, effects on the topology may be more easy to observe for researchers studying changes in real time.

6.2.4 IPv6

IPv4 is still widely used on todays Internet and will likely remain popular for years to come. However, in 2012 IPv6 was launched and since has been adopted steadily throughout the Internet, especially by large ISP who many be concerned about IPv4 address exhaustion. There are issues to take into account regarding the differences between studying IPv4 and IPv6. One issue is IPv4 address space is 2^{32} while IPv6 is 2^{128} . This makes it much more resource intensive and time consuming to study the entirety of an unfiltered IPv6 dataset. CAIDA has an IPv6 routed topology infrastructure, as part of their Ark platform, that probes announced IPv6 prefixes that are /48 or shorter [52]. Therefore this dataset is much more manageable to study than the entirety of the IPv6 space.

There are techniques used in this thesis involving IPv4 that could work in research involving IPv6. Searching for multiple hops or responses in the same /48 or /64 within a single traceroute would be viable and could lead to the discovery of IPv6 artifacts similar to the ones we discovered in IPv4. However, identifying anomalous traceroutes with sequential IPIDs is not possible because IPIDs do not exist in IPv6.

6.2.5 Database Techniques for Traceroute Data

Although we made good use of a structured database for our 5 years of CAIDA traceroute data, in some cases we approached the limits of what it could handle. More complex queries across the entire dataset could take anywhere from hours to days to complete across billions of rows. If a researcher required a database containing all of the years worth of traceroute data, unfiltered, this approach would not suffice. Therefore, a database utilizing a technology such as RocksDB, a high performance key-value store, may be a better solution. CAIDA has such a system for querying traceroute data, given the title Henya. [53] While Henya is currently functional, at the time of this writing it does not have all VPs traceroute data available, with some VPs limited to roughly a year of data. However, Henya is still under development and is expected to eventually have all VPs traceroute data added to it for querying. We strongly recommend future researchers consider using Henya when studying historical traceroute data.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX: Vantage Point Boxplots

The following figures represent VPs that we did not have time to examine in depth.



Figure A.1. Ames, Iowa, United States



Figure A.2. Barcelona, Spain



Figure A.3. Stockholm, Sweden



Figure A.4. Cheongju, South Korea



Figure A.5. Dakar, Senegal



Figure A.6. Geneva, Switzerland



Figure A.7. Helsinki, Finland



Figure A.8. Leipzig, Germany



Figure A.9. Monterey, California, United States



Figure A.10. Narita, Japan



Figure A.11. Oakland, California, United States



Figure A.12. San Diego, California, United States



Figure A.13. Hong Kong



Figure A.14. Sydney, Australia



Figure A.15. Toronto, Canada

- Summary of the Amazon S3 service disruption in the Northern Virginia (US-EAST-1) Region. (2017, March). Amazon. [Online]. Available: https://web.archive.org/ web/20170323115743/https://aws.amazon.com/message/41926/
- [2] Akamai: Our Customers. (2018). Akamai. [Online]. Available: https://www.akamai. com/us/en/our-customers.jsp. Accessed Feb. 1, 2018.
- [3] E. Nygren, R. K. Sitaraman, and J. Sun, "The akamai network: A platform for highperformance internet applications," *SIGOPS Oper. Syst. Rev.*, vol. 44, no. 3, pp. 2– 19, Aug. 2010. Available: http://doi.acm.org/10.1145/1842733.1842736
- [4] N. Phua, Weiyou, "Detection of active topology probing deception," 2015. Available: https://calhoun.nps.edu/handle/10945/47313
- [5] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu, "Detecting prefix hijackings in the internet with argus," in *Proceedings of the 2012 Internet Measurement Conference* (IMC '12). New York, NY, USA: ACM, 2012, pp. 15–28. Available: http://doi.acm. org/10.1145/2398776.2398779
- [6] R. Beverly, "Yarrp'ing the internet: Randomized high-speed active topology discovery," in ACM on Internet Measurement Conference, 2016, pp. 413–420, 10.1145/2987443.2987479.
- [7] K. Claffy, Y. Hyun, K. Keys, M. Fomenkov, and D. Krioukov, "Internet mapping: From art to science," in *Conference For Homeland Security*, 2009, pp. 205–211.
- [8] Q. Lone, M. Luckie, M. Korczyński, and M. van Eeten, "Using loops observed in traceroute to infer the ability to spoof," in *Passive and Active Measurement*. Springer International Publishing, 2017, pp. 229–241.
- [9] R. Bonica, D. Gan, D. Tappan, and C. Pignataro. (2007). RFC 4950: ICMP Extensions for Multiprotocol Label Switching. Internet Engineering Task Force. [Online]. Available: https://tools.ietf.org/html/rfc4950. Accessed May 17th, 2017.
- [10] V. Jacobson. (1989). *Traceroute*. [Online]. Available: ftp://ftp.ee.lbl.gov/traceroute. tar.gz. Accessed Dec. 12, 2016.
- [11] J. Postel. (1981). RFC 792: Internet Control Message Protocol. Internet Engineering Task Force. [Online]. Available: https://tools.ietf.org/html/rfc792. Accessed May 17th, 2017.

- [12] M. Luckie, Y. Hyun, and B. Huffaker, "Traceroute probe method and forward ip path inference," *In Proc. 8th ACM SIGCOMM Conference on Internet Measurement*, pp. 311–324, 2008.
- [13] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, and R. Teixeria, "Avoiding traceroute anomalies with paris traceroute," *In Proc. 6th ACM SIGCOMM Conference on Internet Measurement*, pp. 153–158, 2006.
- [14] V. Giotsas, C. Dietzel, G. Smaragdakis, A. Feldmann, A. Berger, and E. Aben, "Detecting peering infrastructure outages in the wild," in *ACM SIGCOMM*, Aug 2017.
- [15] M. Luckie and R. Beverly, "The impact of router outages on the as-level internet," in *ACM SIGCOMM*, Aug 2017, pp. 488–501.
- [16] P. Ferguson. (2000). RFC 2827: Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing. Internet Engineering Task Force. [Online]. Available: https://www.ietf.org/rfc/rfc2827.txt. Accessed May 10th, 2017.
- [17] F. Baker. (2004). RFC 3704: Ingress filtering for multi-homed networks. Internet Engineering Task Force. [Online]. Available: https://www.rfc-editor.org/rfc/rfc3704. txt. Accessed May 16th, 2017.
- [18] Y. Shavitt and U. Weinsberg, "Quantifying the importance of vantage points distribution in internet topology measurements," in *IEEE INFOCOM 2009*, April 2009, pp. 792–800.
- [19] P. Barford, A. Bestavros, J. Byers, and M. Crovella, "On the marginal utility of network topology measurements," in *Proceedings of the 1st ACM SIGCOMM Workshop* on Internet Measurement (IMW '01). New York, NY, USA: ACM, 2001, pp. 5–17. Available: http://doi.acm.org/10.1145/505202.505204
- [20] M. Luckie, "Resilience of deployed TCP to blind attacks," Center for Applied Internet Data Analysis (CAIDA), Tech. Rep., Oct 2017.
- [21] M. Luckie, "Scamper: a scalable and extensible packet prober for active measurement of the internet," *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement (IMC)*, pp. p. 239–245, November 2010. Available: http: //www.caida.org/tools/measurement/scamper/
- [22] CAIDA Data Overview of datasets, monitors, and reports. (2018). CAIDA. [Online]. Available: http://www.caida.org/data/overview/. Accessed Feb. 1, 2018.

- [23] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani, "Plane: An information plane for distributed services," in *Proceedings of the 7th Symposium on Operating Systems Design and Implementation* (OSDI '06). Berkeley, CA, USA: USENIX Association, 2006, pp. 367–380. Available: http://dl.acm.org/citation.cfm?id=1298455.1298490
- [24] D. Meyer. (2017, March). RouteViews. [Online]. Available: http://www.routeviews. org
- [25] Y. Shavitt and E. Shir, "Dimes: Let the internet measure itself," SIGCOMM Comput. Commun. Rev., vol. 35, no. 5, pp. 71–74, Oct. 2005. Available: http://doi.acm.org/ 10.1145/1096536.1096546
- [26] The IPv4 Routed /24 Topology Dataset. (2016, December). CAIDA. [Online]. Available: http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml
- [27] V. Bajpai, S. J. Eravuchira, and J. Schönwälder, "Lessons learned from using the ripe atlas platform for measurement research," *SIGCOMM Comput. Commun. Rev.*, vol. 45, no. 3, pp. 35–42, July 2015. Available: http://doi.acm.org/10.1145/2805789. 2805796
- [28] Raspberry Pi. (2018). Raspberry Pi Foundation. [Online]. Available: https://www.raspberrypi.org/. Accessed Mar. 1, 2018.
- [29] Archipelago monitor locations. (2018, January). CAIDA. [Online]. Available: http: //www.caida.org/projects/ark/locations/
- [30] S. Trassare, R. Beverly, and D. Alderson, "A technique for network topology deception," *MILCOM*, pp. 13–36, 2013.
- [31] R. Werber. (2013, February). Star Wars Traceroute. [Online]. Available: http: //beaglenetworks.net/post/42707829171/star-wars-traceroute
- [32] RKBicknell. (2015, December). Christmas Carol Traceroute. [Online]. Available: http://www.reddit.com/r/networking/comments/2qb86s/
- [33] W. Blog. (2017, March). The Pirate Bay North Korean hosting? No, it's fake. (P2). [Online]. Available: https://web.archive.org/web/20160816033531/https://rdns.im/ the-pirate-bay-north-korean-hosting-no-its-fake-p2
- [34] J. Renken. (2018, January). Bad.Horse. [Online]. Available: http://bad.horse
- [35] B. Donnet, M. Luckie, P. Mérindol, and J. Pansiot, "Revealing MPLS tunnels obscured from traceroute," ACM SIGCOMM Computer Communication Review (CCR), vol. 42, no. 2, pp. 87–93, Apr 2012.

- [36] E. Rosen, A. Viswanathan, and R. Callon. (2001). RFC 3031: Multiprotocol Label Switching Architecture. Internet Engineering Task Force. [Online]. Available: http: //www.rfc-editor.org/rfc/rfc3031.txt. Accessed May 17th, 2017.
- [37] MPLS FAQ for Beginners. (2016, December). Cisco. [Online]. Available: http:// www.cisco.com/c/en/us/support/docs/multiprotocol-label-switching-mpls/mpls/ 4649-mpls-faq-4649.html
- [38] Archipelago (Ark) Measurement infrastructure. (2018, February). CAIDA. [Online]. Available: http://www.caida.org/projects/ark/
- [39] J. Touch. (2013). RFC 6864: Updated specifications of the IPv4 ID field. Internet Engineering Task Force. [Online]. Available: https://tools.ietf.org/html/rfc6864. Accessed May 15th, 2017.
- [40] A. Bender, R. Sherwood, and N. Spring, "Fixing ally's growing pains with velocity modeling," in *Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement* (IMC '08). New York, NY, USA: ACM, 2008, pp. 337–342. Available: http://doi.acm.org/10.1145/1452520.1452560
- [41] Tools for CAIDA's Scamper Packet Prober. (2018). CMAND. [Online]. Available: https://github.com/cmand/scamper. Accessed Feb. 1, 2018.
- [42] IPv4 Routed /24 DNS Names Dataset. (2018, February). CAIDA. [Online]. Available: https://www.caida.org/data/active/ipv4_dnsnames_dataset.xml
- [43] Socket: Low-Level Networking Interface. Python Software Foundation. Wilmington, DE. [Online]. Available: https://docs.python.org/2/library/socket.html. Accessed Mar. 5, 2017.
- [44] *adns-python*. Google LLC. Mountain View, CA. [Online]. Available: https://code. google.com/archive/p/adns-python/. Accessed Feb. 20, 2017.
- [45] L. Gao, "On inferring autonomous system relationships in the internet," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 733–745, Dec 2001.
- [46] Pickle Python Object Serialization. Python Software Foundation. Wilmington, DE. [Online]. Available: https://docs.python.org/2.7/library/pickle.html. Accessed Mar. 5, 2017.
- [47] Maxmind. (2018, February). Maxmind GeoIP Service. [Online]. Available: http: //dev.maxmind.com/geoip/legacy/downloadable/
- [48] AS Rank. (2018). CAIDA. [Online]. Available: http://as-rank.caida.org/asns/3356. Accessed Feb. 1, 2018.
- [49] R. Miller. (2015, September). You should traceroute bad.horse right now. *The Verge*.
 [Online]. Available: https://www.theverge.com/tldr/2015/9/25/9398889/dr-horrible-traceroute-bad-horse
- [50] nwatson. (2016, December). traceroute bad.horse. *Hacker News*. [Online]. Available: https://news.ycombinator.com/item?id=13122389
- [51] Y. Vanaubel, P. Mérindol, J.-J. Pansiot, and B. Donnet, "Through the wormhole: Tracking invisible mpls tunnels," in *Proceedings of the 2017 Internet Measurement Conference* (IMC '17). New York, NY, USA: ACM, 2017, pp. 29–42. Available: http://doi.acm.org/10.1145/3131365.3131378
- [52] The IPv6 Topology Dataset. (2018). CAIDA. [Online]. Available: https://www.caida.org/data/active/ipv6_allpref_topology_dataset.xml. Accessed Mar. 28, 2018.
- [53] Henya. (2018). CAIDA. [Online]. Available: https://www.caida.org/tools/utilities/ henya/. Accessed Mar. 28, 2018.

THIS PAGE INTENTIONALLY LEFT BLANK

Initial Distribution List

- 1. Defense Technical Information Center Ft. Belvoir, Virginia
- 2. Dudley Knox Library Naval Postgraduate School Monterey, California