Shining Light on Dark Address Space[1]

Craig Labovitz, Arbor Networks, Merit Network, Inc.
Abha Ahuja, Arbor Networks, Merit Network, Inc.
Michael Bailey, Arbor Networks

November 13, 2001

[1] Some of this research was conducted under the auspices of the Internet Performance Measurement and Analysis project, supported by the National Science Foundation grant NCR 9710176.

ABSTRACT

In this report, we explore the degree to which commercial strategies, peering disputes, network failures, misconfiguration, and occasionally, malicious intent, lead to a partitioning of Internet topology. Specifically, we present a three-year study of the one-sided differences in Internet provider reachability. We focus on "dark address space", or the range of topology accessible from one provider, but unreachable via one or more competitor networks. We present active and passive measurements of these differences on time scales ranging from several seconds to multiple months. We show that more than five percent of the currently routed Internet address space lacks global connectivity. Our analysis shows that the majority of this partitioned Internet topology includes broadband customers and United States military networks. Finally, we explore "murky" address space, or the transient announcement of unallocated or reserved address space. We find that the majority of murky address space announcements represent misconfiguration, while a smaller number may represent intentional misuse and cooption of the Internet routing infrastructure.

1 INTRODUCTION

The rapid expansion in Internet infrastructure over the last several years has created a global network of unparallel proportions and promise. Even a low-cost local dial-in account provides the promise of ubiquitous and global network connectivity to the estimated hundred million (and growing) Internet end hosts.

Although Internet failures are commonplace, recent measurement and qualitative end-user experience shows that most outages are transient and persist for periods on the order of hours or even minutes as networks reroute or circuits are repaired [Labovitz:99]. Conventional wisdom holds that modulo these failures, the Internet topology represents a complete graph, or spanning tree of inter-connected networks [Faloutsos:99].

In this paper, we explore the degree to which this spanning tree connectivity model currently holds true for all nodes in the Internet. We base our discussion on three years of continuous passive and active measurements of geographically and topologically diverse core Internet routing tables. On time scales ranging from several seconds to multiple months, we identify the persistent onesided differences in Internet provider topology. We find that commercial strategies, peering disputes, network failures, misconfiguration, and occasionally, malicious intent, lead to a persistent partitioning of Internet topology. Specifically, recent measurements show that more than five percent of the default-free Internet address space includes contains "dark addresses," or networks which lack global interconnectivity.

A number of recent studies, including [RIPE:01,Malan:01] have reported tangential or qualitative evidence of the existence of dark address space. In particular, Uijterwaal et al. in [RIPE:01] identify significant variations in the number of prefixes maintained in multiple concurrent default-free routing tables based on static comparisons of European provider routing table snapshots. The authors found several thousand "dark prefixes," or oneside differences amongst all prefixes available from the monitored peers. In this paper, we focus on the specific range of partitioned topologies. As we discuss later in Section 2, variations in BGP prefix announcements (dark prefixes), may reflect differences in the quality of aggregation and not necessarily partitioned topology (dark address space).

Malan et. al in [Malan:01] analyzed the patterns of IP addresses in the backscatter logs of victims of wide-area DDoS attacks. Malan identified several thousand unique victim IP addresses that did not correspond to validly routed or allocated Internet address space. Malan postulated that these victims represented DDoS attack responses to prior attacks emanating from forged IP source addresses.

Other recent studies have also explored the evolution of the Internet routing table size and constitution. In [Houston:01,Alaettinogul:01], Houston and Alaettinoglu describe significant growth in more specific BGP announcements and argue this trend reflects growth in Internet multi-homing and inter-domain route tuning. Although both authors describe regional differences in routing table size, they attribute these differences primarily to variations in the quality of provider aggregation. Finally, other studies including [Caida:00] have analyzed static Internet routing table snapshots to identify differences in prefix and ASPath distributions. Our work explores a complimentary aspect of Internet topological analysis – the continuous long-term differences in diverse views of Internet topology.

Specifically, our major results include:

- More than five percent of Internet routed prefixes constitute dark address space
- The range of dark address space encompasses tens of millions of possible end hosts
- The majority (78%) of dark prefix blocks contain hosts that respond to active measurements
- Partitioned networks include cable modems and United States military networks
- The primary sources of dark address space include aggressive prefix length and IRR filtering, as well as misconfiguration.

The remainder of this report is organized as follows: Section 2 provides a description of our experimental methodology. In Section 3, we present the results and analysis of our experimental measurements. Finally, we conclude with a discussion of murky address space and misuse of the routing infrastructure.

2 METHODOLOGY

Past studies of Internet topology have typically focused on either active measurements or analysis of static routing tables from a small number of backbone vantage points. As we show in this paper and [Labovitz-Infocom:01], local measurements of Internet topology often prove unrepresentative of the Internet as a whole. In particular, we found Internet routing tables exhibit significant variation based on provider, and geographic and network topology. In this work, we base our measurements on three years of local and multihop eBGP peering sessions with tier-1 and tier-2 providers in the United States, Asia, and Europe. All providers configured the peering sessions with our probe machines in the same manner as their default-free customer sessions.

Our passive instrumentation included several “RouteViews” probe machines that maintained continuous state on the differences between all the monitored BGP sessions. In addition, the RouteViews probes logged all updates received from the peers to disk for subsequent post analysis. Once the system identified potential dark address blocks, we began active instrumentation of these prefixes. In particular, we used a collection of several topologically diverse proprietary active measurement hosts and publicly available looking glass servers [Kernen:01]. We used the publicly available looking glasses to verify the reachability via

traceroute and ping of hosts classified as dark address space from a broad range of providers. We used the proprietary measurement platforms to determine the host type and owner via whois, dns and nmap.

We measured a broad cross section of providers and timescales to provide a statistically valid lower bound on the degree of partitioned topology and Internet address space. Our system identified all long-lived differences amongst the monitored peering session tables. Specifically, the software automatically discovered prefixes which exhibited long-term persistence and for which all or a portion of the prefix's address space was never available in one or more of the other monitored tables. By portion of address space, we mean that the dark prefix did not match any exact, less or more specific prefixes in other concurrently monitored routing tables. We define "long-lived" prefix as one that persisted for more than 80% duration of the study period, or three months for the majority of our analysis. We focus on long-lived prefixes so as to distinguish truly partitioned addresses space from transient failures or misconfiguration.

Finally, we used logs from email spam traces and DoS backscatter analysis to identify murky prefixes. For each spam or DoS IP addresses we traced the availability of the corresponding covering routing table entry. We classify an IP address as "murky" if the corresponding routing table entry was available in routing table less five percent of the study period and if the entry was announced/withdrawn in the 12-hour period surrounding the spam or DDoS event.

3 RESULTS

In this section, we present the results from our three-year analysis of dark and murky Internet address space. We begin by describing measured differences in BGP routing tables and quantifying the size of partitioned address space. The next section explores some of the origins of these partitions. We conclude by presenting ongoing work in the exploration of intentional misuse of Internet routing table space.

3.1 Differences in BGP Routing Tables

As described earlier, Internet routing tables exhibit significant variation. In particular, provider filtering and peering policies have significant impact on the size of the default free routing table. In Figure 1, we provide a graph showing the routing table size in five-minute bins for 12 Internet providers on October 1, 2001. We observe that the routing tables differ by as much as 25,000 routes. Analysis of the data shows that the majority of these differences reflect the quality of aggregation. For example, in the below graph one provider announced the large 65.0.0.0/8 aggregate while other providers announced several dozen more specifics of this range, including 65.2.0.0/16 and 65.0.10/24. Still, as we describe below, some of these variances represent differences in reachability.

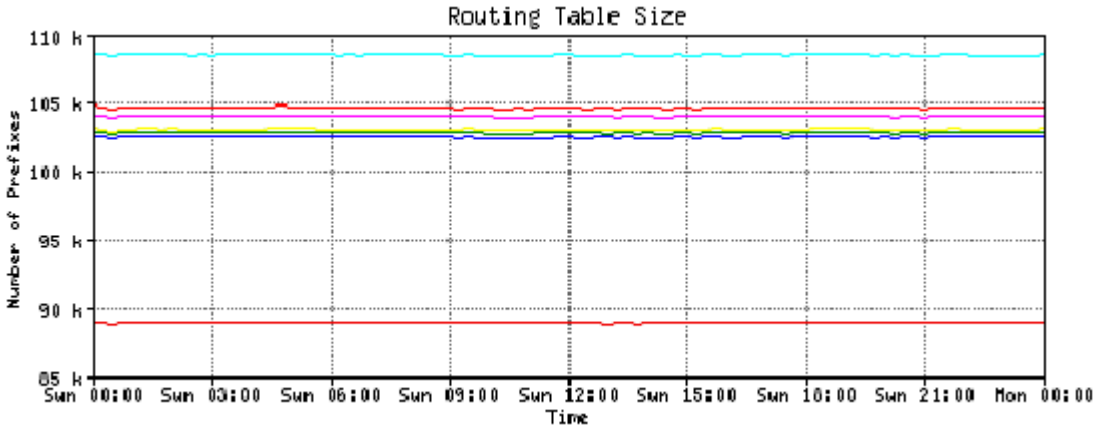

Figure 1 Routing Table Size (5 minute Bins)

We next explore the topological differences between routing tables. In figure 2, we graph the onesided differences in the number of unique autonomous systems. In other words, we show the ASes that were present in at least one provider's table and missing in at least one provider. All data is collated in five-minute bins. We note that measurements from a variety of sources including [Huston:01, RIPE:01] indicate some ten to fifteen thousand autonomous systems in current routing tables. In Figure 2, we see approximately 1500 of these ASes, or ten percent, lack global visibility.

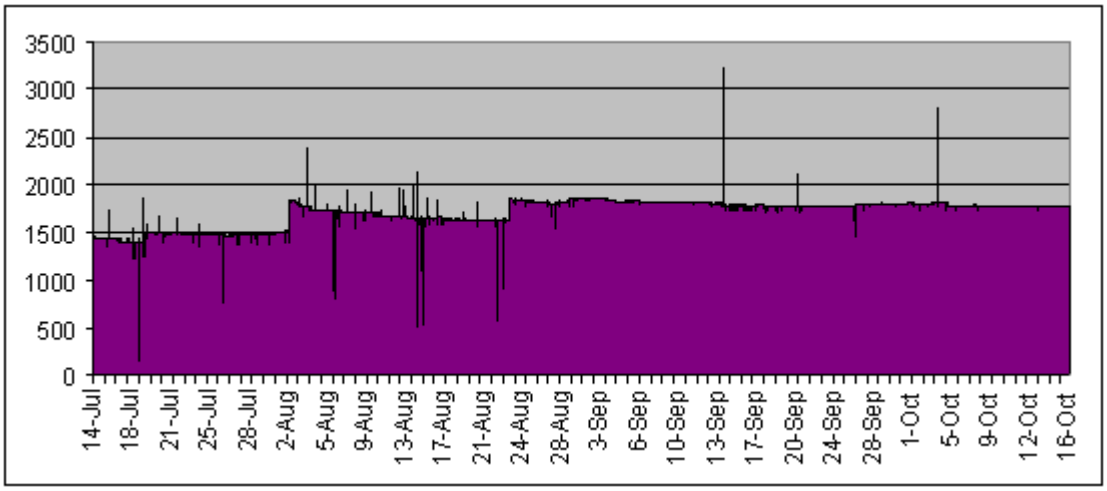

Figure 2 One-sided differences in the number of unique autonomous systems

Further analysis shows that the majority of these ASes appear solely as origin ASes. Only two to three hundred of non-globally visible ASes also appear in transit paths. These 1500 ASes lacking of global visibility represent either a high degree of coordination in aggregation between providers, or partitioned topology as a result of routing policy.

In figure 3, we measure the total number of onesided partitioned prefixes. As discussed earlier, we define a partitioned prefix as one present in at least one provider's table and for which no exact, less or more specific prefix exists in one or more other tables. We graph the number of dark addresses in union of 20 provider tables in five-minute bins over the course

two-month study period. We see that more than 5000 routes, or up to five percent of routed Internet address lack global visibility. We repeated the analysis on time periods ranging from several days to months and found the majority of these dark addresses were persistent. In other words, only several hundred represent transient events such as failures, or re-routing due to policy changes.

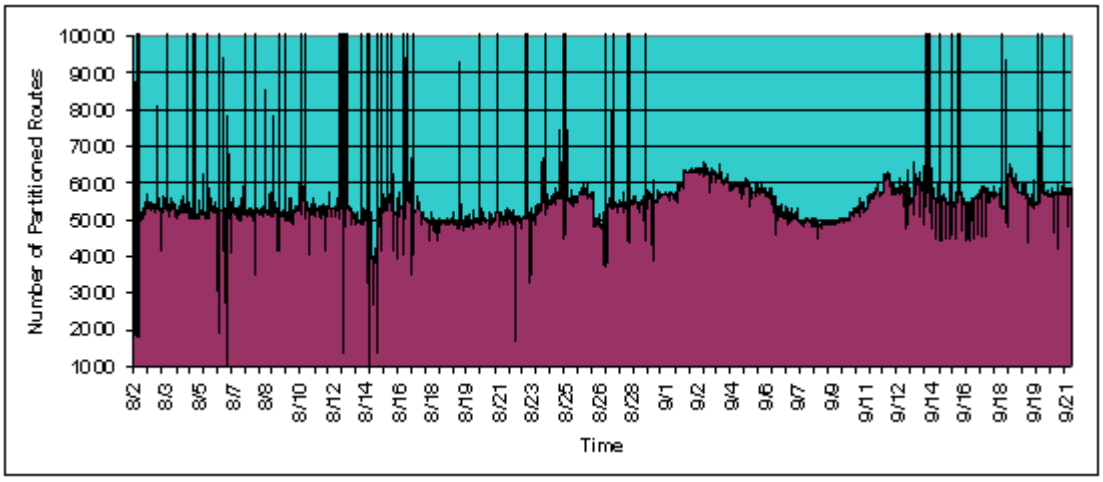

Figure 3 Total number of one-sided partitioned prefixes

3.2 Characterization of Dark Address Space

After identification of each dark prefix, we performed active measurements using publicly available look-glass servers. Approximately 78 percent of these dark addresses contained at least on IP address which responded to active measurement. We cannot distinguish whether the remaining 20 percent consisted of filtered (i.e. behind firewall) networks, or if these address space represented unallocated (i.e. did contain active hosts) space. Due to the intrusive nature of these probes, we randomly selected approximately 1% of the routes to measure.

In an effort characterize these dark address segments we used DNS inverse lookup, RIR, Internet Routing Registry information and NMAP. Thirty percent of the in these segments hosts had reverse DNS entries and 76 percent had covering entries in the Internet Routing Registry. Further analysis of the available data shows most of the segments were cable and ISDN modem pools as well as US military networks. The latter likely represents underutilized historical allocations. An ongoing area of research is the classification of the 24% of hosts that responded to active availability measurements, but had neither addressing, nor routing information.

3.3 Origins of Dark Address Space

In this subsection we turn our attention to the classification of the origins of dark address space. In effort to isolate and identify specific origins of dark regions, we compared each pair of providers. In Figure 4, we graph comparison of one such pair that represents two geographically close and topologically similar providers. Specifically, both providers represent North American regional networks and share many of the same commodity upstream transit providers. Figure 4 shows the number of prefixes available in the provider but lacking a corresponding less, exact or more specific entry in the other provider.

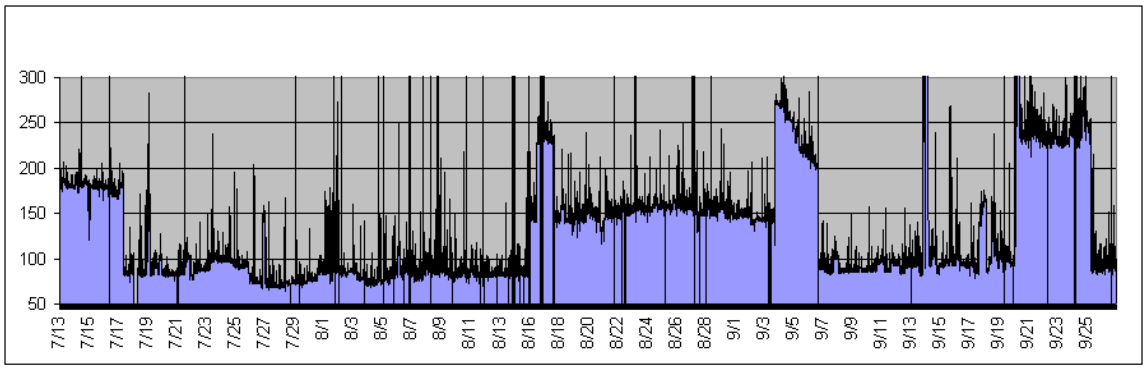

Figure 4 Missing prefixes between two geographically close and topologically similar providers

Discussions with both providers indicate that the plateau beginning August 16 represents a specific router misconfiguration. Analysis of remaining graph artifacts remains an area of ongoing research. In comparisons of provider pairs over a year period (data not shown) we identify a number of significant difference which correspond to peering disputes, including the January 2001 C&W and PSI dispute [NANOG:01]

Finally, we used analysis of provider comparisons, registry information, ISP surveys and discussions with providers to classify dark address space. We found the majority of dark address space reflects aggressive prefix length filters. For example, at least one tier-one provider filters /xxx. Some of largest onesided differences due to combination of prefix length filtering with use of IRR for access lists.

In addition, we identified follow as source dark address space:

- Test routes: One provider announced two statically configured test routes to each peer. Did not carry route in IGP
- IXP space. Policies decision whether to announce Internet exchange point, such as Mae-East or PAIX address blocks.
- Cisco configuration examples
- Misconfiguration: includes leaking of RFC-1918 space, default, and other private address ranges.

3.4 Misuse

In the previous sections we explored the scope of persistent, or long-lived differences in Internet address space. In this section, we explore the degree to which short-lived routing announcements indicate misuse of the routing infrastructure.

A growing number of articles have explored the theoretic possibility of Internet routing misuse, including intentionally "black holing" a target networks traffic [batz 1999].

In ongoing work, we are comparing short-lived BGP route announcements with spam DoS backscatter logs and mail logs from a diverse collection of large smtp servers. In preliminary results, we found that of the 40,000 thousand unique mail sources in the logs, approximately 30 of these entries correspond closely to brief BGP announcements

4 CONCLUSION

In this report the authors have attempted to answer several important questions regarding the increasing large and complex world of Internet routing. In particular:
Is the Internet actually a fully connected graph?
If it is not, to what degree is the Internet partitioned?
Do these partitioned networks represent actual disenfranchised hosts or are they underutilized allocations? Who is being disenfranchised?
What are the origins of these partitions?

The results of a three-year study of the one-sided differences in Internet provider reachability are presented. Our findings agree with others such as [RIPE:01] that the Internet is indeed partitioned and there exists “Dark Address Space”, or prefixes that are not reachable for one provider but that are available from other providers for long periods of time. Data is presented which shows that this space is very large, 5% of the total number of prefixes in the Internet or tens of millions of possible end hosts. These prefixes represent a large number of disenfranchised hosts; over 70% of the prefixes had hosts that responded to reachability tests. These hosts were mostly cable/ISDN pools as well as US military hosts. The origins of Dark Address Space are discussed and the common reasons including misconfigurations as well as policy decisions are reported. Other interesting causes included test routes as well as intentional misuse.

REFERENCES

[Zhao:01] X. Zhao, D. Pei, D. Massey, A. Mankin, S. Wu, L. Zhangm “An Analysis of BGP Multiple Origin AS (MOAS) Conflicts.

[Huston:01] G. Huston, “BGP Table Statistics,” IETF Plenary Presentation.

[Alaettinolgu:01] C. Alaettinolgu, “RIPE/RIS Project BGP Analysis: CIDR at Work,” NANOG 23 presentation, October 2001.

[Labovitz-FTCS:99] C. Labovitz, A. Ahuja, F. Jahanian, “Experimental Study of Internet Stability and Wide-Area Network Failures,” in the Proceedings of FTCS99, Madison, WI, June 22, 1999.

[Labovitz-Infocom:01] C. Labovitz, R. Wattenhofer, S. Venkatachary, A. Ahuja, "The Impact of Internet Policy and Topology on Delayed Routing Convergence", in the Proceedings of INFOCOM 2001, Anchorage, Alaska, April 26, 2000.

[Faloutsos:99] [Michalis Faloutsos, Petros Faloutsos, Christos Faloutsos, “On Power-Law Relationships of the Internet Topology,” ACM SIGCOMM'99.

CAIDA Routing Table Analysis, http://www.caida.org/analysis/routing/.

[Kernen:01] Public TRACEROUTE Server List, http://www.traceroute.org/

[RIPE:01] Henk Uijterwaal, Antony Antony, “Routing Information Service Status and Plans,” http://www.ripe.net/ris/

[NANOG-PSI:01] http://www.cctec.com/maillists/nanog/historical/0106/msg00313.html

[NANOG 2000] NANOG Mailing List Archive, July 2000, http://www.cctec.com/maillists/nanog/index.html

[Batz:1999] "batz", Security Issues Affecting Internet Transit Points and Backbone Providers, Blackhat Briefings, 1999. http://www.blackhat.com/html/bh-usa-99/bh3-speakers.html

[Rmalan:2001] Rob Malan, Farnam Jahanian, Jon Arnold, Matthew Smart, Paul Howell, Russell Dwarshius, Jeff Ogden, Jon Poland, Observations and Experiences Tracking Denial-Of-Service Attacks Across a Large Regional ISP, NANOG 22, May 20-22, 2001, Scottsdale, AZ

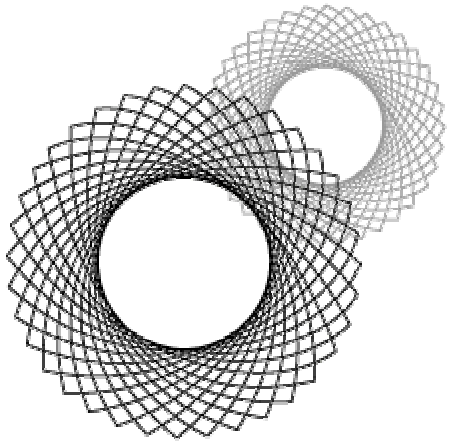

ABOUT ARBOR NETWORKS

Arbor Networks™ builds availability solutions for companies whose success depends on highly available and fully optimized networks.

Arbor Networks Peakflow DoS is a distributed, non-intrusive, scalable availability solution that detects, traces and recommends filters to counter availability threats, such as DoS attacks, improving network uptime, performance and security for large enterprises and service and hosting providers.

Arbor Networks' patent-pending technology is based on three years of pioneering research in the availability, reliability and security of networks and distributed systems, conducted at the University of Michigan by Arbor Networks' founders.

Funded by Battery Ventures, a leading venture capital firm, and Cisco Systems, Arbor Networks has been recognized in Red Herring's "Ten to Watch," and Network World's "Ten Start-Ups to Watch in 2001."

CONTACT US

research@arbornetworks.com

Arbor Networks, Research and Development
625 E. Liberty Street
3rd Floor
Ann Arbor, MI 48104
T: +1.734.327.0000
F: +1.734.327.9048
www.research.arbornetworks.com

Arbor Networks, Headquarters
610 Lincoln Street
Waltham, MA 02451
T:+1.781.684.0900
F: \+1 781.768.3299
www.arbornetworks.com