The RIPE NCC Internet Measurement Data Repository

Tony McGregor[1,2], Shane Alcock[1], and Daniel Karrenberg[2]

[1] University of Waikato, Hamilton, New Zealand
[2] RIPE NCC, Amsterdam, The Netherlands
(tonym, salcock)@cs.waikato.ac.nz
daniel.karrenberg@ripe.net

Abstract. This paper describes datasets that will shortly be made available to the research community through an Internet measurement data repository operated by the RIPE NCC. The datasets include measurements collected by RIPE NCC projects, packet trace sets recovered from the defunct NLANR website and datasets collected and currently hosted by other research institutions. This work aims to raise awareness of these datasets amongst researchers and to promote discussion about possible changes to the data collection processes to ensure that the measurements are relevant and useful to the community.

1 Introduction

A core requirement of any Internet measurement project is to acquire appropriate measurement data. However, privacy and security concerns often prevent researchers from being able to collect the data themselves. It is very important, therefore, that organisations that collect useful measurement data are able to share it with the research community. Time that would otherwise be spent conducting measurements can instead be dedicated to the analysis of existing data. Shared access to measurement resources also promotes collaboration between researchers and allows validation studies to be performed.

One common problem when sharing Internet measurement data is cultivating awareness of data sets amongst the research community. At present, publicly available data is typically scattered amongst a large number of hosting locations, meaning that it can be difficult for researchers to locate suitable datasets and keep informed of new datasets as they are released. DatCat [1] has helped in this regard, but it is not yet a comprehensive resource.

Maintenance of repositories hosted by research groups that depend on competitive grants for funding is also a significant concern. This was evidenced by the recent disappearance of the NLANR website which had hosted many passive trace sets, including the popular Auckland and Abilene traces. In this instance, it was fortunate that the contents of the site were salvaged by the University of Waikato with the support of CAIDA before they became inaccessible. However, the data could easily have been lost to the research community permanently if that intervention had not taken place.

As a Regional Internet Registry (RIR), the RIPE Network Coordination Centre (RIPE NCC) has the ability to collect a large quantity of measurement data that would be extremely difficult for researchers based in academic institutions to acquire themselves. Some of the RIPE NCC data is already publicly available, but each RIPE NCC project shares data independently in a variety of ways. Therefore, the RIPE NCC is developing a common and consistent platform for hosting and sharing Internet measurement data. While the primary goal is to streamline the mechanisms by which the RIPE NCC datasets can be accessed, the data repository will be open to other collectors who wish to share their measurement data with the research community. By grouping the datasets into a single repository, finding and accessing appropriate measurement data will be easier and awareness of the datasets that are available to researchers will be increased. One advantage of the RIPE NCC founding and operating a measurement data repository is that the continued existence of the repository does not depend solely on research grants and the likelihood of the repository disappearing is much smaller.

One of the most significant issues that arises when sharing Internet measurement data is that of anonymisation. Datasets that are being published for the first time will need to be anonymised in some fashion and agreements with users must be developed in order to prevent inappropriate disclosure of personal and commercial information. Such decisions will need to be made on a case-by-case basis, as each dataset can contain different types of sensitive information, e.g. policies for protecting IP addresses will not be applicable to personal information such as names and contact details.

Similarly, the structure and scope of a system for providing useful metadata and annotation of the shared datasets is yet to be completely determined. We expect that entries for datasets shared through the RIPE repository will be added to existing sites such as DatCat [1] and WITS [2], as well as a site that is developed and hosted as part of the RIPE repository itself. We hope to elicit ideas and thoughts from the wider research community with regard to the information that should be provided through such a system and the best format in which to present it. The RIPE NCC also plans to identify all users of the repository and keep contact with them during their research. It is hoped that this will encourage researchers to engage with the data collectors about how the measurements can be improved to be more useful and relevant.

The remainder of this paper describes the datasets that are currently under consideration for sharing through the RIPE repository. For each dataset, a brief overview of the dataset, its associated research project and the measurement techniques employed is presented.

2 RIPE Datasets

2.1 K-root

The K-root service is an Internet root name service operated by the RIPE NCC [3]. The server consists of seventeen nodes located both inside and outside of

Europe. Six of the nodes are global instances and are announced with an anycast 23-bit prefix. The remaining eleven nodes are local instances announced with a 24-bit prefix using the Border Gateway Protocol (BGP) no-export community tag [4]. Each node operates three distinct data collection systems.

Firstly, tcpdump [5] is used to capture passive traces of incoming traffic on port 53, i.e. Domain Name Server (DNS) queries. The trace files are rotated hourly and retained on disk for five days. Each node generates between 300 and 500 megabytes of compressed traces per hour. The total amount of data produced daily through this system is approximately 300 gigabytes.

As part of the "Day in the Life of the Internet" (DITL) project organised by CAIDA [6], traces for a 50 hour period have been fetched and archived each year. The 2008 DITL traces contain 1.46 billion packets and are 600 gigabytes in size [7]. At present, this data is hosted by DNS-OARC [8] and is available under their terms and conditions. Meta-data about the traces, including query rates per node and known issues that may affect analysis such as clock skew, has been documented publicly on DatCat, a measurement data catalogue [1].

Secondly, each node operates a DNS Statistics Collector (DSC) [9] that captures DNS traffic and summarises it into one minute bins. This data is used to generate graphs that are shown on the K-root website. In addition, the raw DSC output is transferred to RIPE NCC and archived indefinitely. The archive extends back to the beginning of 2008. The amount of data collected is estimated to be approximately 1 megabyte per day, with the entire 2008 dataset being less than 200 MB. The raw data has also been exported to DNS-OARC where it can be accessed by members.

Finally, Simple Network Management Protocol (SNMP) statistics are collected from the last-hop router serving each K-root node. The SNMP queries originate from RIPE NCC in Amsterdam. If connectivity to a queried router is lost, data will not be collected and the resulting dataset will not be contiguous. The statistics are summarised and exported into a round-robin database (RRD) where they can be queried and analysed using Cacti [10]. The RRD is configured to retain the SNMP data for a year.

2.2 Reverse DNS

RIPE NCC also hosts reverse DNS (rDNS) services for its delegated address space. There are four servers providing rDNS and other associated services. These servers process approximately 50 thousand queries per second which is more than triple the load of the K-root server. Because of the high query rate, it is not feasible to regularly collect passive traces directly on these servers. Occasional tcpdump traces have been collected when there was a specific need, such as during an attack, but the traces are short and irregular. However, if there was sufficient need, it may be possible to collect a sample or summary of the traffic using a dedicated collector on a mirrored switch port.

DSC is used on each of the rDNS servers and the raw data is kept indefinitely. At present, this data and the DSC graphs are only available internally at RIPE

NCC but could be made available to researchers if there was a need amongst the research community.

2.3 AS112

RIPE NCC hosts an Autonomous System (AS) number 112 [11] reverse DNS and dynamic DNS update server for the RFC 1918 private address space [12]. The server processes about 2000 transactions per second. As anyone can announce the AS112 prefix, there is no definitive list of AS112 servers. There are more than 50 servers listed at [11] but there are almost certainly others.

A passive tcpdump trace is collected from the RIPE NCC AS112 server annually and contributed to the DITL project [6]. More frequent passive captures could be scheduled from this server if required. In addition, DSC data is collected and used to generate graphs that are publicly available from the RIPE NCC AS112 website [13].

2.4 RIS

The Routing Information Service (RIS) is a set of 16 route collectors running quagga [14] that peer with approximately 600 BGP routers. Most of the collectors are located within the RIPE region, but there are a few elsewhere including the United States and South America. The routes learned are not used for routing traffic but are instead collected and published to provide a resource for understanding Internet routing and diagnosing routing problems. Around 100 of the peers provide a complete routing table and the others provide partial tables. The BGP sessions include both IPv4 and IPv6.

The collectors export route updates every five minutes and perform a full table dump every eight hours. It is normal that, at any one time, not all of the peers are actively peered with a collector. There are many causes explaining these peering gaps including configuration changes, equipment failure, network failure and human error. There are also some gaps in the dataset due to system failures. To reduce the impact of these errors, an automated system detects results that are much smaller than expected and informs an administrator who will investigate the fault. Also, it should be noted that some Interior Gateway Protocol (IGP) routes are leaked to RIS, meaning that there will be some single-bit prefixes advertised in the data.

The data collected from the RIS is stored in the Multi-threaded Routing Toolkit (MRT) format [15]. All data collected since the system started in 2000 is retained indefinitely. A month of data is approximately 22 gigabytes compressed and the entire dataset is close to one terabyte. The last three months of data is exported to a MySQL database which is also one terabyte in size. Quagga logs and a selection of meta-data is supplied with the RIS data. The logs show when peers have started or ended a BGP session and when timers expired, for example.

The raw data is currently publicly available through the RIS website. This is accessed by around 1000 distinct hosts each month, including the BGPmon

[16] and Cyclops [17] websites which use the data to offer route announcement and bogon notification services. In addition, RIPE NCC also publish weekly statistical reports, have released a variety of tools for querying and visualising the RIS data and have enabled Looking Glass queries to be sent directly to the collectors. Users can also subscribe to a notification service that will inform them of interesting events, such as a change in the advertisement for a particular AS.

2.5 Hostcount

The Hostcount project [18] generates statistics from a monthly DNS scan of approximately 100 top level domains (TLDs) within the RIPE region. The scan is performed by conducting a zone transfer on the DNS tree rooted at each TLD. During the walk, counts of A and PTR records are maintained for both forward and reverse IPv4 addresses as well as forward AAAA (but not reverse) IPv6 addresses.

As public zone transfers are disabled by most DNS administrators, the scan is not exhaustive. Some administrators permit the RIPE scan, but this is often under the condition that only statistics, rather than raw data, are published. The blocking of zone transfers has increased over time so the data from earlier years is a better reflection of the total number of hosts at that time.

Currently, the statistics are published via the Hostcount website [18]. These include the number of distinct hosts found at different levels of the DNS tree for each TLD, the number of zones discovered and the number of zone transfers that were permitted and successful. The raw data from between 1990 and 2007 has been archived and is currently in off-line storage. However, the present policy is to discard the raw data after statistics have been extracted. This could be reversed if there was sufficient need amongst the research community.

2.6 TTM

Test Traffic Measurements (TTM) is an active measurement system consisting of 105 operational probes [19] that has been operating since 1999. The majority of probes are located at ISPs and academic institutions within the RIPE region. Other probe locations include the United States, South America, Asia, the Middle East, Australia and New Zealand. The clocks on the probes are synchronised using GPS to give a 10 microsecond accuracy to real-time. With the exception of some probes that are in private meshes, the TTM project conducts full mesh measurements.

The probes regularly perform a series of active tests including a UDP one-way delay test, traceroute, a multicast performance test (limited to sites that have enabled multicast measurements), DNSMON measurements (see Section 2.7 below) and IPv6 pMTU discovery. In addition, the TTM probes support conducting ad hoc measurements that are initiated by authorised users. The ad hoc tests include a ping test and an HTTP page fetch. It is also possible to develop and run arbitrary ad hoc tests. There are limits on the range of destinations and the probing rate for ad hoc tests and the results are not released

to other sites or to the public. To maintain the integrity of the system, the probes are managed solely by the RIPE NCC. However, it may be possible to request special tests be performed using the TTM framework, provided there is no significant impact on the existing measurements.

At present, performance graphs are available via the TTM website to users who accept an electronic license agreement. Bulk data is also published using the CERN ROOT format [20] to researchers who sign a paper license agreement. The total dataset is approximately 0.7 terabytes in size, with the dataset from 2008 being 110 gigabytes.

2.7 DNSMON

The DNS Monitor project (DNSMON) collects data regarding the reachability and latency for some of the top levels of the DNS system [21]. The data is collected using 60 TTM probes that are not located in private meshes. The root domain, `.com`, `.net`, `.org`, `e164.arpa` and 24 country code TLDs (mostly from within the RIPE region) are measured. Performance over both IPv4 and IPv6 is measured for probes that have IPv6 connectivity. Name Server Identifier [22] information returned in the DNS response is retained and may be used in the future to generate per anycast instance graphs.

Summary statistics are available dating back to the commencement of the project in April 2003, although not all servers have been monitored since this time. Graphs are publicly available to anyone, but only paying subscribers can access graphs for the last two hours. Raw data is retained for approximately a year. Each probe collects between 10 and 20 megabytes of uncompressed data per day and the total dataset is 121 gigabytes in size. The raw data is available on request to country code TLD administrators and researchers, although non-subscribers are restricted from accessing the most recent two days of data.

2.8 RIPE DB

The RIPE DB is an open shared database that is maintained by the RIPE community [23], containing 3.2 million public object records. There are three classes of data stored in the database: Internet number registration objects, reverse DNS domain objects and route registry objects.

Internet number registration objects store details about IP addresses and AS numbers including references to administrative and technical contact information. The RIPE DB only contains records for the RIPE region, not the entire Internet. The vast majority of the 2.6 million registration objects are for IPv4 addresses, but there are also 24,000 IPv6 address and 18,000 AS number objects in the database.

At present, there are 410,000 reverse DNS objects and eight thousand forward domain objects in the RIPE DB. Historically, both forward and reverse domain information was stored but forward domains are no longer encouraged except as community support for small, emerging domains. It is likely that all forward

domain support will cease in the future. The reverse DNS records are used to create the zone files for the RIPE NCC reverse DNS service.

The route registry objects are used to provide an Internet Routing Registry (IRR), enabling organisations that participate in Internet routing to store and publish their routing policy. The RIPE DB contains approximately 100,000 route registry objects. The structure of the routing data in the database conforms to the Routing Policy Specification Language [24] and the structure and use of the route registry conforms to RFC 2650 [25]. Standard tools exist that can be used to check the policies stored in the IRR data for consistency and to generate router configurations from the IRR records. Unlike the other classes of object in the RIPE DB, the route registry is synchronised to other IRRs and copies of the information from the other IRRs is retained.

Public queries of the RIPE DB are supported through the use of both command-line and web `whois` tools. A daily limit is imposed on the number of queries that include personal information attributes. Bulk data is also available via FTP. The bulk data files are generated daily, including both a file with the complete database and files split by object type. Personal details, such as the person and maintainer objects, are not included in these files. The complete database file is approximately 150MB in size. In addition, it is possible to subscribe to a near real time mirror feed of the database for an annual fee.

Access to personal data within the RIPE DB is restricted for both legal and practical (e.g. limiting abuse) reasons. At present, the restrictions are applied at a very broad level which sometimes results in limitations that are inappropriate. For example, an ISP that has entered a large number of person objects may not be able to access all the objects that they have created. This problem will be resolved as part of the development of the common RIPE data sharing platform.

3 External Datasets

3.1 Auckland

The Auckland dataset consists of a series of trace sets collected by the WAND group at the University of Waikato. The traces were collected at the University of Auckland in New Zealand, measuring the link between the University and the Internet. All of the traces were captured using DAG hardware capture cards [26], although the card model was upgraded on several occasions. Each of the Auckland trace sets is briefly described in Table 1, which is based on detailed summaries provided by the Waikato Internet Trace Storage (WITS) project [2].

There have been some significant variations in the capture configuration between each trace set. Some of the changes were necessitated by the network infrastructure being upgraded, meaning that the measurement point was no longer capable of capturing all of the traffic it had previously. In other cases, the amount of packet header data that was retained is varied, e.g. the Auckland VI traces captured the first 64 bytes of every TCP, UDP and ICMP packet whereas the Auckland VII traces are limited to only the ATM cell header for all packets regardless of protocol.

Table 1. The Auckland trace sets

Name	Format	Year	Duration	Packets	Bytes	Size
I	ERF	1999	7 days	169 M	8 GB	2 GB
II	Legacy ATM	2000	24 days	996 M	359 GB	26 GB
IV	Legacy ATM	2001	45 days	3,157 M	1,269 GB	64 GB
V	ATM Cell	2001	7.5 hours	2,710 M	133 GB	8 GB
VI	Mixed Legacy	2001	4.5 days	844 M	345 GB	17 GB
VII	ATM Cell	2001	15.5 hours	6,040 M	297 GB	19 GB
VIII	ERF	2003	13 days	1,654 M	698 GB	68 GB

Most of the Auckland traces were publicly released by NLANR [27] and frequently feature in measurement literature such as [28] and [29]. With the recent demise of the NLANR site, the traces have become difficult for researchers to acquire. In addition, the RIPE repository will include the Auckland I and Auckland V trace sets, which were not available from NLANR.

3.2 Waikato

Table 2. The Waikato trace sets

Name	Format	Years	Duration	Packets	Bytes	Size
I	ERF	2003-2005	620 days	53,263 M	21,434 GB	1,329 GB
II	ERF	2005-2006	301 days	34,712 M	15,789 GB	839 GB
III	ERF	2006-2007	160 days	21,984 M	9,144 GB	545 GB
IV	ERF	2007	56 days	10,128 M	4,588 GB	255 GB
V	ERF	2007	99 days	19,710 M	9,740 GB	491 GB
VI	ERF	2007-2008	135 days	20,886 M	11,092 GB	495 GB

The Waikato dataset is a collection of six very long-duration trace sets captured at the border of the University of Waikato network by the WAND group. The capture point is located between the University network infrastructure and the commodity Internet, allowing access to all traffic entering and exiting the University but excluding any internal traffic. All of the traces were captured using software that was specifically developed for the Waikato capture point [30] and a DAG 3 series hardware capture card [26]. All IP addresses within the traces are anonymised using Crypto-Pan AES encryption [31], with the encryption key being changed on a weekly basis. A brief description of each of the trace sets is shown in Table 2.

The location and hardware of the capture point has remained unchanged since the first capture began in 2003. The software has been upgraded between

each trace set, resulting in some minor variations. For example, in Waikato I packets are truncated at the end of the transport header but subsequent trace sets retained four bytes of unanonymised application payload for all packets except in the case of DNS packets where twelve bytes were kept.

The Waikato I trace set is currently available for public download from the WITS archive [2]. The other Waikato trace sets will be made available through the RIPE data repository.

3.3 NLANR Datasets

The NLANR project [27] collected both active and passive datasets. These data sets have been the focus of a significant amount of research and many papers have been published based on them. Although the project is complete, the datasets collected are still in demand. They have been preserved by the WAND network research group and are available from the WITS [2] repository. The traces will also be hosted on the RIPE repository.

4 Conclusion

This paper catalogues and describes the large quantity of Internet measurement data that will be shared with the research community through a data repository hosted by the RIPE NCC. While much of the data described here is already publicly available, it is scattered amongst a variety of hosting organisations or, in the case of the Auckland traces, is no longer available from the original source. By creating a common portal for sharing and accessing all of the data, it will become easier for researchers to locate and download suitable measurement data for their particular project.

The primary aim of the repository is to bridge the gap between the organisations capable of conducting Internet measurements and the researchers who analyse the measurement data. This is evidenced by the partnership with the University of Waikato to enable the Auckland and Waikato datasets to be mirrored on the RIPE repository. However, the relationship must function in both directions to ensure the collected data is relevant and useful. As a result, we encourage submissions from the community regarding the collection and format of any of the aforementioned datasets that would improve their utility to researchers.

References

1. Cooperative Association for Internet Data Analysis (CAIDA): DatCat: Internet Measurement Data Catalog `http://imdc.datcat.org/`.
2. WAND Network Research Group: WITS: Waikato Internet Traffic Storage `http://www.wand.net.nz/wits/`.
3. RIPE NCC: K-root `http://k.root-servers.org/`.

4. Chandra, R., Traina, P., Li, T.: RFC 1997 - BGP Communities Attribute (August 1996).
5. Jacobson, V., Leres, C., McCanne, S.: tcpdump http://www.tcpdump.org/.
6. Cooperative Association for Internet Data Analysis (CAIDA): A Day in the Life of the Internet http://www.caida.org/projects/ditl/.
7. Nagele, W., Buddhdev, A., Wessels, D.: K-root DNS traces DITL 2008 (collection) http://imdc.datcat.org/collection/1-0690-J=K-root+DNS+traces+DITL+2008.
8. DNS-OARC: Domain Name System Operations, Analysis and Research Center https://www.dns-oarc.net/.
9. The Measurement Factory: DSC: A DNS Statistics Collector http://dns.measurement-factory.com/tools/dsc/.
10. Cacti: http://www.cacti.net/.
11. The AS112 Project: http://www.as112.net/.
12. Rekhter, Y., Moskowitx, B., Karrenberg, D., de Groot, G.J., Lear, E.: RFC 1918 - Address Allocation for Private Internets (February 1996).
13. RIPE NCC: RIPE NCC AS112 http://www.ripe.net/as112/.
14. Quagga: http://www.quagga.net/web/quagga.html.
15. Blunk, L., Karir, M., Labovitz, C.: MRT routing information export format (IETF Draft) http://tools.ietf.org/html/draft-ietf-grow-mrt-04.
16. BGPmon: http://bgpmon.net/.
17. Oliveira, R.V., Lad, M., Zhang, L.: Cyclops http://cyclops.cs.ucla.edu/.
18. RIPE NCC: Hostcount http://www.ripe.net/is/hostcount/stats.
19. RIPE NCC: Test Traffic Measurements http://www.ripe.net/ttm/.
20. CERN: Root http://root.cern.ch/drupal/.
21. RIPE NCC: RIPE NCC DNS Monitoring Services http://dnsmon.ripe.net/dns-servmon/.
22. Austein, R.: RFC 5001 - DNS Name Server Identifier (NSID) Option (August 2007).
23. RIPE NCC: RIPE Database http://www.ripe.net/db/.
24. Alaettinoglu, C., Villamizar, C., Gerich, E., Kessens, D., Meyer, D., Bates, T., Karrenberg, D., Terpstra, M.: RFC 2622 - Routing Policy Specification Language (RPSL) (June 1999).
25. Meyer, D., Schmitz, J., Orange, C., Prior, M., Alaettinoglu, C.: RFC 2650 - Using RPSL in Practice (August 1999).
26. Endace Measurement Systems, Ltd: http://www.endace.com.
27. McGregor, A., Braun, H.W., Brown, J.: The NLANR NAI Network Analysis Infrastructure. IEEE Communication Magazine: Special Issue on Network Measurement (May 2000) 122–128
28. Erman, J., Arlitt, M., Mahanti, A.: Traffic Classification Using Clustering Algorithms. In: MineNet '06: Proceedings of the 2006 SIGCOMM workshop on Mining network data, New York, NY, USA, ACM (2006) 281–286
29. McGregor, A., Hall, M., Lorier, P., Brunskill, J.: Flow Clustering Using Machine Learning Techniques. In: Passive and Active Measurement. (2004) 205–214
30. WAND Network Research Group: WDCap http://research.wand.net.nz/software/wdcap.php.
31. Fan, J., Xu, J., Ammar, M.H., Moon, S.B.: Prefix-preserving IP address anonymization: measurement-based security evaluation and a new cryptography-based scheme. Computer Networks **46**(2) (2004) 253 – 272