

IPv6 Routing Table Analysis

FH- Diplomstudiengang
für Präzisions- System- und Informationstechnik
Vertiefung: Informationstechnik

Name: Attila Miklos

Firma: RIPE NCC

Matrikelnummer: 0210016147

Firmenbetreuer: Dr. H.A.J.R. Uijterwaal

Datum: 07 April 2006

FH- Betreuer: Wilfried Wöber

1	Introduction	- 1 -
1.1	RIPE.....	- 1 -
1.2	RIPE NCC.....	- 1 -
2	The Internet.....	- 2 -
2.1	History.....	- 2 -
2.2	Routing.....	- 3 -
2.3	Autonomous Systems (AS).....	- 3 -
2.4	Routing Concepts.....	- 3 -
2.5	The BGP Protocol	- 4 -
3	IP Version 6.....	- 4 -
3.1	IPv6 vs. IPv4	- 4 -
3.1.1	Expanded Addressing Capabilities	- 5 -
3.1.2	Header Format Simplification.....	- 5 -
3.1.3	Improved support for Extensions and Options.....	- 5 -
3.1.4	Flow Labelling Capability	- 5 -
3.1.5	Authentication and Privacy Capabilities	- 6 -
3.2	Extension Headers.....	- 6 -
3.3	Addressing	- 7 -
3.3.1	Address Format	- 7 -
4	The Routing Information Service.....	- 8 -
4.1	The Idea behind RIS	- 8 -
4.2	Goals of the Routing Information Service.....	- 9 -
4.3	Remote Route Collector.....	- 10 -
4.4	The Evolution of RIS	- 11 -
4.4.1	Classic RIS	- 11 -
4.4.2	RISng.....	- 12 -
4.5	BGP Updates	- 13 -
4.6	Tools of the RIS	- 14 -
4.6.1	RISreport	- 14 -
4.6.2	MyASn	- 14 -

4.6.3	RISwhois.....	- 14 -
4.6.4	ASInuse	- 15 -
4.6.5	Looking Glass	- 15 -
5	The RIS Database.....	- 16 -
5.1	Recent Changes on the Database [11]:	- 17 -
6	What is the RISreport?	- 17 -
6.1	Development so far	- 17 -
6.2	IPv6 version of the RISreport	- 18 -
6.2.1	Table Definitions [10]	- 18 -
6.2.2	Perl Files	- 20 -
6.2.3	SQL Queries	- 21 -
6.2.3.1	BGP Count.....	- 21 -
6.2.3.2	BGP Unique Announcements Count	- 21 -
6.2.3.3	BGP Withdrawals.....	- 22 -
6.2.3.4	BGP State Changes	- 22 -
6.2.4	BGP Sample Plots	- 23 -
6.2.4.1	BGP Count Plot	- 23 -
6.2.4.2	BGP Unique Announcements Count	- 24 -
6.2.4.3	Normalized BGP Count	- 24 -
6.2.4.4	IPv6 Prefix Distribution	- 25 -
7	Conclusion	- 26 -
8	Further Work	- 27 -
	List of Acronyms.....	I
	Bibliography.....	II

1 Introduction

1.1 RIPE

RIPE (Reseaux IP Europeans)[1] is a collaborative forum open to all parties interested in wide area IP networks. RIPE has the objective to ensure the administrative and technical coordination, to allow the operation and expansion of a pan-European IP network. The activities of RIPE are on voluntary basis and the decisions are made by consensus. The work is carried out from different working groups which meet two times a year throughout the RIPE Meetings. For participation in RIPE there are no membership requirements.

1.2 RIPE NCC

The RIPE-NCC is one of the 5 independent, not-for-profit membership Regional Internet Registries (RIRs) that exist in the world today. RIRs are recognized by ICANN (Internet Corporation for Assigned Names and Numbers) to serve and represent large geographical regions in managing and distributing address space. The RIPE-NCC has its members in Europe, the Middle East and parts of Central Asia. In addition to the allocation of Internet address space it also manages BGP AS numbers, the reverse domain space and the RIPE database. The RIPE NCC also provides services for the benefit of the Internet community at large. These services include development and maintenance of the RIPE Whois Database, administrative support for the RIPE community.

Other activities include [12]:

- Outreach activities with governments and industry-related organisations
- Management of one of the 13 root name servers (K-root)
- Deployment of a routing database
- Co-ordination support for ENUM delegations

- Neutral measuring network that provides publicly accessible and authoritative statistics on the operation of the Internet.

2 The Internet

2.1 History

In August 1962, J.C.R. Licklider from the MIT had a vision about a globally connected set of computers through which everyone could quickly access data and programs from any site. In theory, the concept was very much like the Internet of today. Licklider was the first head of the computer research program at DARPA [2], starting in October 1962. He convinced his successors, Ivan Sutherland, Bob Taylor, and MIT researcher Lawrence G. Roberts, of the importance of this networking concept and of the theoretical feasibility of communications using packets rather than circuits. That was a big step along the path towards computer networking. The other key step was to make the computers talk to each other. To explore this, in 1965 working with Thomas Merrill, Roberts connected the TX-2 computer in Massachusetts to the Q-32 in California with a dial-up telephone line creating the first wide-area computer network ever built.

The result of this experiment was the realization that the computers could work well together, running programs and retrieving data as necessary on the remote machine. In late 1966 the MIT researcher Roberts went to DARPA to develop the computer network concept and quickly put together his plan for the "ARPANET", which was published in 1967. Afterwards computers were added quickly to the ARPANET during the following years, and work proceeded on completing a functionally complete Host-to-Host protocol and other network software.

Today's Internet became a very large distributed Network and it is not possible to operate it centrally. Therefore 5 RIRs provide registration services to their respective regions around the globe.

2.2 Routing

Routing is a technique to transport information from one network to another on the Internet. This transportation can happen through various Routing Protocols e.g. BGP. This Protocol is discussed in chapter 2.5

2.3 Autonomous Systems (AS)

An autonomous system (AS) is a network that is controlled by a common network administrator on behalf of a single administrative entity (e.g. university). An AS is also sometimes referred to as a routing domain. A globally unique number, obtained from one of the RIRs, called an Autonomous System Number (ASN), identifies the AS. Networks within an autonomous system communicate routing information to each other using an Interior Gateway Protocol (IGP). By using the Border Gateway Protocol (BGP), an autonomous system shares routing information with other autonomous systems. The internet consists of thousands of subnets, which are called autonomous system (AS). For details see RFC1930.

2.4 Routing Concepts

Usually there are two routing protocols. Routing within an autonomous system (IGP) and routing between autonomous systems (EGP). IGP use routes, defined by an administrator and default routes. EGPs cannot have default routes. Routers announce information about the connectivity to their reachable network. All of these networks have to be listed in the routing table. Because of that, routers which are placed on the border of an autonomous system are labelled *default-free* routers.

The standard in the global Internet is the *Border Gateway Protocol Version 4* (BGP4) for inter domain routing. The RIS, which is discussed in chapter 4, deals only with routing information between autonomous systems.

2.5 The BGP Protocol

The primary function of a BGP [3] speaking system is to exchange network reachability information with other BGP systems. This network reachability information includes information on the list of Autonomous Systems (ASs) that reachability information traverses.

BGP is meant to be used between autonomous systems to provide an interdomain loop-free topology. BGP is used within an AS, as a pipe between border routers running external BGP to other ASs. A neighbour connection also called a peer connection, between two routers can be established within the same AS, in which case BGP is called internal BGP (IBGP). A peer connection can also be established between two routers in different ASs. BGP is then called external BGP (EBGP).

3 IP Version 6

3.1 IPv6 vs. IPv4

IPv4 was defined in the 1970s when the structure of the protocol was sufficient for the existing networking infrastructure at that time. The exhaustion of the address space, the need for Quality of Service (QoS) and encryption, were some of the reasons to develop a new version of IP. IPv6 is a new version of the Internet Protocol, developed as the successor to IPv4.

The following paragraphs can be found in [4], describing the main differences between the two protocols.

3.1.1 Expanded Addressing Capabilities

The IP address size is increased from 32 bits to 128 bits. Therefore it is possible to use more levels of addressing hierarchy, a greater number of nodes and a simpler auto configuration of addresses. The adding of a “scope” field to multicast addresses also improves scalability. A new type of address, the “anycast address” was introduced to send packets to any one of a group of hosts.

3.1.2 Header Format Simplification

Some IPv4 header fields have been dropped or made optional, to reduce the common case processing cost of packet handling and to limit the bandwidth cost of the IPv6 header. The header has a fixed size of 40 bytes. A detailed description of the IPv6 header and its fields can be found in [4].

3.1.3 Improved support for Extensions and Options

Changes in the way IP header options are encoded allows more efficient forwarding, less stringent limits on the length of options and greater flexibility for introducing new options in the future.

3.1.4 Flow Labelling Capability

A new capability is added to enable the labelling of packets belonging to particular traffic “flows” for which the sender requests special handling, such as non-default quality of service or “real-time” service.

3.1.5 Authentication and Privacy Capabilities

Extensions to support authentication, data integrity, and data confidentiality are specified for IPv6. IPv6 requires that every link in the Internet has an MTU of at least 1280 octets. On any link that cannot convey a 1280-octet packet in one piece, link-specific fragmentation and reassembly must be provided at a layer below IPv6. For the transfer of packets that are larger than the MTU size, Fragment headers have to be used. Therefore it is strongly recommended by [4] that IPv6 nodes implement Path MTU Discovery, in order to discover and take advantage of path MTUs greater than 1280 octets. The IPv6 header was reduced to a basic functionality and set to a fixed size of 40 bytes to minimize the processing time at hops in between the source and the destination host. Any additional options are implemented in “extension headers”.

3.2 Extension Headers

In IPv6, optional internet-layer information is encoded in separate headers that may be placed between the IPv6 header and the upper-layer header in a packet.

There are a small number of such extension headers, each identified by a distinct Next Header value. An IPv6 packet may carry zero, one, or more extension headers, each identified by the Next Header field of the preceding header.

A full implementation of IPv6 includes the following extension headers:

- IPv6 header
- Hop-by-Hop Options header
- Destination Options header
- Routing header
- Fragment header
- Authentication header
- Encapsulation Options header
- Destinations Options header
- Upper-layer header

3.3 Addressing

There are three types of IPv6 addresses:

- Unicast:* An identifier for a single interface. A packet sent to a unicast address is delivered to the interface identified by that address.
- Anycast:* An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to an anycast address is delivered to one of the interfaces identified by that address (the "nearest" one, according to the routing protocols' measure of distance).
- Multicast:* An identifier for a set of interfaces (typically belonging to different nodes). A packet sent to a multicast address is delivered to all interfaces identified by that address.

IPv6 addresses of all types are assigned to interfaces, not nodes. An IPv6 unicast address refers to a single interface. Since each interface belongs to a single node, any of that node's interfaces' unicast addresses may be used as an identifier for the node.

3.3.1 Address Format

IPv6 addresses are normally written as eight groups of four hexadecimal digits. These are often followed by a slash and the prefix length (called a CIDR range), which turns them into a range of IPv6 addresses.

For example: 2001:0db8:85a3:08d3:1319:8a2e:0370:7334 /64

If a four-digit group is 0000, the zeros may be omitted.

For example: 2001:0db8:85a3:0000:1319:8a2e:0370:7344 /64

can be shortened: 2001:0db8:85a3::1319:8a2e:0370:7344 /64

A detailed description of the IPv6 address format can be found in [5].

4 The Routing Information Service

4.1 *The Idea behind RIS*

The traffic all over the Internet is packed into IP packets and shipped to the final destination. In most cases, an IP packet passes through several autonomous systems on its way. This means for these packets that multiple independent administrative authorities are involved in transporting the data. The end-user only notices whether the communication is working or not.

If the communication is not working properly the operator can only examine his local network. A look into the routing table will show whether a network is reachable at present. Normally there is no chance to observe other border routers. A workaround is to provide so-called *looking glasses* allowing an observer how the traffic is routed back from the remote AS into the observer's network. As mentioned above it is only possible to observe one's own network that has installed some tools. The observation is restricted to the current configuration. It is not possible to track a problem over time. Therefore it is time for a tool that collects route information over time. This is the Routing Information Service. [7]

4.2 Goals of the Routing Information Service

The goal of the Routing Information Service (RIS) is to collect routing information between Autonomous Systems (AS) and their development over time.

The RIS collect and store default free BGP announcements as a function of time from several locations and provide that information to the users of the service, allowing them to see the full picture with all routes that are currently anywhere and their development over time. In other words it can be regarded as one integrated looking-glass for the entire Internet that includes history information. [6]

1. The routing table and its development over time can be used to check for local and global convergence of the table and routing flaps.
2. The routing table reflects the policies announced by the sites operating the routers. At a local level, the data can be used to verify the setup of the routers and correct any errors.
3. On a global level, the RIS data can be used to compare policies registered in the Routing Registry against the policies that are actually implemented.
4. Related to this is, RIS will provide information about fake routes inserted into the network, for example by spammers or attackers.
5. The routing table also reflects the addressable IP-networks that can be reached, which AS-numbers are actually being used and such. These are valuable statistics.

4.3 Remote Route Collector

A Remote Route Collector is a software router, running on the FreeBSD platform, which only collects default free BGP routing information. RIPE NCC uses the BGP daemon from Zebra/Quagga[8] project to interface with BGP peers. The collected raw data is regularly transferred via rsync to a central storage area at the RIPE NCC in Amsterdam.

All of our Route Collectors use UTC (Coordinated Universal Time) as their time standard and NTP (Network Time Protocol) for synchronization.

Currently there are 13 RRCs:

rrc00	RIPE NCC, Amsterdam
rrc01	LINX, London
rrc02	SFINX , Paris
rrc03	AMS-IX / NL-IX / GN-IX, Amsterdam
rrc04	CIXP, Geneva
rrc05	VIX, Vienna
rrc06	NSPIXP2, Otemachi
rrc07	Netnod, Stockholm
rrc10	MIX, Milan
rrc11	NYIIX, New York
rrc12	DE-CIX, Frankfurt
rrc13	MSK-IX, Moscow
rrc14	PAIX, Palo Alto

4.4 The Evolution of RIS

4.4.1 Classic RIS

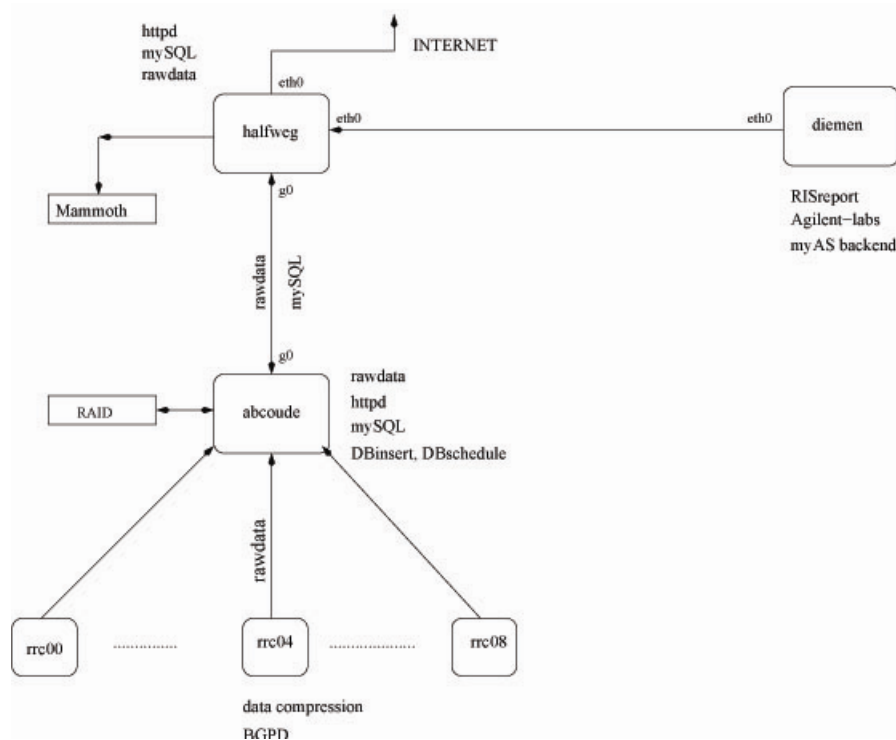


Figure 4.1: Classic RIS Infrastructure

The RRCs only run Zebra bgpd. All other processing and all database insertions are done on abcoude, the central server. Abcoude is running Solaris / Sun SPARC Ultra Enterprise 420R .The MySQL insertion scripts are written in Perl.

Halfweg, the front-end server, handles the user interface and store a replicated copy of the database.

Problems:

- Database insertion of data from route collectors on a single central machine is slow. It can sometimes take more than 24 hours to insert a single day's data!
- Little capacity to add more RRCs or full BGP feeds
- Limited attributes are stored in the database

E.g. only first 255 characters of AS Path are stored

4.4.2 RISng

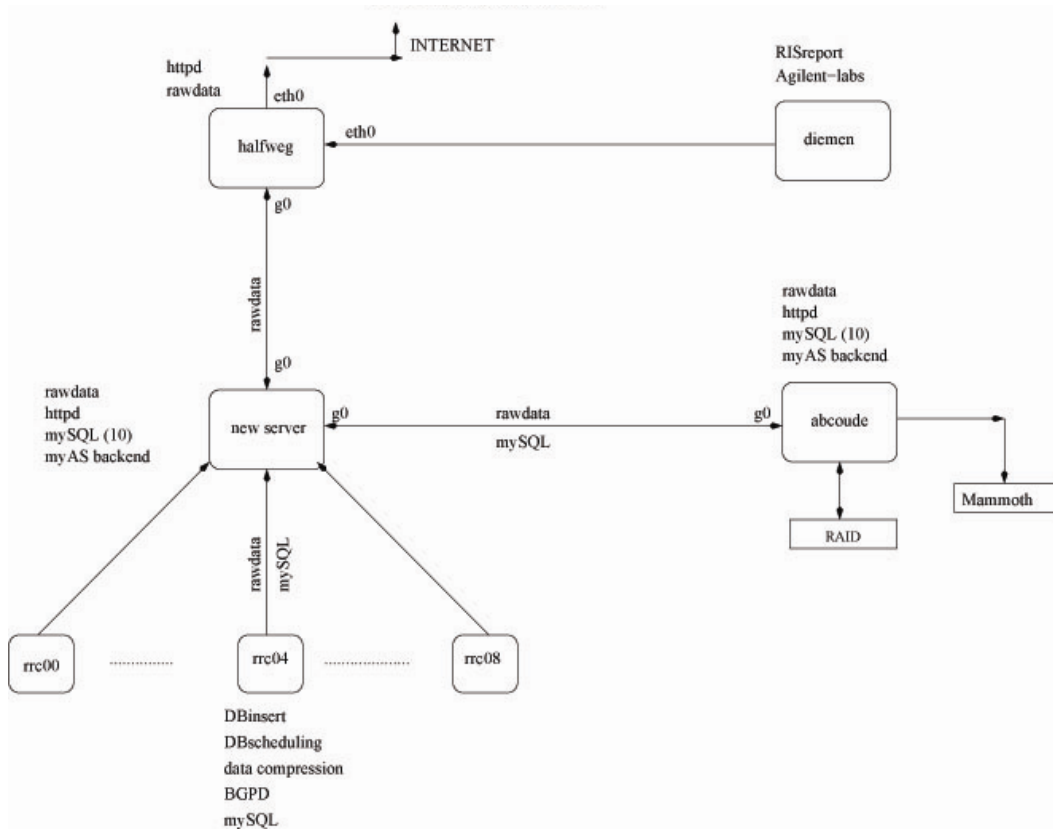


Figure 4.2: RISng Infrastructure

The goals of the RISng were to improve the scalability, easier software maintenance and to store more complete route attributes. The new database structures remove arbitrary limits on AS Path, store additional attributes and perform database insertion locally on route collectors. The database insertion process is now written in C instead of Perl.

4.5 BGP Updates

One BGP update can contain multiple announcements and withdrawals. These so-called sub-updates are stored separately in the RIS database. Each sub-update consists of a set of data fields. RIS stores these fields in the database:

AS path: This is a sequence of intermediate AS between source and destination routers. It reflects the path that an IP packet will travel through the Internet to its destination.

Origin AS: The AS that originated this BGP update. This must not necessarily be the AS where the destination host is located. In this case of route aggregation, which is used to reduce the number of routing entries in the routing table, the AS that has aggregated is the new origin AS.

Type of BGP updates: This field keeps track of the type of the sub-update: announcement, withdrawal or state change.

Time stamp: RIS operates on UTC. It is important that all parts of RIS operate using the same time standard. Otherwise the output would not be useful. Each sub-announcement is stored with its time stamp. The time is stored in the so-called UNIX time stamp format, which are the number of seconds since 0:00:00 January 1, 1970 UTC.

IP prefix and length: Each update contains an IP prefix. RIS applies CIDR Notation. Each prefix is announced with the prefix length used as network mask.

Remote Route Collector: This field stores which RRC has collected this BGP message.

4.6 Tools of the RIS

The RIS has several tools which are made available for everybody on the RIS website. These are the most popular and important ones:

4.6.1 RISreport

This tool was designed by Thomas Franchetti for his master thesis [9] to create statistical summaries of the global IPv4 routing environment, seen by the RIS. The information is presented in graphs and the results are published on the RIS website. Additional to this website, a daily email report will be generated. This part of the RIS-report contains only text information. This email is for the operator of the RISreport. The content of the email is a confirmation whether the daily RISreport is created correctly or not.

4.6.2 MyASn

The objective of myASn is to provide an alarm notification system that monitors the propagation of eBGP routes. The system allows users to 'lock down', on a per-prefix-basis, the parts of the AS path one would expect to be announced to the RIS. If 'unexpected' routing information (events) is detected by myASn, the user will be notified of that deviation via the user interface. Email and syslog are optional alarm notification features.

4.6.3 RISwhois

RISwhois providing a higher level view over the most recently collected set of routing tables from the Remote Route Collectors (RRCs). Given an IPv4 or IPv6 address,

RISwhois will tell which prefixes and origin ASes on which RRCs match that particular IP. The RIPE Whois database provides a mechanism for finding contact and registration information for networks in the RIPE NCC service region by querying the RIS Whois Database.

4.6.4 ASInuse

Show the last time that a particular AS or Prefix was seen by the RIS in the global routing table and is mainly used by RIPE NCC Registration Services to check use of assigned resources. The query is made throughout an AS number. It is also possible to choose which RRC should be queried.

4.6.5 Looking Glass

Looking Glasses provide access to the routers of the remote AS, allowing an observer to see how traffic from the remote AS is routed back to his network (similar to a trace route on the remote AS) and thus determine the reachability of the observer's network as seen from the remote AS (prefix and AS-PATH).

5 The RIS Database

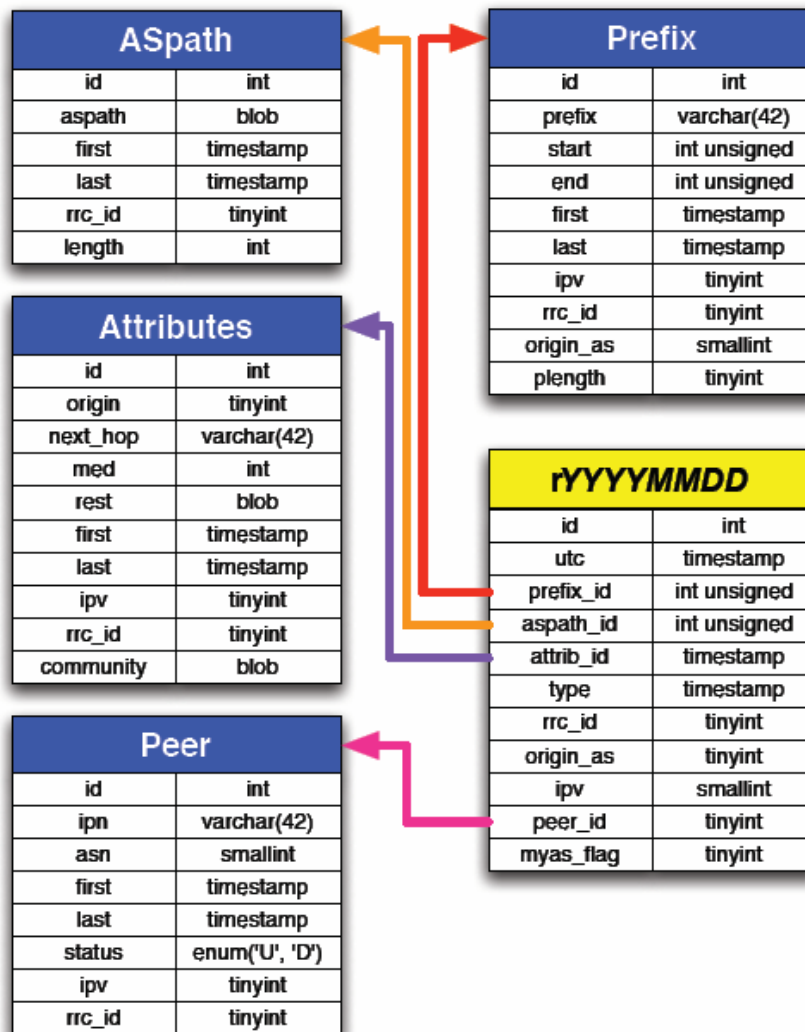


Figure 6.5: RIS Database Tables

Each table contains an id field which is a unique value and is used as primary key. All foreign keys are named by their type of reference and followed by the suffix *_id*. Tables containing BGP updates are called r followed by the specific date: rYYYY-YMMDD (Figure 6.5) The ASpath table contains all AS paths seen by the RIS. An AS path is the sequence of autonomous systems that BGP announcements has passed on its way through the Internet. The Peer table contains information about all BGP peers. Routing announcements are collected from these peers. The Prefix table is

used to store all prefixes seen by the RIS. The prefix field contains the announced prefix, the start and end field contains the decimal value of the IP address.

5.1 Recent Changes on the Database [11]:

The changes effects that more BGP attributes are stored, arbitrary restrictions are removed. The distributed architecture enable off-loading the central server and use “spare” CPU cycles on RRCs themselves.

Now data is inserted as soon as dump file is ready into the database. We don’t have to wait for the file to be transferred to the central machine. Dump interval reduced from 15 minutes to 5 minutes. The database is now ready for IPv6, which means that the Database has IP Version fields and all string fields large enough to hold both IPv4 and IPv6 addresses and prefixes where necessary.

6 What is the RISreport?

6.1 Development so far

The RISreport from Thomas Franchetti was designed in 2001 for IPv4. Since 2001 the traffic in IPv6 data has increased nearly exponential. An IPv6 version of the RIS-report became necessary. In the last four years also the number of RRCs was increased from two to thirteen. That was necessary to collect the data from different places all over the world and cover the global Internet as good as possible. To be able to handle the largely increased amount of data, some changes have to be done on the RIS database structure, mentioned in chapter 4.4. A basic subject of my internship at RIPE NCC was the development of an IPv6 version of the RISreport, which is discussed in chapter 6.2. This report is now in production and the results are published for the user community on the RIS Website [13]. The basic process flow is as follows.

The data collection is done by several RRCs. This data is then inserted into the RIS database, where MYSQL is used as database server. It holds at least the BGP data of the last three month. The data in the RIS database is the source of the RISreport. For the report, the data is queried by a universally usable tool called SQL2RRD. This program extracts the data from the database and fed it into the RRD data files. RRD stands for *round robin database*. The RRD tool was written by Tobi Oetker [11]. Round robin is a technique that works with a fixed amount of data and a pointer to the current element. Like a circle with some dots plotted on the edge, these dots are the places where data can be stored. With this technique the RRD will not grown in size and therefore requires no maintenance.

6.2 IPv6 version of the RISreport

6.2.1 Table Definitions [10]

The following two tables store all IPv6 peers and prefixes. These tables are queried by the IPv6 RISreport.

peer6 table:

Field	Type	Null	Key	Default	Extra
id	Int(11)		Pri	NULL	Auto_increment
ipn	Varchar(39)		Mul	0	
asn	smallint(5)			0	
first	Int(11)	Yes		2147483647	
last	Int(11)	Yes		0	
status	Enum('U','D')			U	
num_prefixes	Int(11)			0	

Figure 5.3: peer6 Table Definition

- id** Assigned automatically when a new row of data is inserted into the table, this uniquely identifies the row and is referenced by the 6Updates and 6RIB tables (as peer_id) below
- ipn** The IP address of this peer
- asn** The autonomous system number of this peer

- first** The time of first activity referencing this peer seen by the RIS. Generally this will be when the peering was first configured on the route collector.
- last** The time of the most recent activity from the peer seen by the RIS status Whether the peer is up (we're receiving updates) or down (only receiving state changes)
- num_prefixes** Number of prefixes received from this peer (only updated when processing RIB dumps).

prefix6 table:

Field	Type	Null	Key	Default	Extra
id	Int(11)		Pri	NULL	Auto_increment
prefix	varchar(43)		Mul		
start0	int(11) unsigned		Mul	0	
start1	int(11) unsigned		Mul	0	
start2	int(11) unsigned			0	
start3	int(11) unsigned			0	
end0	int(11) unsigned		Mul	0	
end1	int(11) unsigned		Mul	0	
end2	int(11) unsigned			0	
end3	int(11) unsigned			0	
first	Int(11) unsigned	Yes		2147483647	
last	Int(11) unsigned	Yes		0	
aspath_id	Int(11)		Mul	0	
plength	Tinyint(3) unsigned	Yes		Null	

Figure 5.4: prefix6 Table Definition

- id** Assigned automatically when a new row of data is inserted into the table, this uniquely identifies the row and is referenced by the Updates and RIB tables (as prefix_id) below.
- prefix** The text representation of the prefix and prefix length.
- start0...start3** The numerical value of the first IP address in the range (most significant word first). Note: only the first 64 bits of the address are indexed.
- end0...end3** The numerical value of the last IP address in the range (most significant word first). Note: only the first 64 bits of the address are indexed.
- first** The time that this prefix was first seen by the RIS

last	The time of the last activity for this prefix seen by the RIS
aspath_id	Reference into the aspath table of the most recently seen AS Path originating this prefix.
Plength	The prefix length.

6.2.2 Perl Files

These files create the RISreport. Each of them had to be changed for the IPv6 version of the RISreport. The changes affect the SQL query. For detailed information about the SQL Queries see chapter 6.2.3

ipv6_run_mysql2rrd.sh the main file which is executed automatically once a day and calls the following files

	Responsible for:
ipv6_rrd_updates_bgpcount.pl	"bgp count"
ipv6_rrd_updates_bgpcount_total.pl	"total bgp count"
ipv6_rrd_updates_bgpcount_a.pl	"announcement"
ipv6_rrd_updates_bgpcount_a_total.pl	"total announcement"
ipv6_rrd_updates_bgpcount_sc.pl	"state change"
ipv6_rrd_updates_bgpcount_sc_total.pl	"total state change"
ipv6_rrd_updates_bgpcount_ua.pl	"unique announcements"
ipv6_rrd_updates_bgpcount_ua_total.pl	"total unique announcements"
ipv6_rrd_updates_bgpcount_w.pl	"withdrawals"
ipv6_rrd_updates_bgpcount_w_total.pl	"total withdrawals"
ipv6_rrd_updates_template.pl	"template"
ipv6_rrd_updates_total_template.pl	"template"
ipv6_mysql2rrd_gen.conf	"config file for the rrd BD files"
ipv6_prefix_n.pl	"responsible for prefix distribution"
prefix_n.plt	"config file for gnuplot prefix graph"

ipv6_mysql2rrd "for drawing graphs",
"for creating html files",
"for creating index files"

Finally the created html files and graphs for the RISreport are synchronized back to *Halfweg*, the front-end server, where the user can see the completed RISreport.

6.2.3 SQL Queries

6.2.3.1 BGP Count

The following query is stored in *ipv6_rrd_updates_bgpcount.pl*.

```
SELECT peer_id, 60*FLOOR (utc/60) AS timestamp,
COUNT (*) AS count, ipn AS peer_ip FROM $param {table}, peer6
WHERE peer6.id = peer_id
AND utc >= $param {from_time} AND utc <= $param {to_time}
GROUP BY timestamp, peer_id
ORDER BY timestamp, peer_id
```

The result will be a list of the number of BGP updates per minute. Here the count interval is 60 seconds. A sample plot from the result is shown in Figure 6.5

6.2.3.2 BGP Unique Announcements Count

The following query is stored in *ipv6_rrd_updates_bgpcount_ua.pl*.

```
SELECT peer_id, 60*FLOOR (utc/60) AS timestamp,
COUNT (DISTINCT prefix_id) AS count, ipn AS peer_ip FROM $param {table}, peer6
WHERE peer6.id = peer_id
AND type='U' AND utc >= $param {from_time} AND utc <= $param {to_time}
```


GROUP BY timestamp, peer_id
ORDER BY timestamp, peer_id

A sample plot from the result is shown in Figure 6.6

6.2.3.3 BGP Withdrawals

The following query is stored in *ipv6_rrd_updates_bgpcount_w.pl*.

```
SELECT peer_id, 60*FLOOR (utc/60) AS timestamp,
COUNT (prefix_id) AS count, ipn AS peer_ip FROM $param {table}, peer6
WHERE peer6.id = peer_id
AND type='W' AND utc >= $param {from_time} AND utc <= $param {to_time}
GROUP BY timestamp, peer_id
ORDER BY timestamp, peer_id
```

6.2.3.4 BGP State Changes

The following query is stored in *ipv6_rrd_updates_bgpcount_sc.pl*.

```
SELECT peer_id, 60*FLOOR (utc/60) AS timestamp,
COUNT (*) AS count, ipn AS peer_ip FROM $param {table}, peer6
WHERE peer6.id = peer_id
AND type='S' AND utc >= $param {from_time} AND utc <= $param {to_time}
GROUP BY timestamp, peer_id
ORDER BY timestamp, peer_id
```

6.2.4 BGP Sample Plots

6.2.4.1 BGP Count Plot




Figure 6.5: BGP Update Count

The thick line in the plot shows the average number of BGP updates per minute. The shaded area, around the line, shows the spread around the average value. That means that the absolute minimum and maximum values, within the average interval, are used as lower and upper border. Figure 6.5 is one of the Plots which describes one peer or the whole RRC. There are four different Plots listed on the Website, created in various intervals.

Daily	3 minutes Interval
Weekly	15 minutes Interval
Monthly	90 minutes Interval
Yearly	12 hours Interval

URL: <http://www.ris.ripe.net/risreport/ipv6/rrc05/peer1/bgpcount.html>

6.2.4.2 BGP Unique Announcements Count




Figure 6.6: BGP Update Count (Unique Announcements)

Figure 6.6 shows the count of the unique announcements per minute. If a prefix is announced multiple times within a minute, it is counted as one.

URL: http://www.ris.ripe.net/risreport/ipv6/rrc05/peer1/bgpcount_ua.html

6.2.4.3 Normalized BGP Count




Figure 6.7: BGP Update Count (normalized)

This Plot shows the announcements, withdrawals and state changes normalized to the total number of updates per minute. That means the percentage of these tree

types of BGP updates. You can also see the difference between the total announcements and the unique announcements.

URL: http://www.ris.ripe.net/risreport/ipv6/rrc05/peer1/bgpcount_normalized.html

6.2.4.4 IPv6 Prefix Distribution

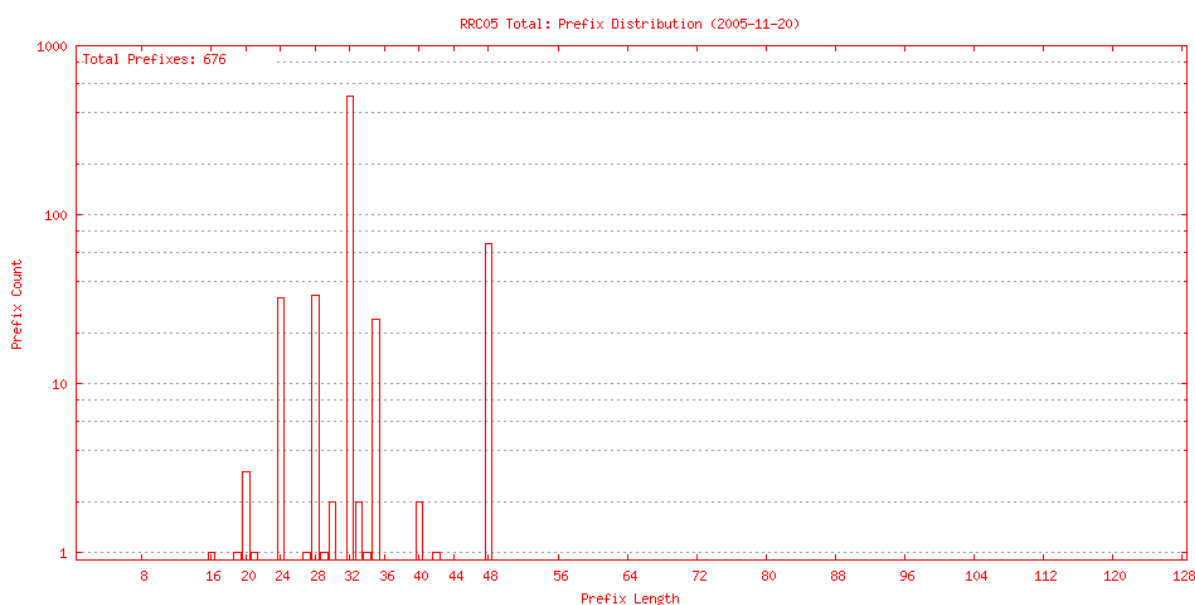


Figure 6.8: IPv6 Prefix Distribution

Figure 6.8 shows the distribution of the prefixes per prefix length seen by the RIS for each peer or the whole RRC (like in this case).

URL: http://www.ris.ripe.net/risreport/ipv6/rrc05/total/prefix_n.html

7 Conclusion

A basic subject of my internship at RIPE NCC was the development of an IPv6 version of the RISreport. This report was designed by Thomas Franchetti for his master thesis [9] to create statistical summaries of the global IPv4 routing environment, seen by the RIS. The information is presented in graphs and the results are published on the RIS website. Additional to this website, a daily email report will be generated.

The BGP data collection is done by several RRCs. This data is then inserted into the RIS database, where MYSQL is used as database server. It holds at least the BGP data of the last three month. The data in the RIS database is the source of the RIS-report. For the report, the data is queried by a universal usable tool called SQL2RRD. This program extracts the data from the database and fed it into the RRD data files. The source code for the report is written in Perl. Changes had to be done on the SQL Queries, interpretation of IPv6 Addresses and creation of the HTML files, in order to get the right output. The IPv6 version of the RISreport is now in production and the results are published on the RIS Website [13].

8 Further Work

Now the RISreport is in production and collecting data since 10.10.2005. A part of my further work will be to observe and analyse the collected date. The start of the Observation will be the early January 2006. I think this is a good starting time, because then the RISreport has enough data to enable an Observation and draw a useful conclusion for my master thesis.

List of Acronyms

AS	Autonomous System
BGP	Border Gateway Protocol
EGP	Exterior Gateway Protocol
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
IGP	Interior Gateway Protocol
IP	Internet Protocol
ERD	Entity Relationship Diagram
ISP	Internet Service Provider
LIR	Local Internet Registry
RIPE	Réseaux IP Européens
RIPE-NCC	RIPE Network Coordination Centre
RIR	Regional Internet Registry
RIS	Routing Information Service
RRD	Round Robin Database
UI	User Interface
RFC	Request for Comments

Bibliography

- [1] *RIPE* (Réseaux IP Europeans)
<http://www.ripe.net/ripe/about.html>

- [2] *DARPA* (Defence Advanced Research Projects Agency)
<http://www.darpa.mil/>

- [3] Bassam Halabi
"Internet Routing Architectures"
Cisco Systems 1997, ISBN 1-56205-652-2

- [4] S. Deering
RFC 2460 *"Internet Protocol, Version 6 (IPv6) Specification"*
<http://www.faqs.org/rfcs/rfc2460.html>

- [5] R. Hinden & S. Deering
RFC 2373 *"IP Version 6 Addressing Architecture"*
<http://www.faqs.org/rfcs/rfc2373.html>

- [6] Henk Uijterwaal & Antony Antony
"Routing Information Service R.I.S. Design Note"
Document: ripe-200.txt
<http://www.ripe.net/ripe/docs/alltitle.html>

- [7] RIPE NCC New Projects Group
"Routing Information Service"
<http://www.ripe.net/projects/ris/index.html>

- [8] *“Zebra Routing Software”*
www.zebra.org

- [9] Thomas Franchetti
Master Thesis, *“Internet Routing Analysis Provided by the RISreport”*
RIPE NNC 2001

- [10] RIPE NCC intranet
“RISng Database Overview”

- [11] James Aldridge Presentation
Megabit 2003, *“The Routing Information Service”*
<http://www.ripe.net/projects/ris/presentations/index.html>

- [12] RIPE NCC
“About the RIPE NCC”
<http://www.ripe.net/info/ncc/>

- [13] RIS Website
“RISreport for IPv6”
<http://www.ris.ripe.net/risreport/ipv6/>