BGPuma

Border Gateway Protocol Update Metric Analysis

Leigh B. Metcalf, Michael Duggan,
Mark Thomas
lbmetcalf@cert.org
NANOG 65, Montreal

.

This material is based upon work funded and supported by Department of Homeland Security under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN “AS-IS” BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Statement of Problem

So Leigh…

…Can you tell me what routing announcements and withdrawals affected this list of CIDR blocks?

…And can you make it run fast?

Solution

BGPuma!

Combines the speed of bgpdump from RIPE and SiLK from CERT to look through BGP update files quickly and find not only direct matches for CIDR blocks, but CIDR blocks that contain the initial set and are contained by the initial set.

$$\text{CIDR}_A \subseteq \text{CIDR} \subseteq \text{CIDR}_B$$

SiLK

System for Internet Level Knowledge

A collection of traffic analysis tools designed to facilitate security analysis of large networks

http://tools.netsa.cert.org/silk/index.html

SiLK

SiLK supports collecting, storage, and analysis.

It works with network flow and stores the data in its own format, working with the file system for efficient usage of space and allowing for fast queries.

It also has a Python plugin, PySiLK

Documentation:

http://tools.netsa.cert.org/silk/docs.html

SiLK Tools vs Non-specialized Tools

Note log scales

rwset

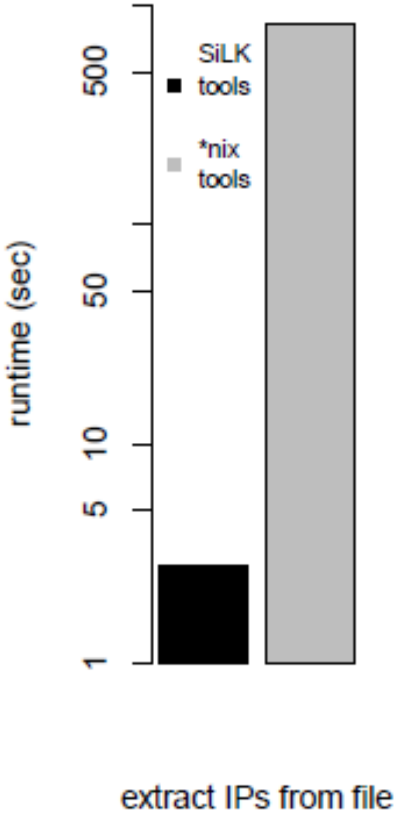

rwp2yaf2silk

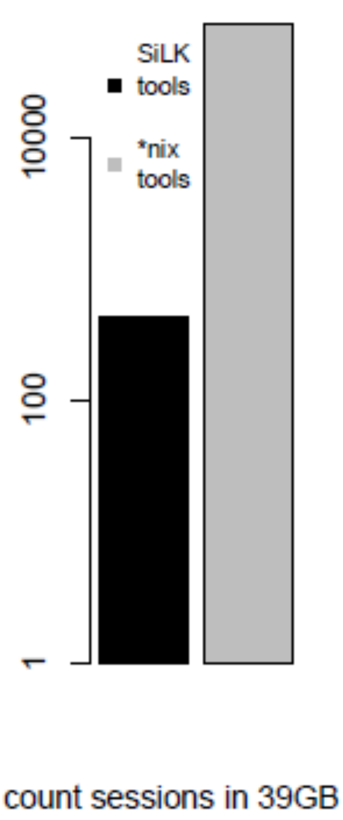

SiLK

Beyond its abilities for storing network data…

…it also contains efficient methods for storing and analyzing IP Sets

SiLK and IP Sets

The IP Sets are incredibly fast.

Name	Description	#of IPs
Odd	All Odd IPv4 addresses	2,147,483,648
Even	All Even IPv4 addresses	2,147,483,648
All	All IPv4 addresses	4,294,967,296
Rand	Randomly selected	43096
10k	Randomly selected	10,000
1M	Randomly selected	1,000,000
RFC1918	Unrouted Ips	336,723,712

Silk and IP Sets

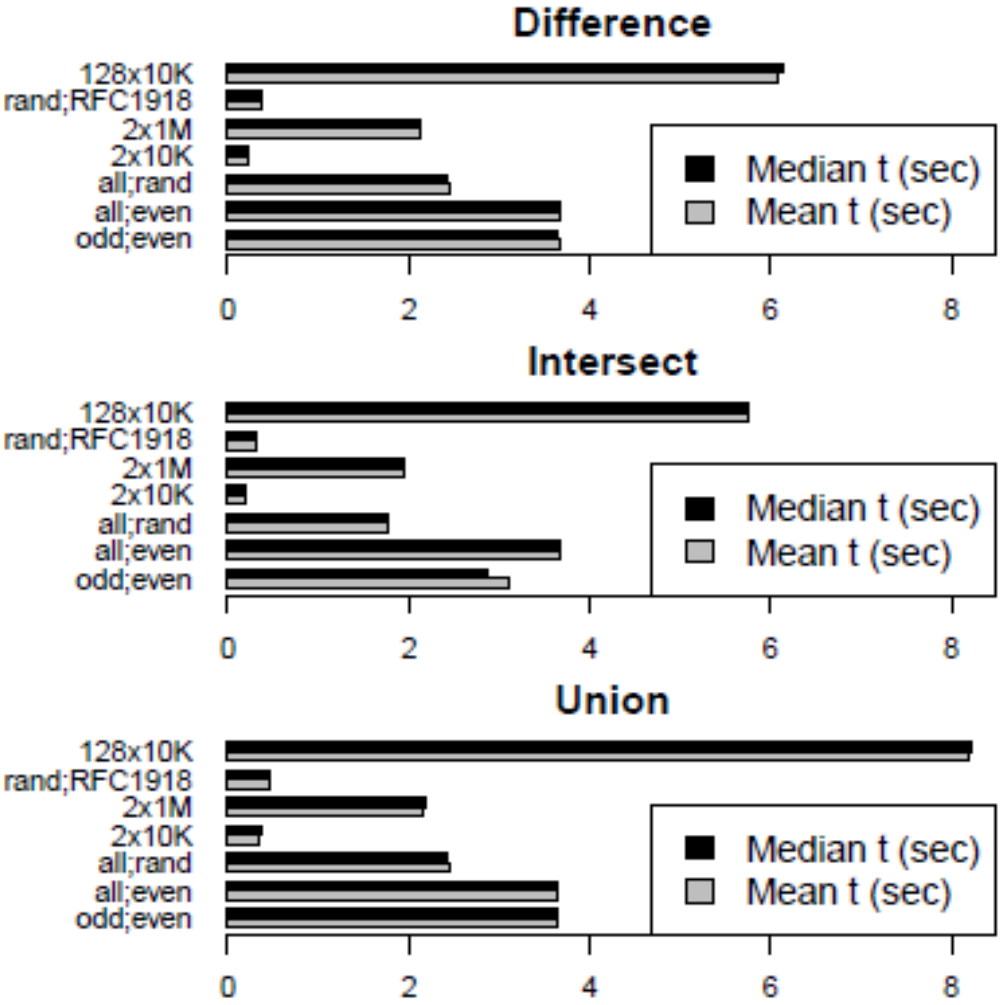

SiLK

The Sets are available by themselves, without the entire SiLK installation:

http://tools.netsa.cert.org/silk-ipset/index.html

SiLK -- libsilk

SiLK has a library, libsilk

Allows access to all of SiLK functionality

Introduced in version 0.9.5

SiLK – libsilk

Extraordinarily well documented source code:

```
/**
*  Write the IPset at 'ipset' the output stream 'stream'. 'stream'
*  should be bound to a file and open. The caller may add headers
*  to the file and set the compression method of the stream before
*  calling this function. If not set, the default compression
*  method is used. *
*  The skIPSetSave() function is a wrapper around this function.
*/
```

Bgpdump

- Library written by and supported by RIPE
- https://bitbucket.org/ripencc/bgpdump/wiki/Home
- Includes a library and a program for reading BGP update files

bgpuma

libsilk + libbgpdump = bpguma

bgpuma

First pass:

Given a list of CIDR blocks and a list of update files, find those announcements and withdrawals that could have affected that list. The –file flag contains a list of BGP update files to search and the –cidrfile contains a list of cidr blocks that you are looking for.

```
bgpuma –file=FILE –cidrfile=FILE
```

bgpuma

For example, if I have 0.0.0.0/8, 127.0.0.0/8 and 10.0.0/8 in one file and a list of update files downloaded from RouteViews in another, the results would include:

```
1393657324|A|206.126.236.120|AS41095|206.126.236.142|AS6447|10.12.4.106/32|41095
1393656402|A|206.126.236.120|AS41095|206.126.236.142|AS6447|10.11.10.0/24|41095 65000
```

bgpuma

It also understands the directory structure from the RouteViews repository and the RIPE routing repository. So if you have those available to mount as file systems in some way, you can give:

A startdate YYYYMMDD

An enddate YYYYMMDD

One or more directories that contain the repositories.

bgpuma output

```
1422812814|W|193.203.0.134|AS39912|193.203.0.123|AS12654|
10.10.25.0/24

14.22812814|W|193.203.0.134|AS39912|193.203.0.123|AS12654|
10.10.30.0/24

1422815529|A|193.203.0.134|AS39912|193.203.0.123|AS12654|
10.10.25.0/24|39912 8513

1422815529|A|193.203.0.134|AS39912|193.203.0.123|AS12654|
10.10.30.0/24|39912 8513

1422826126|A|193.203.0.134|AS39912|193.203.0.123|AS12654|
10.10.25.0/24|39912 8513

1422826126|A|193.203.0.134|AS39912|193.203.0.123|AS12654|
10.10.30.0/24|39912 8513

1422815529|A|206.126.236.24|AS11666|206.126.236.142|AS6447|
10.10.25.0/24|11666
```

bgpuma output decoded

1. Time stamp
2. A (Announcement) or W (Withdrawal)
3. Source IP of the data
4. Source ASN of the data
5. Destination IP of the data
6. Destination ASN of the data
7. CIDR block
8. If an Announcement, the actual path

bgpuma

bgpuma will also look for Autonomous Systems

Put the Autonomous Systems in a file, one per line.

```
 bgpuma –asnfile=FILE
```

bgpuma

The asnfile and cidrfile options work together and find all results by default that match both sets.

If you’d rather find all results that match either set, use:

```
--or
```

bgpuma

Other options:

--outfile FILE

File to put the output from bgpuma in. It defaults to standard out.

--help --verbose

Prints out each file it analyzes

bgpuma

Speed!

I took the worst case for a particular day, a single file that was 40M.

bgpdump: Average time – 39.032s

bgpuma: Average time – 53.805s

Over 5 runs of each.

bgpuma

Using RIPE data from 20150201 the average speed for bgpuma was 5m27.67s to run bgpuma over the entire day.

bgpuma

Currently it outputs everything it finds.

Future plans:
Summary options! (What should I summarize?)

bgpuma

Most important!

It is available at github:

https://github.com/cmu-sei/bgpuma

CERT

Questions/comments?