BGP, some Python, and a DB

Theodore Baschak

BSides Winnipeg 2015-11-15

Slides

Reveal.js slides
ciscodude.net/presentations/bsideswpg2015-mb-bgplogger
Reveal.js PDF slides
ciscodude.net/presentations/bsideswpg2015-mb-bgplogger.pdf

exabgp-logger daemon project at
github.com/tbaschak/exabgp-logger

Who Is Theo?

- Network Architect @ Daemon Defense Systems
- Avid Open Source Software user/fanatic, and contributor
- New form of addiction: Border Gateway Protocol
- Obsessed with network monitoring and routing
- Often written about on ciscodude.net
- Involved with several not-for-profits in Winnipeg
  1. Board of the Manitoba Internet Exchange (MBIX)
  2. Board of Coldhak
     - Coldhak is a not-for-profit dedicated to furthering privacy, security and freedom of speech
     - Coldhak runs a handfull of Tor relays and exits
     - Coldhak also produces coldkernel, grsec-enabled kernel for Debian/Ubuntu

Run a Tor Relay

- If you've got bandwidth and IP address(es) to spare, Coldhak would love to run/manage a Tor relay/exit for you.
  - coldhak.ca
  - twitter.com/coldhakca

Inspiration

- Talked about BGP hijacking in my BSidesWpg 2013 talk
- Dyn Research blogs, formerly Renesys Corporation:
  1. 2 days after BSidesWpg 2013, 2013-11-19, Targeted Internet Traffic Misdirection
  2. 2015-03, UK traffic diverted through Ukraine
- BGPmon & BGP Stream

Blizzard

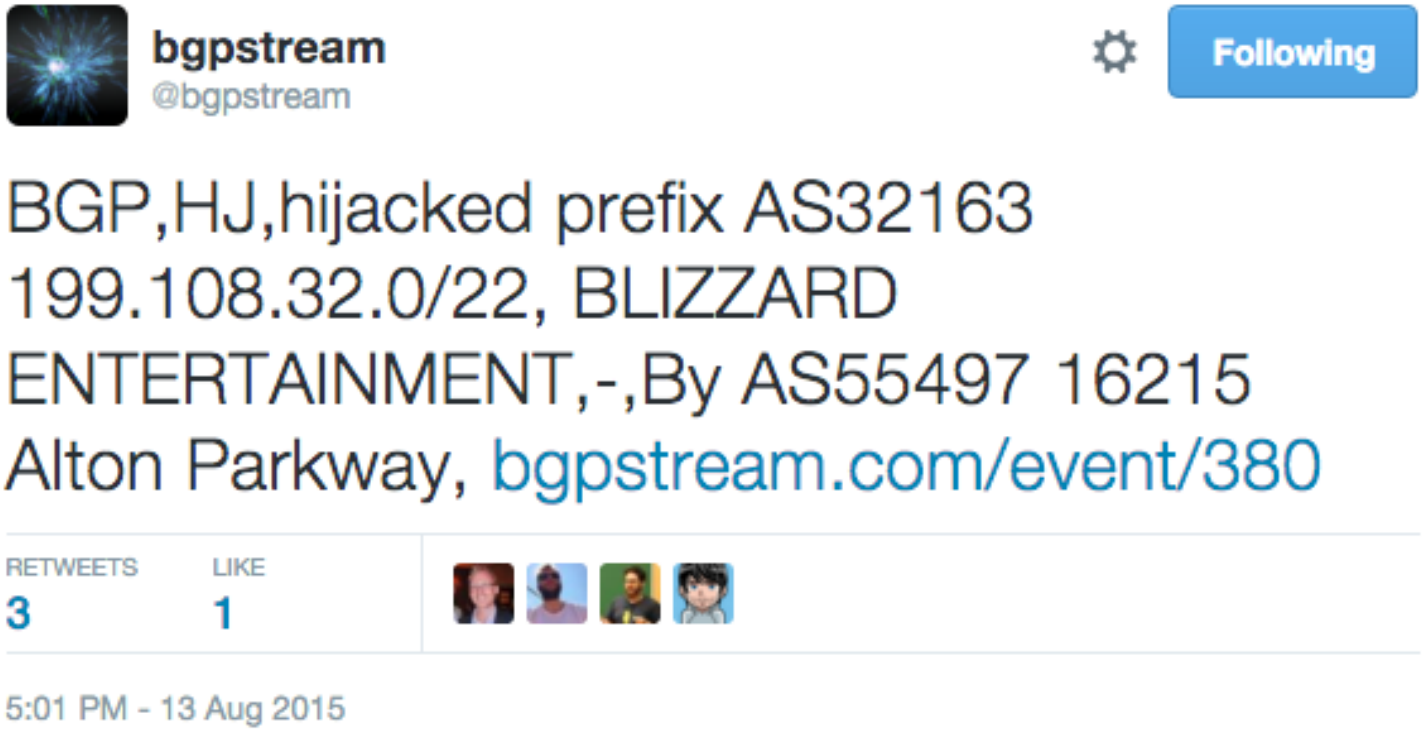

Figure 1: Blizzard IP space Hijack, 2015-08-13

- Tweet
- BGP Stream
- False positive after investigation (Asia Pacific region AS)

hmmmmmmm

Figure 2:

The Idea

- To log BGP updates from various Manitoban sources
  - exabgp seemed like a good starting place
  - a mailing list post w/ a simple shell collector script
- To examine routes for peering relationships between Manitoban ASNs
- To look at the effects of BGP leaks on Manitoban routing

Components

- Exabgp (the BGP, and the Python)
  - one process, per AF, per peer (v4 and v6 separate)
  - outputs BGP updates as JSON
- Couchdb (the DB)
  - stores JSON objects
  - one database per peer (v4 and v6 combined)
  - replicate each peer into common DB

Building

- Uses exabgp
  - Exabgp config defines actions for route input and output
  - Just input in this case, not advertising anything
  - Bash while loop to read line by line and POST to couchdb

exabgp-logger

- Able to run without a peer for basic config checks
- Needed a live BGP peer to start viewing BGP update record format – which turned out to be beautiful JSON
- First tests
  - MFNERC, alpha testing only – 8 days, IPv4 only, 1 million updates
  - First version just cat appended the JSON to a flat text file so I could look at the records
- Growing Pains

Use Cases

- Personal Interest
- Single-AS route logger
  - Security conscious enterprises
- IXP route logger

Peering

0. MFNERC/AS62758
   - v4: Mon, 07 Sep 2015 18:29:28 GMT
1. MERLIN/AS16796
   - v4: Mon, 12 Oct 2015 21:52:46 GMT
   - v6: Mon, 12 Oct 2015 22:24:07 GMT
2. 3T Systems/AS20291
   - v6: Thu, 15 Oct 2015 22:04:15 GMT
   - v4: Thu, 15 Oct 2015 22:05:54 GMT
3. Les.net/AS18451
   - v4: Tue, 20 Oct 2015 19:44:29 GMT
   - v6: Tue, 20 Oct 2015 19:44:52 GMT
4. Swift High Speed Internet/AS393445
   - v4: Thu, 05 Nov 2015 05:15:12 GMT

Manitoban Routing

- Many common carriers
  - AS6327, AS6939, AS7122, AS10965, AS14866, AS18451
  - Some locally connected
  - Links between these out of the province; Calgary, Toronto, Chicago
- Other carriers
  - some less common picked up in Toronto, Alberta and Chicago

Scaling Up

- Moved VM
- Offsite replication
- Increased size of VM and went 64bit
  - Couchdb now using 16GB of memory
  - Exabgp using 55MB memory per process, 385MB total

DB Size

- 17 million records after just over a month
- 20-40GB diskspace used

Name	Size	Number of Documents	Update Seq
_replicator	4.1 KB	1	1
_users	4.1 KB	1	1
bgplog	11.6 GB	16886335	16886357
les	2.4 GB	3257077	3257083
merlin	3.8 GB	5336847	5336847
mfnerc	0.6 GB	1001381	1001381
swift	1.5 GB	1446303	1446303
threetsystems	4.2 GB	5844730	5844738

Showing 1-8 of 8 databases ← Previous Page | Rows per page: 10 | Next Page →

Figure 3: Database sizes

Finding Events

- Currently, create couchdb views to search for as-paths
  - Can take a while to make indexes
  - Validated several events from BGP Stream
- Searching by CIDR not efficient

Automating

- Create a RIB for each peer & AF
- Need to compare updates with each RIB
  - Prefix length
  - AS-PATH
- Need ability to replay JSON BGP updates from couchdb
  - Would make a useful lab testing tool as well
- False postives

Thanks

- MERLIN, Jared and his management who supported the project and idea
- 3T Systems, Jordan
- Les.net, Les and Jonathan
- Swift High Speed Internet, Evan
- BSides

Questions

Reveal.js slides
ciscodude.net/presentations/bsideswpg2015-mb-bgplogger
Reveal.js PDF slides
ciscodude.net/presentations/bsideswpg2015-mb-bgplogger.pdf

exabgp-logger daemon project at
github.com/tbaschak/exabgp-logger