

A study of BGP misconfiguration

Ratul Mahajan
David Wetherall
Tom Anderson

University of Washington

Motivation

- ◆ BGP is considered the weakest link in Internet infrastructure
 - Widespread impact
 - Enough anecdotal evidence about BGP screw-ups
 - Very little hard data
- ◆ Exactly how bad is it?
 - An attempt to attach some numbers to the malaise feeling
- ◆ What are the underlying causes?

Misconfiguration

- ◆ Configuration ¹ intent
- ◆ Configuration in violation of specifications

Prefix AS_k, AS_{k-1},, AS₀

- Invalid prefix, ASN, AS-path
- Origin misconfiguration
- Route export misconfiguration

Methodology

- ◆ BGP tables from RouteViews
 - 60+ peers, 40+ AS's
 - 2 hourly snapshots
- ◆ Use heuristics to identify potentially bad announcements
- ◆ Validate using an email survey of ISP's

Origin misconfiguration

- ◆ Accidental origination of routes
 - insertion (propagation) of routes into global tables
 - address space hijacks
- ◆ Identify short-lived announcements in the table
 - with no past or future (weed out failures)
 - both new prefixes, and new origins for a prefix

Definitions of short-lived changes

	Stable Announcements		Short-lived Announcements	
Self Deaggregation	a.b.0.0/16	X-Y-Z	a.b.c.0/24	X`-Y`-Z
Stripping	a.b.0.0/16	X-Y-Z	a.b.0.0/16 a.b.0.0/16	X`-Y` X`-Y`-Z-O
Strip Deaggregation	a.b.0.0/16	X-Y-Z	a.b.c1.0/24 a.b.c2.0/24	X`-Y` X`-Y`-Z-O
Foreign Origination	a.b.0.0/16	X-Y-Z	a.b.0.0/16	X`-Y`-O
Foreign Deaggregation	a.b.0.0/16	X-Y-Z	a.b.c1.0/24	X`-Y`-O

Results (7 days)

	Total	Replies	Misconfigs	Connectivity	False +ve
Self Deaggragation	1979	902 (45%)	822 (91%)	39 (4.3, 4.7%)	80 (9%)
Stripping	797	738 (92%)	731 (99%)	6 (8.1, 8.2%)	7 (1%)
Strip Deaggregation	798	464 (58%)	446 (96%)	14 (3.0, 3.1%)	18 (4%)
Foreign Origination	180	62 (34%)	45 (72%)	24 (39, 53%)	17 (28%)
Foreign Deaggregation	313	101 (32%)	89 (88%)	35 (34, 39%)	12 (12%)
All	4067	2267 (56%)	2133 (94%)	118 (5.2, 5.5%)	134 (6%)

What's going on?

◆ Misconfigurations

- Buggy filters
 - ACL's, route-maps
- Redistribution
- Old Router/Config
- Hijacking
- Forgetting auto-summary
- Incorrect Summarization
- Communities
- Relying on upstream

◆ False +ves

- Testing
- Load Balancing
- Migration
- Failures

Route export misconfiguration

- ◆ Paths that are in violation of an AS's commercial policies (customer, peer, provider,...)
 - e.g., a customer providing transit between its two providers
- ◆ AS relationships are closely guarded
 - get a lower bound
 - understand why it happens

Methodology

- ◆ Use Gao's analysis to estimate AS relationships
 - paths are "valley free"
 - providers are more likely to have higher degrees
- ◆ Identify AS-sequences that
 - don't obey those relationships
 - are not long-lived

Results (6 days)

	Total	Replies	Misconfigs	False +ve
Incidents	466	149 (32%)	133 (89%)	16 (11%)
Prefixes	8111	6871 (85%)	6822 (99%)	49 (1%)

What's going on?


◆ Misconfigurations

- prefix based configuration
- typos
- community

◆ False +ves

- alternate networks
- backup paths
- special arrangements

Configuration at *C*:
announce prefixes
from *CC* to *P1* & *P2*



On email validation

- ◆ 30% of emails never make it
- ◆ I have placed you in my /dev/null filter, goodbye...
- ◆ Don't worry. That was a configuration error of our upstream ISP.
- ◆ I think although i don't have much time to do R&D on BGP, i do believe every secs there are routes appear and disappear from internet routing table.
- ◆ Yes, we know this is not a recommended way of doing things; but the packet monster of the internet must be fed.
- ◆ Hope you enjoy living in Seattle; it's a beautiful city.
- ◆ I am writing to thank you for your letter and say that I am glad that someone apart from me is interested in our BGP announcements.
- ◆ Despite our complaints and RFO, further clarification was not given by AS-XXXX. Such is life of a small ISP barking like a Chihuahua at 'big' AS-XXXX

Speculative conclusions

- ◆ BGP is a rich ground for misconfigurations
- ◆ The causes are diverse
 - configuration seems too complex
- ◆ Connectivity is robust to most misconfiguration incidents
- ◆ There are research-friendly network operators out there

Feedback:

<http://www.cs.washington.edu/homes/ratul/bgp>

Acknowledgments

- ◆ Operations community
- ◆ RouteViews (David Meyer)
- ◆ Skitter project (CAIDA)
- ◆ Folks at UW-CSE support and UW-C&C
- ◆ Lixin Gao, Ramesh Govindan, Brad Volz,