

Data Analysis and Design of a BGP Monitoring and Alarm System

Gunnar Bornemann

Diploma Thesis

Research Unit VIII - Network Architectures


Technische Universität München

borneman@net.in.tum.de

23.05.2007


Contents

- Motivation
- Background: Data Collection
- Step 1: Data Analysis
- Step 2: Conclusions
- Step 3: BGP Monitoring and Alarm System Design
- Testing
- Future Work




Contents

- **Motivation**
- Background: Data Collection
- Step 1: Data Analysis
- Step 2: Conclusions
- Step 3: BGP Monitoring and Alarm System Design
- Testing
- Future Work




Motivation

- Routing is an important part in the Internet
 - Routing can become very complex...
 - ... especially when a lot of routers and/or neighbor networks are involved
- Design a system to aid network administrators in managing routing
- Data analysis needed first!




Research Aspects

- Previous approaches: Collect data from the *whole* Internet, results provide information about certain networks
- Problems:
 - Huge amount of data needs to be processed
 - Results may not be precise without inside knowledge of a network




Research Aspects

- Different approach:
 - Start by analyzing data from a *single* network
 - Learn what information can be abstracted
 - Add data containing an outside view of the network
 - Learn how the outside view complements the inside view




Benefits for Network Administrators

- Provide automated way of monitoring (many) routers (in large networks)
- Detect abnormal or critical situations
- ISP: Monitor customer influence on BGP




Contents

- Motivation
- **Background: Data Collection**
- Step 1: Data Analysis
- Step 2: Conclusions
- Step 3: BGP Monitoring and Alarm System Design
- Testing
- Future Work


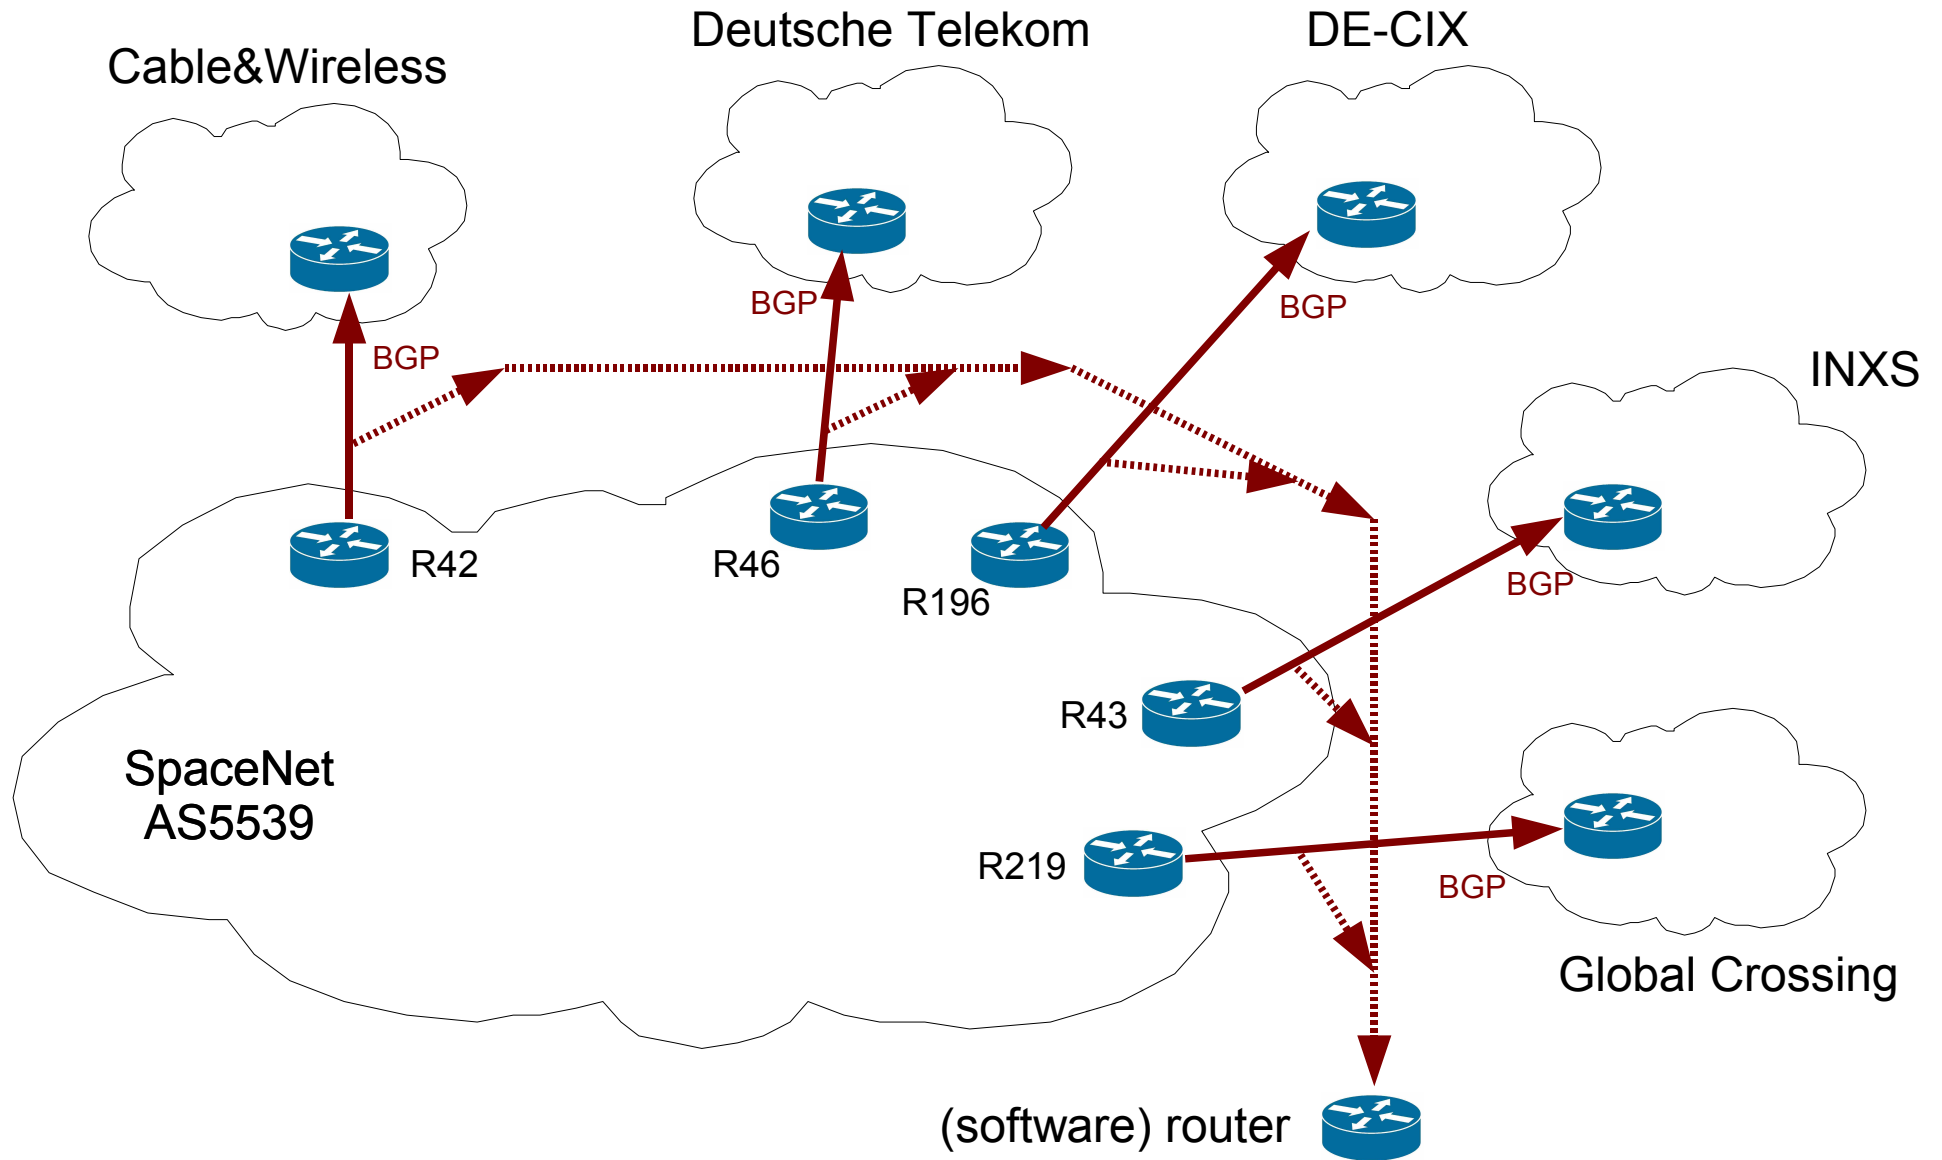


Data Collection Points

- Internal collection points:
 - Border routers providing BGP updates to their neighbors
 - Software router receives copy of each BGP update
 - Do not interfere with normal BGP operation!
 - Export BGP updates to disk




Background




Data Collection Points

- External collection points:
 - RIPE Routing Information Service
 - Similar setup as internal collection
 - Software routers collect data from many networks
 - Data offered on RIPE home page
 - Download data to disk




Background

Software Router	Location	# IPv4 Full Feeds	# IPv6 Full Feeds	# Peerings
RRC01	London	11	5	110
RRC02	Paris	0	0	25
RRC03	Amsterdam	16	11	141
RRC10	Milan	2	1	19
RRC11	New York	8	3	38
RRC12	Frankfurt	11	13	73



Contents


- Motivation
- Background: Data Collection
- **Step 1: Data Analysis**
- Step 2: Conclusions
- Step 3: BGP Monitoring and Alarm System Design
- Testing
- Future Work




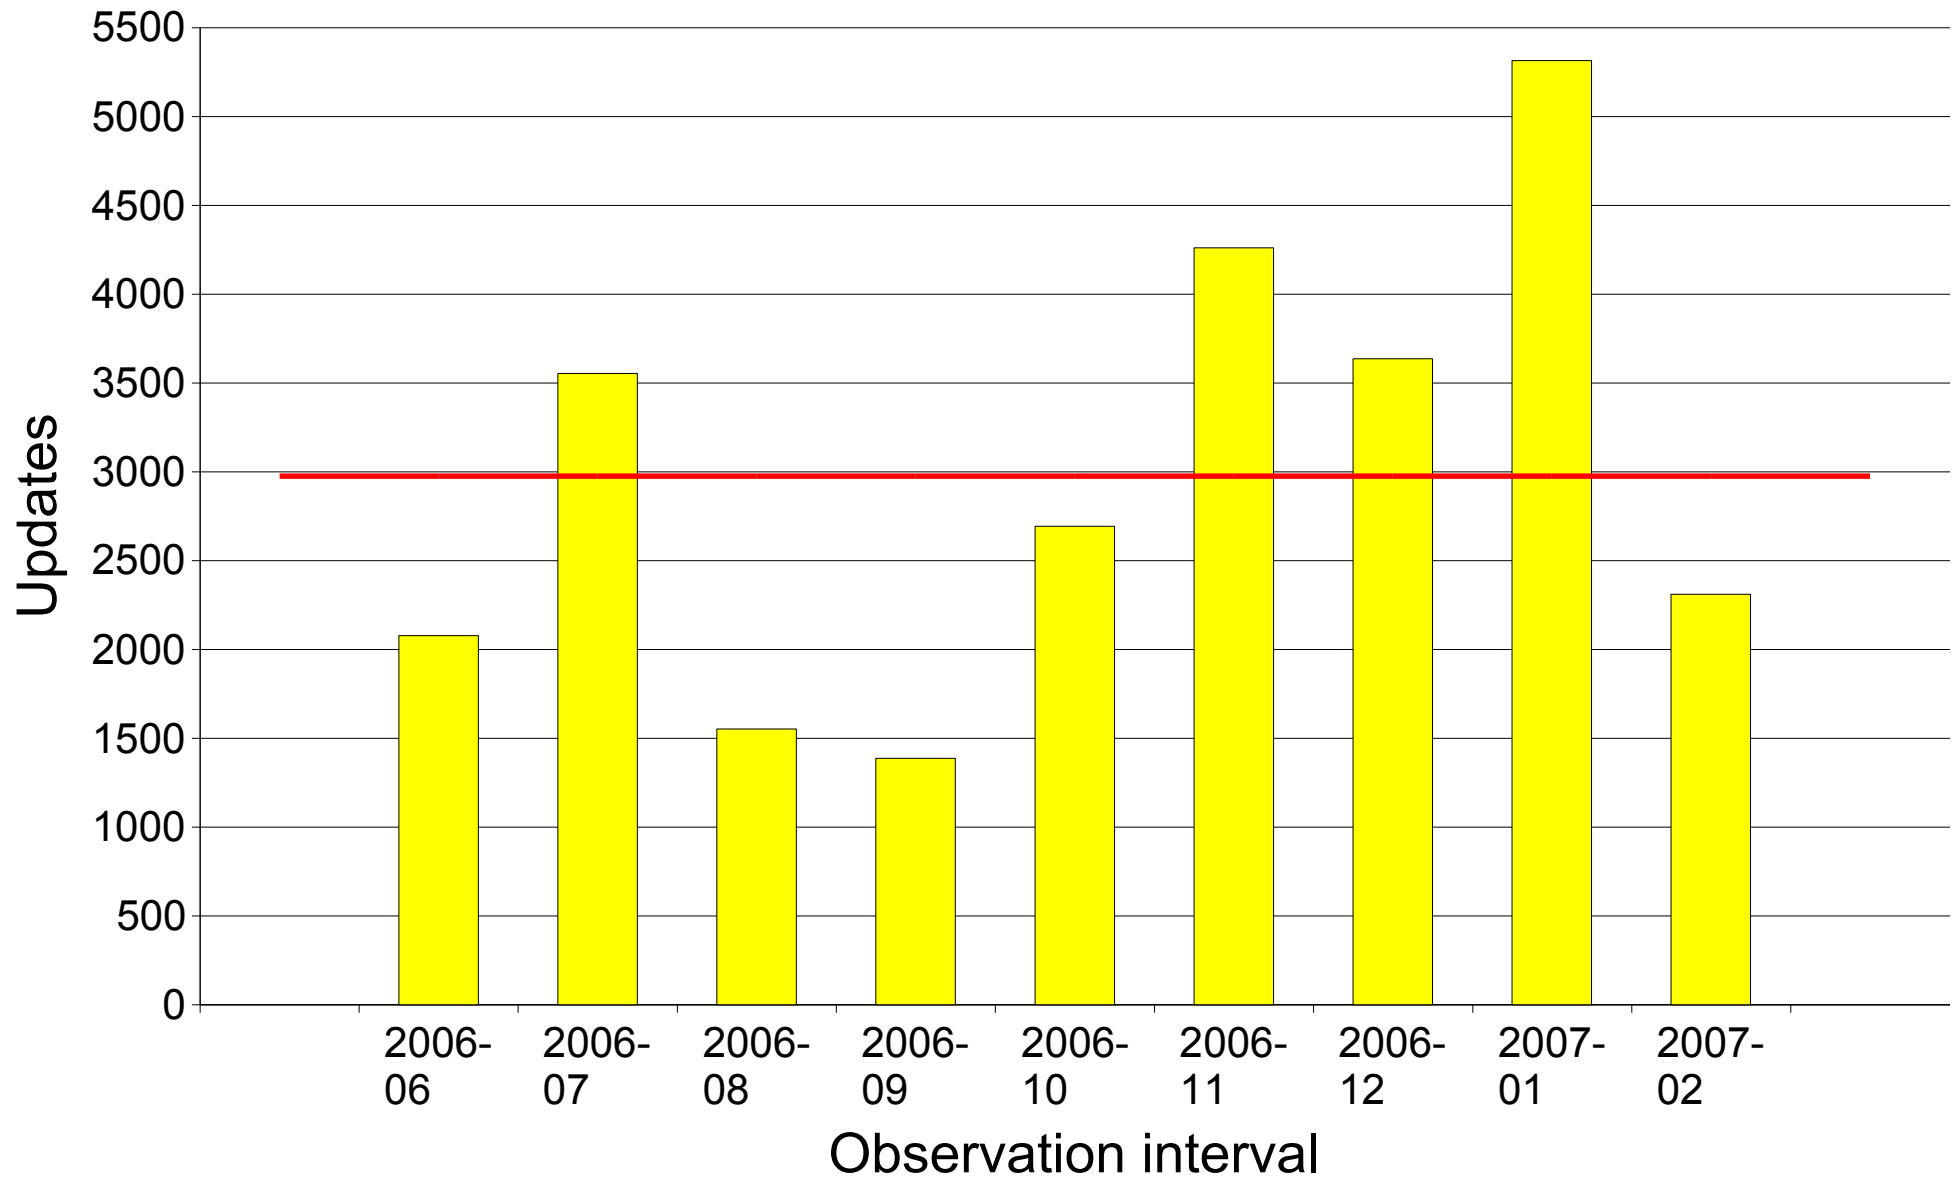
Step 1: Data Analysis

“How much data do we receive?”

- Internal collection points:
 - Total of 100 prefixes
 - Average of 3000 updates per month
 - Average of 100 updates per day




Step 1: Data Analysis



Step 1: Data Analysis


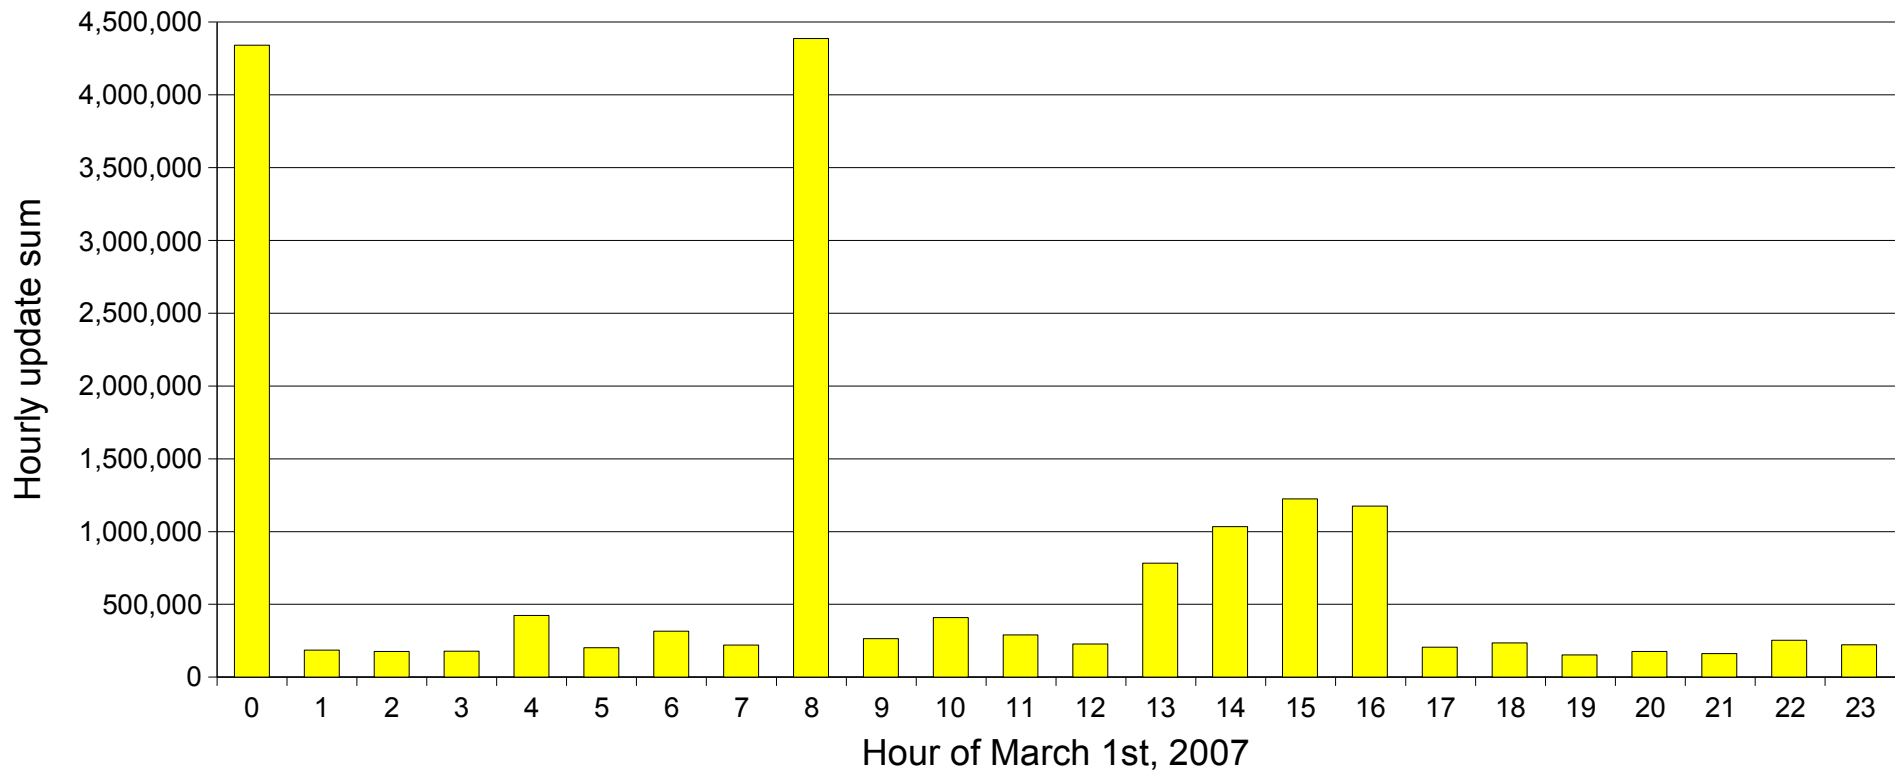
“How much data do we receive?”

- External collection points:
 - Depends on software router
 - Depends on type of download (full table dump or incremental update)
 - Incremental updates: Depends on time of day!



Step 1: Data Analysis


Software Router	Full Table Dump		Incremental Update (avg)	
	Download Size	# Updates	Download Size	# Updates
RRC03	37.4 MB	4 million	350 KB	60000
RRC10	4.5 MB	500000	14 KB	3000



Step 1: Data Analysis


Further Observations

- Not all SpaceNet prefixes show up in RIPE data
- ASN attribute of internal data is always 5539
- Only one distinct LOCAL-PREF value used by SpaceNet routers (SpaceNet dislikes usage of this attribute)
- MED attribute not useful (no two routers monitored at SpaceNet have the same neighbor)




Contents

- Motivation
- Background: Data Collection
- Step 1: Data Analysis
- **Step 2: Conclusions**
- Step 3: BGP Monitoring and Alarm System Design
- Testing
- Future Work




Step 2: Conclusions

- Amount of data definitely requires automation
- Internal prefixes not found in external data require additional processing (check for aggregation)
- Amount of internal data does not pose processing problems...
- ... but aggregation checks on external data could




Step 2: Conclusions

- Because of aggregation the AS-PATH attribute in external data may be incomplete
- COMMUNITY attribute in external data may be heavily filtered or empty
- BGP attributes needed:
 - TIMESTAMP, ROUTER, PREFIX
 - From internal data: COMMUNITY
 - From external data: AS-PATH



Contents


- Motivation
- Background: Data Collection
- Step 1: Data Analysis
- Step 2: Conclusions
- **Step 3: BGP Monitoring and Alarm System Design**
- Testing
- Future Work



Step 3: Design

Goals


- Run as close to real-time as possible
 - RIPE software routers export at 5 minute intervals
 - Our software router is freely configurable
 - Run every 5 minutes
- Easy to access data store
 - Make present work easier
 - Allow future work
 - SQL database (MySQL)



Step 3: Design

Goals


- Inform administrator quickly
 - Do not require a special tool
 - Messages should be easy and fast to read and understand
 - Simple e-mail
 - Alarm message in subject



Step 3: Design

Periodic Program Run (every 5 minutes)


- Import part
 - For RIPE: Download latest data to disk
 - Convert data from binary to ASCII format
 - Parse data and write to database
- Processing part
 - Merge changes with current status of routers and prefixes



Step 3: Design

Periodic Program Run (every 5 minutes)


- Processing part
 - Check for defined “interesting” changes or situations
 - Depending on interest:
 - Write *log* about change or situation to database
 - Send *alarm* to administrator via e-mail



Step 3: Design

Daily Program Run (shortly after midnight)

- Build statistics (e.g. calculate router update averages)
- Read logged information from database
- Send daily summary via e-mail
- Cleanup (e.g. delete “old” routers and prefixes)




Step 3: Design

Logging vs. Alarming

Log:


- Event or situation of interest
 - New router
 - New prefix
 - Router has session reset
 - Changes in COMMUNITY attribute of a prefix
 - All alarms!



Step 3: Design


Alarm:

- Critical or abnormal for BGP operation
- Alarms are *always* logged
- Alarms are sent via e-mail only *once per day*
- Router has too many session resets
- Router is sending too many BGP updates
- Router's last session reset is too old
- Router is sending repeating updates
- Router is sending alternating updates
- *Internal* prefix is not seen by *external* software routers



Contents

- Motivation
- Background: Data Collection
- Step 1: Data Analysis
- Step 2: Conclusions
- Step 3: BGP Monitoring and Alarm System Design
- **Testing**
- Future Work




Facts:

- Internal data collection since May 24th, 2006 (One year)
- Almost 80000 updates collected from internal routers

Our system reacted during this period to:

- Session resets
- Alternating and repeating updates
- Prefixes unobservable in the external data
- ...




E-Mail Examples

Subject: bgpmon: prefix 2001:868::/32 has alternating updates
Date: Wed, 23 May 2007 13:15:17 +0200
From: Gunnar Bornemann <borneman@net.informatik.tu-muenchen.de>
To: borneman@net.in.tum.de, wolfgang@net.in.tum.de

FYI

Subject: bgpmon: prefix 217.25.64.0/20 not on collectors
Date: Sat, 31 Mar 2007 06:40:29 +0200
To: borneman@net.in.tum.de, wolfgang@net.in.tum.de
From: Gunnar Bornemann <borneman@net.informatik.tu-muenchen.de>

FYI




E-Mail Examples

Subject: bgpmon: daily summary
Date: Fri, 20 Apr 2007 00:04:01 +0200
From: Gunnar Bornemann <borneman@net.informatik.tu-muenchen.de>
To: borneman@net.in.tum.de, wolfgang@net.in.tum.de

timeline

=====

08:20:17 router 193.149.44.43 exceeded update count threshold
08:20:17 router 193.149.44.219 exceeded update count threshold
13:10:21 router 193.149.44.42 exceeded update count threshold
14:15:17 prefix 2001:bf0::/32 on router 193.149.44.43 has alternating updates
14:15:17 prefix 2001:bf0::/32 on router 193.149.44.42 has alternating updates
18:35:17 prefix 2a01:78::/32 on router 193.149.44.42 has alternating updates
18:35:17 prefix 2a01:78::/32 on router 193.149.44.43 has alternating updates
18:35:17 prefix 2a01:78::/32 on router 193.149.44.219 has alternating updates
19:50:17 new prefix 2001:14e0::/32 on router 193.149.44.42
19:50:17 community changes for prefix 2001:14e0::/32 on router 193.149.44.42: 1239:90(+)
1273:30651(+) 1273:39951(+) 5539:500(+)
19:50:17 new prefix 2001:14e0::/32 on router 193.149.44.43
19:50:17 community changes for prefix 2001:14e0::/32 on router 193.149.44.43: 5539:500(+)
19:50:17 new prefix 2001:14e0::/32 on router 193.149.44.219
19:50:17 community changes for prefix 2001:14e0::/32 on router 193.149.44.219:
3549:8243(+) 5539:500(+)



E-Mail Examples

router status

=====


193.149.44.196: prefixes	0 updates (0.0/day)	0 session resets (last 2007-02-07 22:20:16)	46
193.149.44.219: prefixes	96 updates (10.5/day)	0 session resets (last 2007-03-27 08:15:17)	44
193.149.44.42: prefixes	96 updates (16.0/day)	0 session resets (last 2007-03-27 08:15:17)	67
193.149.44.43: prefixes	96 updates (10.4/day)	0 session resets (last 2007-03-27 08:15:17)	44
193.149.44.46: prefixes	4 updates (1.1/day)	0 session resets (last 2007-03-27 08:15:17)	25

prefix status

=====


62.208.31.12/30:	on 1 router	0-0 (0-17) at collectors
62.208.31.16/30:	on 1 router	0-0 (0-17) at collectors
81.91.160.0/20:	on 5 routers	5-18 (0-0) at collectors
96.0.16.0/20:	on 5 routers	20-0 (0-0) at collectors
134.247.0.0/16:	on 5 routers	11-11 (0-0) at collectors
193.97.129.0/24:	on 5 routers	21-0 (0-0) at collectors
193.149.32.0/19:	on 5 routers	21-0 (0-0) at collectors

[...]



Contents


- Motivation
- Background: Data Collection
- Step 1: Data Analysis
- Step 2: Conclusions
- Step 3: BGP Monitoring and Alarm System Design
- Testing
- **Future Work**



Future Work

- Improve functionality:
 - Add more internal collection points (especially *all* border routers)
 - Add more external collection points (more RIPE software routers, RouteViews, etc.)
 - Session reset detection

- Extend functionality:
 - On-demand view of system state
 - Detect origin of routing instabilities



Thank you for your attention!

