BGP Data Collection and Organization (OBGP) Tool

Colorado State University
University of Arizona
UCLA

BGP Data Collection

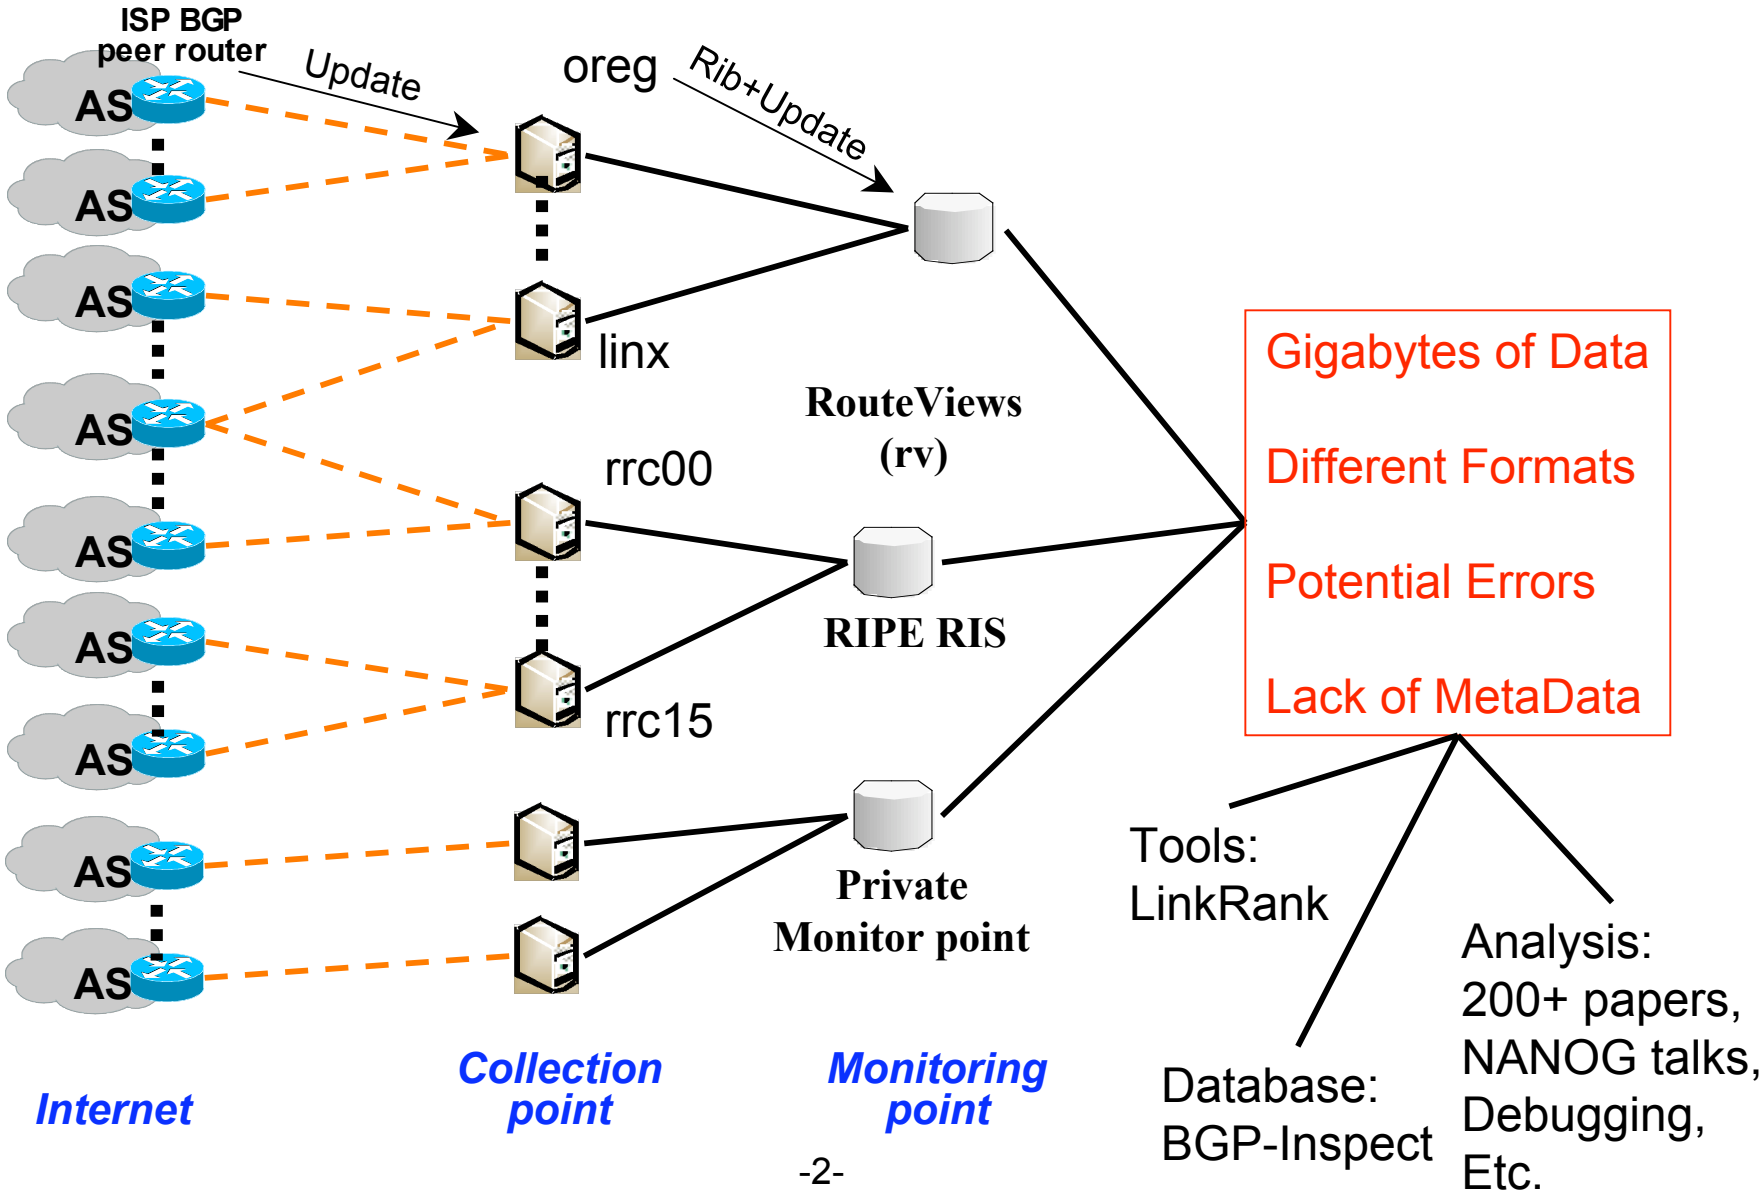

OBGP Motivation

- Large Volume of Data
  - Data from many sources (RIPE, RV, private data)
  - Long time scales and very recent (real-time?) data
- Slightly Different Formats
  - RIPE/RV use different naming conventions
  - Different dump intervals
  - Different time zones (for older data)
- Lack of MetaData
  - Would like to only desired peers and desired update types
- Possible Errors in the Data
  - Are updates missing due to (log??) errors?
  - What is lost due to session failures?

OBGP Overview

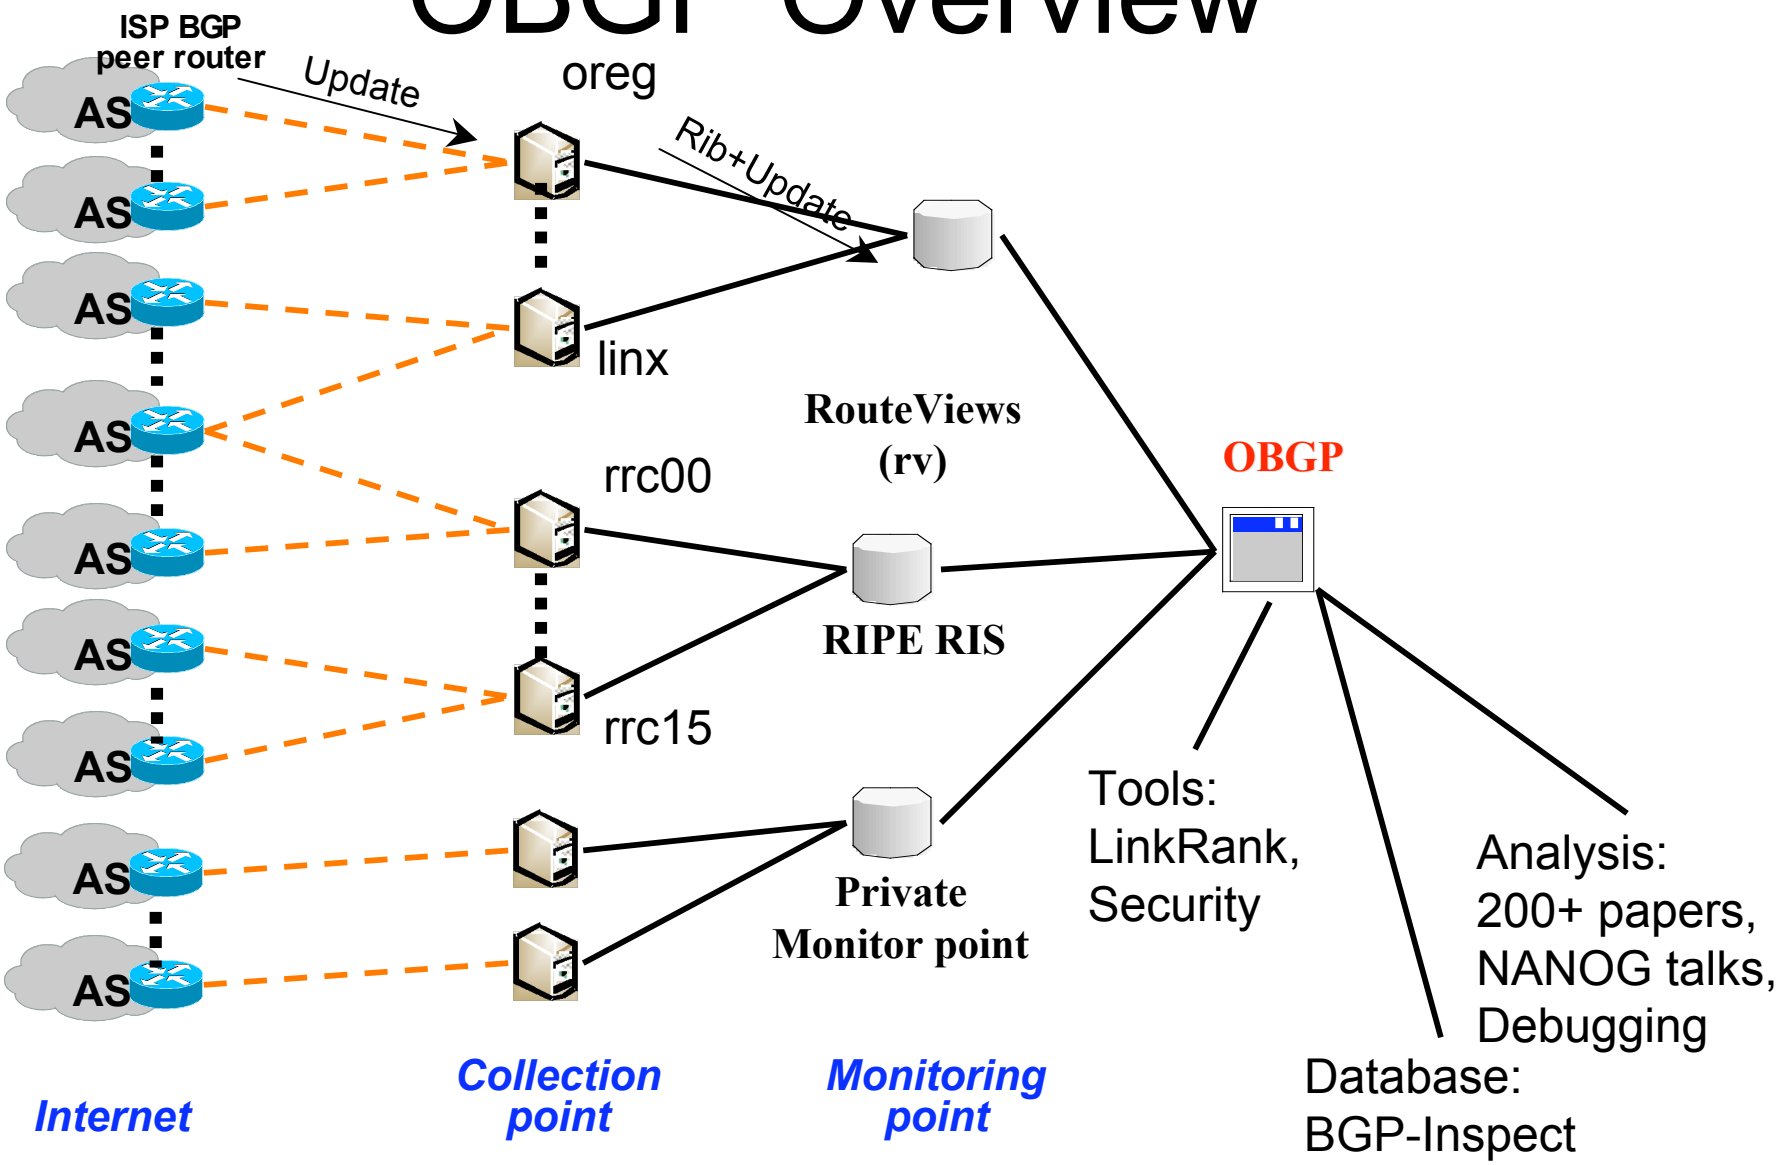

OBGP Features

- Uniform Data Organization
  - Consistent and easy to use for scripts
- Consistent View of Multiple Monitoring Points
- Annotations/Labels
  - Easy access to custom views of the large data
- Table Transfer Detection
  - Distinguish updates from data collection peering
- Data Inconsistency Detection and Correction
  - Understand and fix (??) possible data errors

Uniform Data Organization

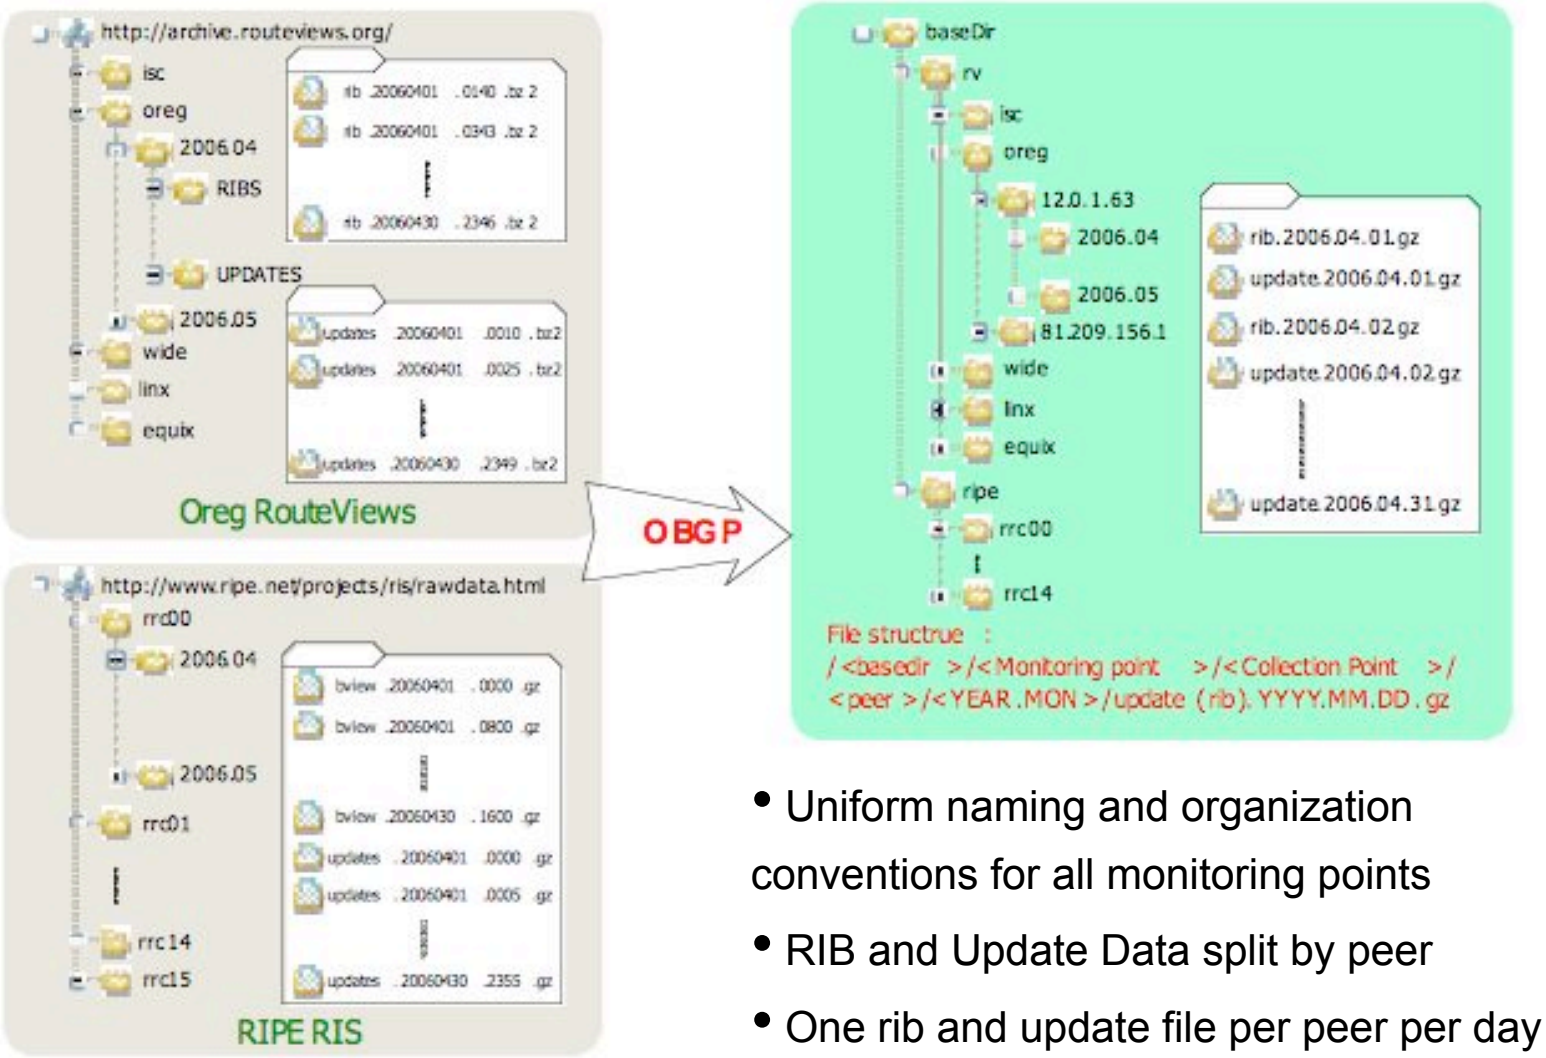

- Uniform naming and organization conventions for all monitoring points
- RIB and Update Data split by peer
- One rib and update file per peer per day

Labels And Annotations

- Existing Format Labels Updates As:
  - Announce (A) or Withdraw (W)
  - Also includes some STATE messages
- OBGP Enhances the Labels
  - Adds a Status Message
  - Adds an Update Type
  - More STATE Messages
    - Route table dump
    - Table Transfers

```
BGP4MP|1136076348|RTS|…|
BGP4MP|1136076350|E:RIB:NANN| …
BGP4MP|1136076365|RTE|…|
BGP4MP|1136073679|A:INC:DPATH| …
```

OBGP Added Labels

|<original update type>:<status information>:<OBGP update type>|

- <original update type>
  - Add E for Error Correction
  - (more on errors soon….)
- <status information>
  - INC incremental update
  - TT table transfer update
  - RIB: correction update
- <OBGP update type>
  - New Announcement
  - Duplicate Announcement
  - Change in AS Path
  - Change in other attribute
  - Withdraw
  - Duplicate Withdraw

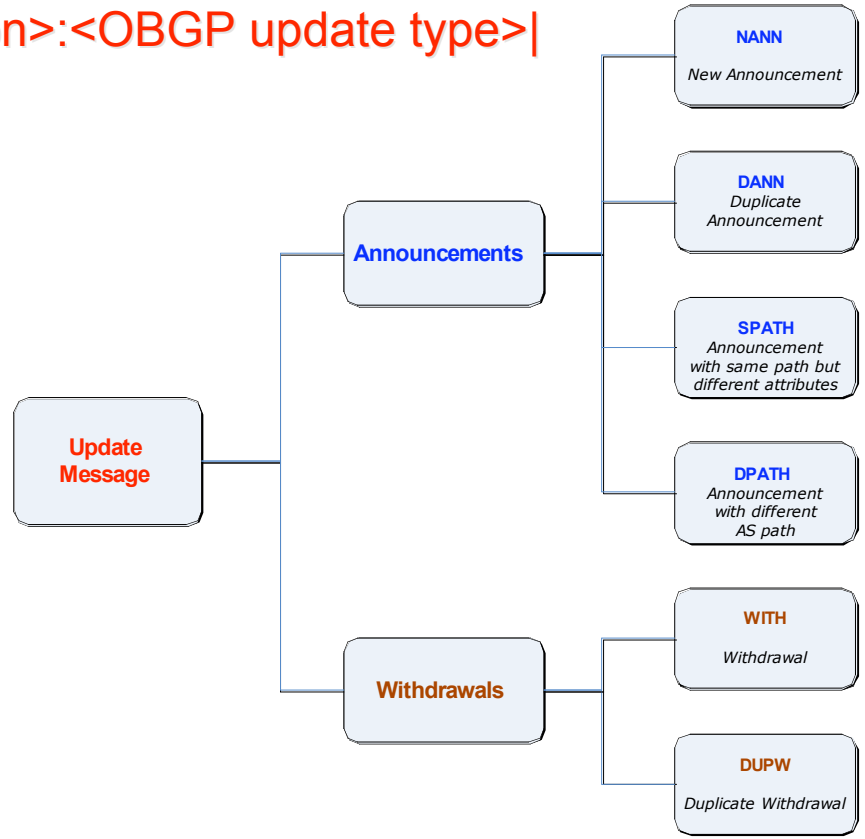

BGP4MP|1136076350|E:RIB:NANN| …
BGP4MP|1136073679|A:INC:DPATH| …

Using Labels to Filter Data

- Example: Find suballocation hijacks.
- Only need new announcements and withdraws

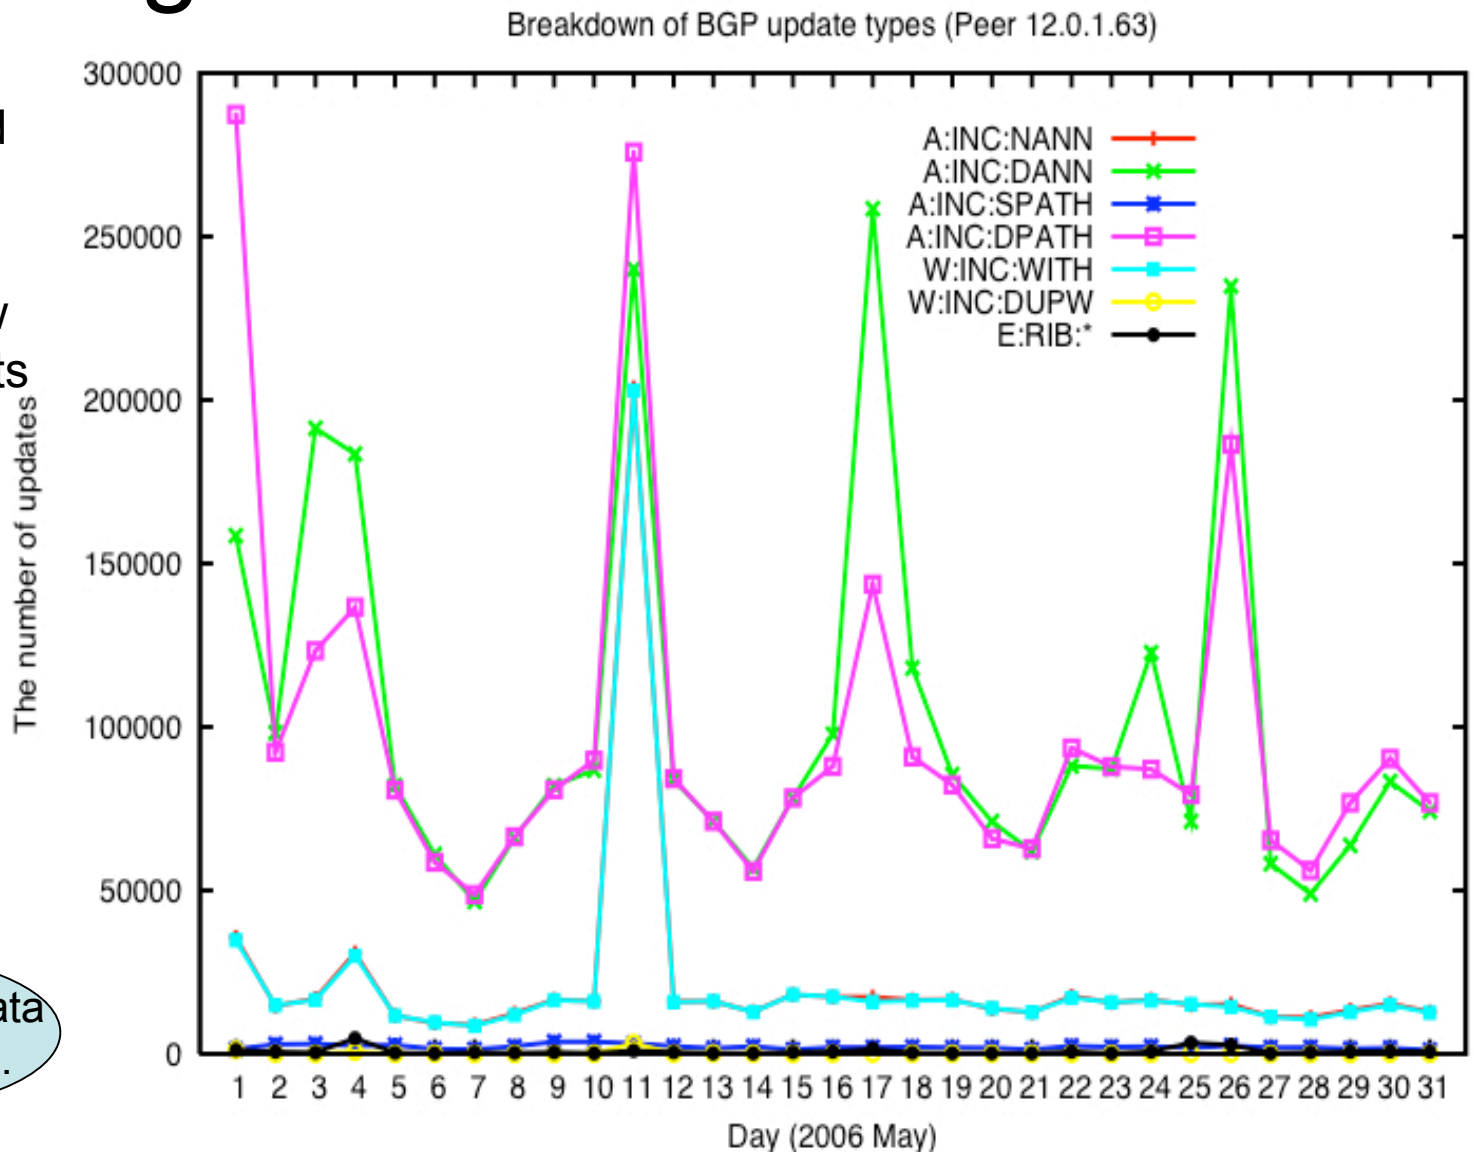

83% of update data can be ignored.

Is The Collected Data Accurate?

- May lose updates due to data collection errors
  - Start with an accurate RIB
  - Apply updates in log
  - Should match the next RIB dumped by router
    - modulo some race conditions near dump time
  - Does this really work on RouteViews/RIPE data?
- May miss dynamics when session is down
  - Must clearly label when session fails

Comparing Updates and RIBs

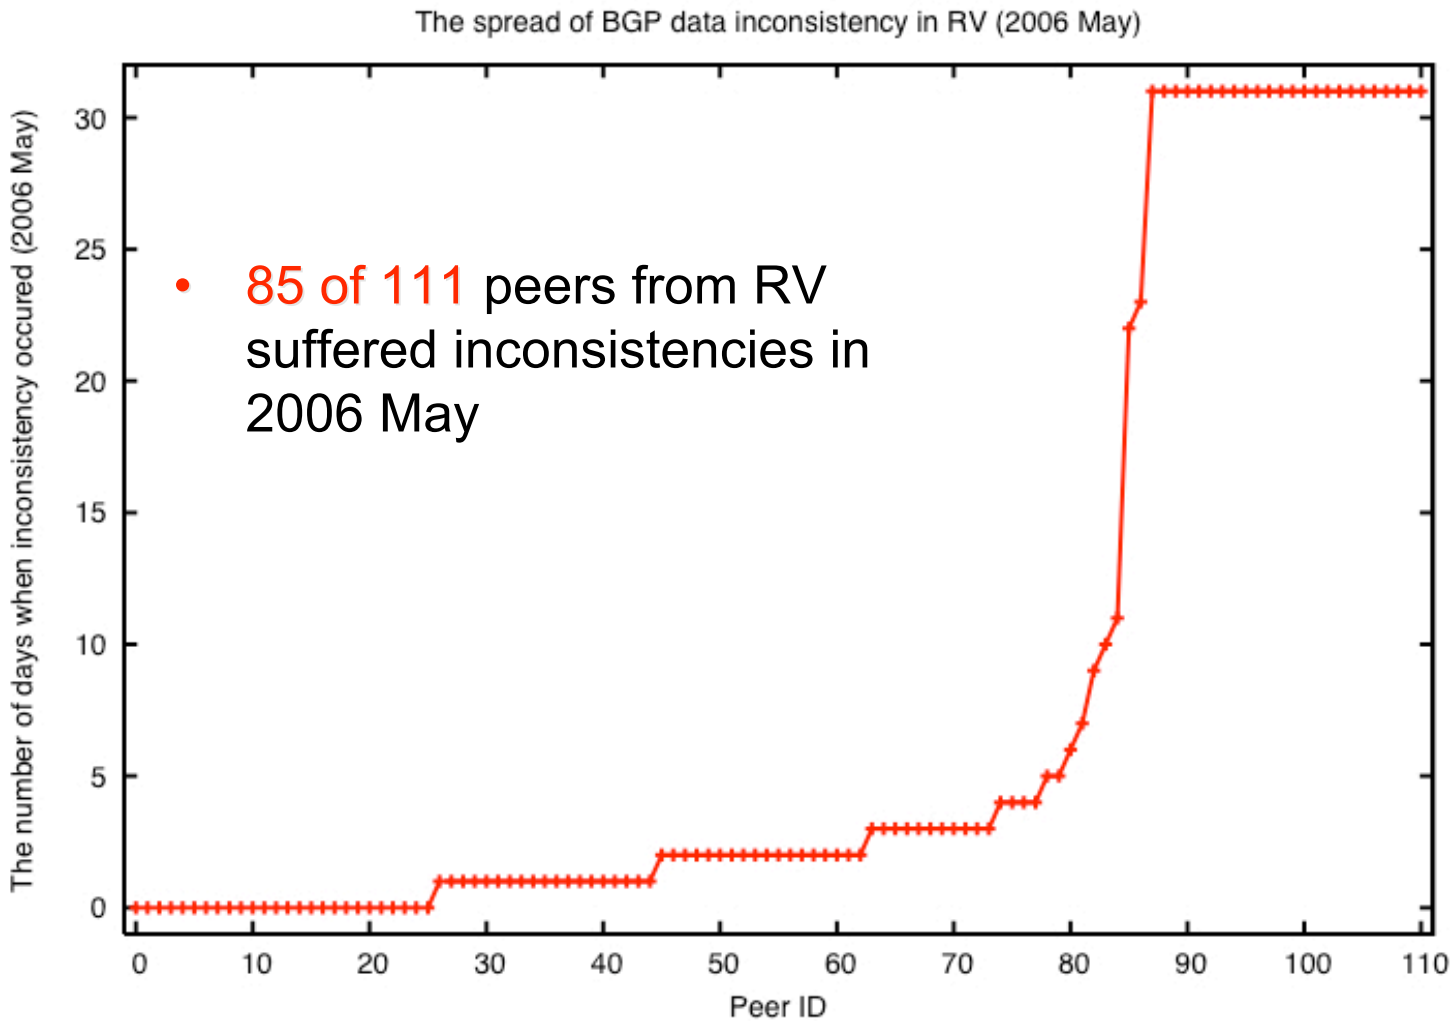

- 85 of 111 peers from RV suffered inconsistencies in 2006 May

How Many Differences Occur

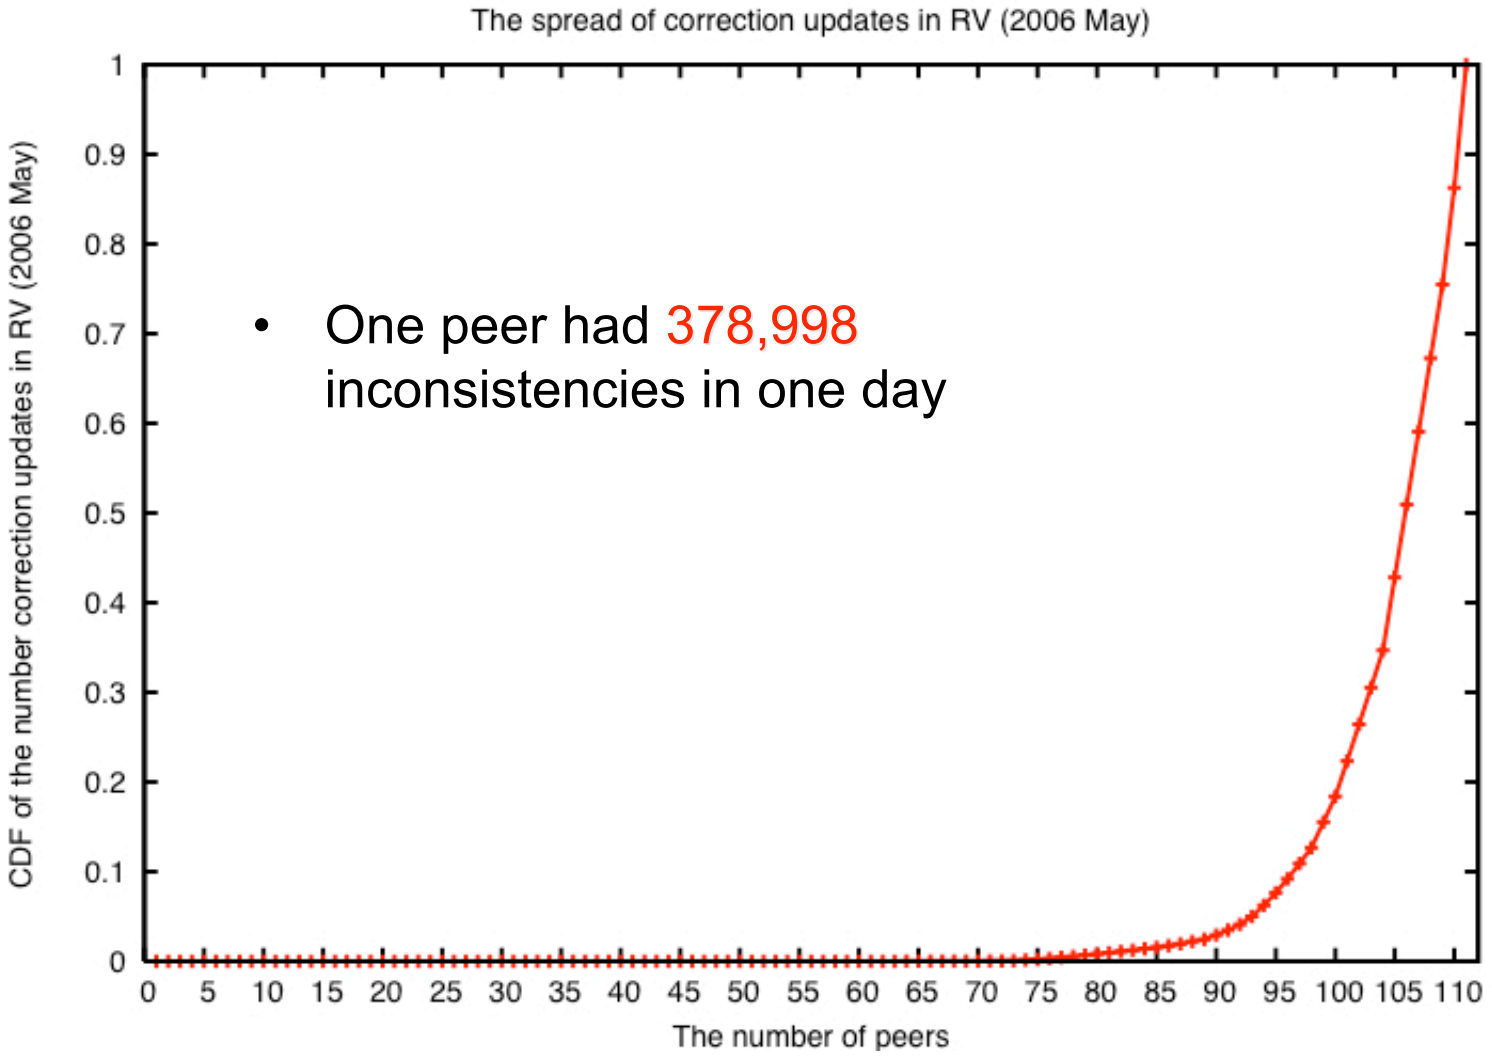

- One peer had 378,998 inconsistencies in one day

Inconsistencies and Session Failures

- Session down: RIB-IN drops to empty
- Session up: Table transfer

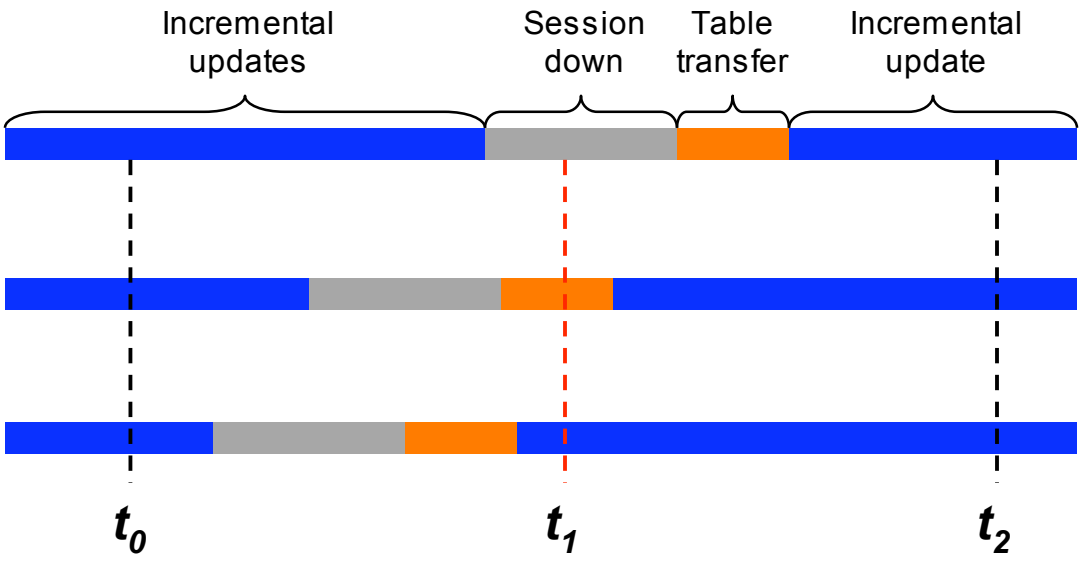

rib dumping time

How Long Does An Error Persist?

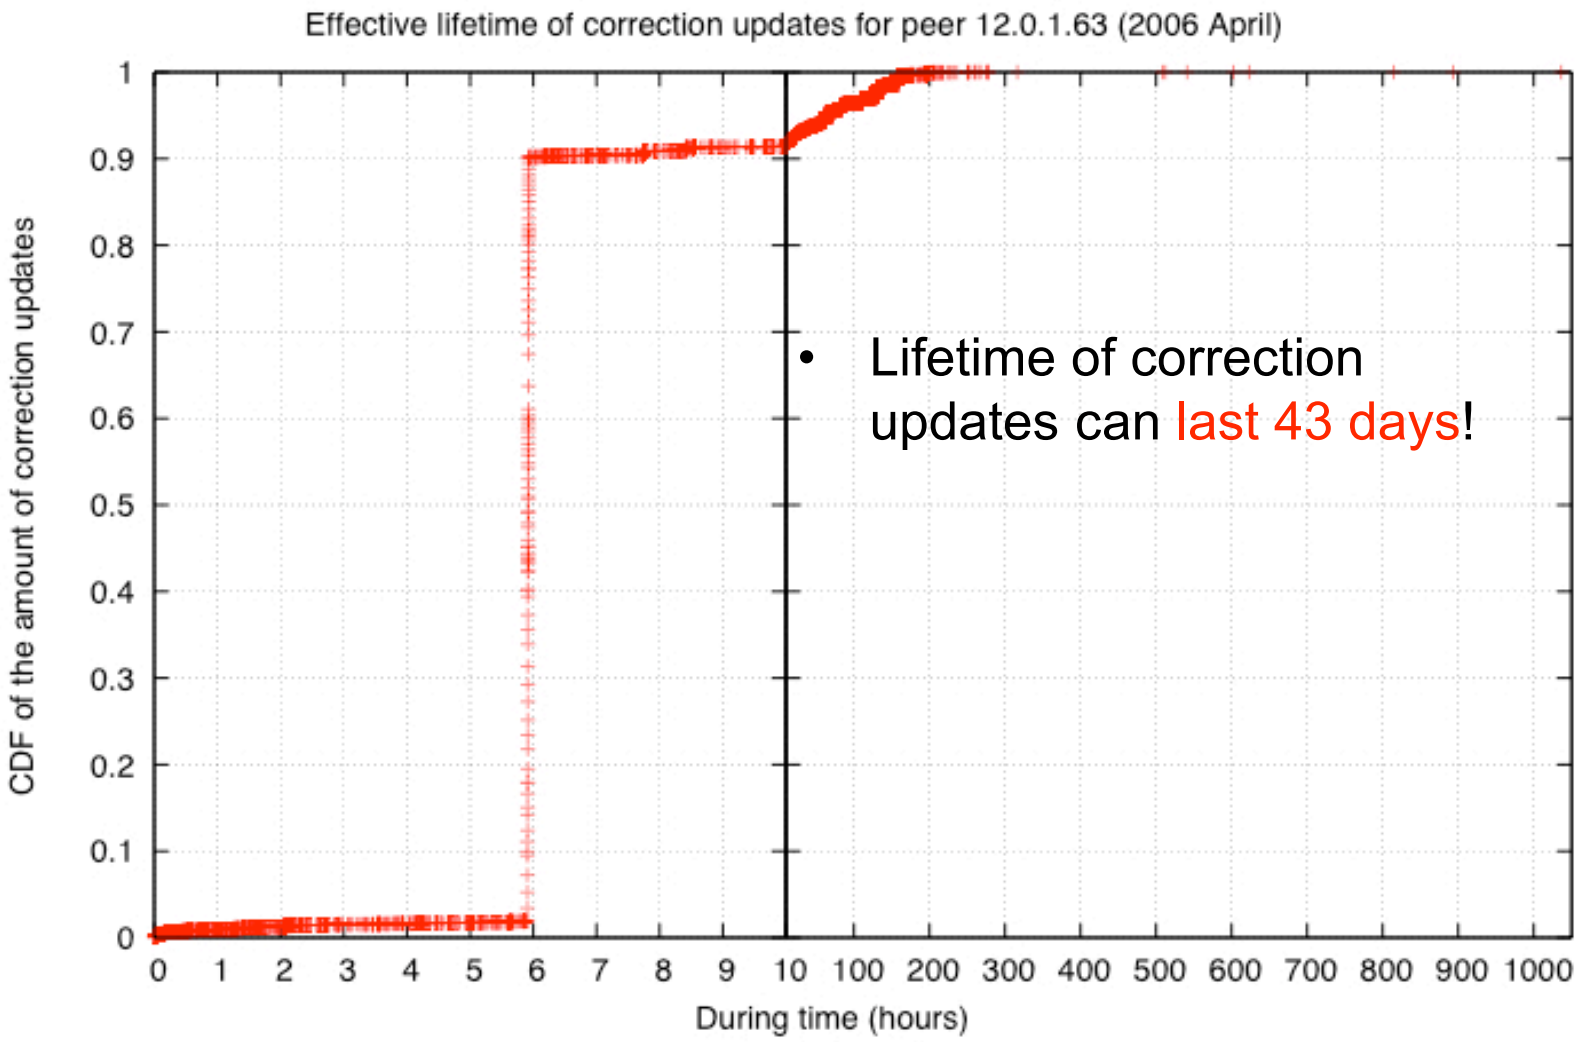

- Lifetime of correction updates can last 43 days!

Correction Updates Added by OBGP

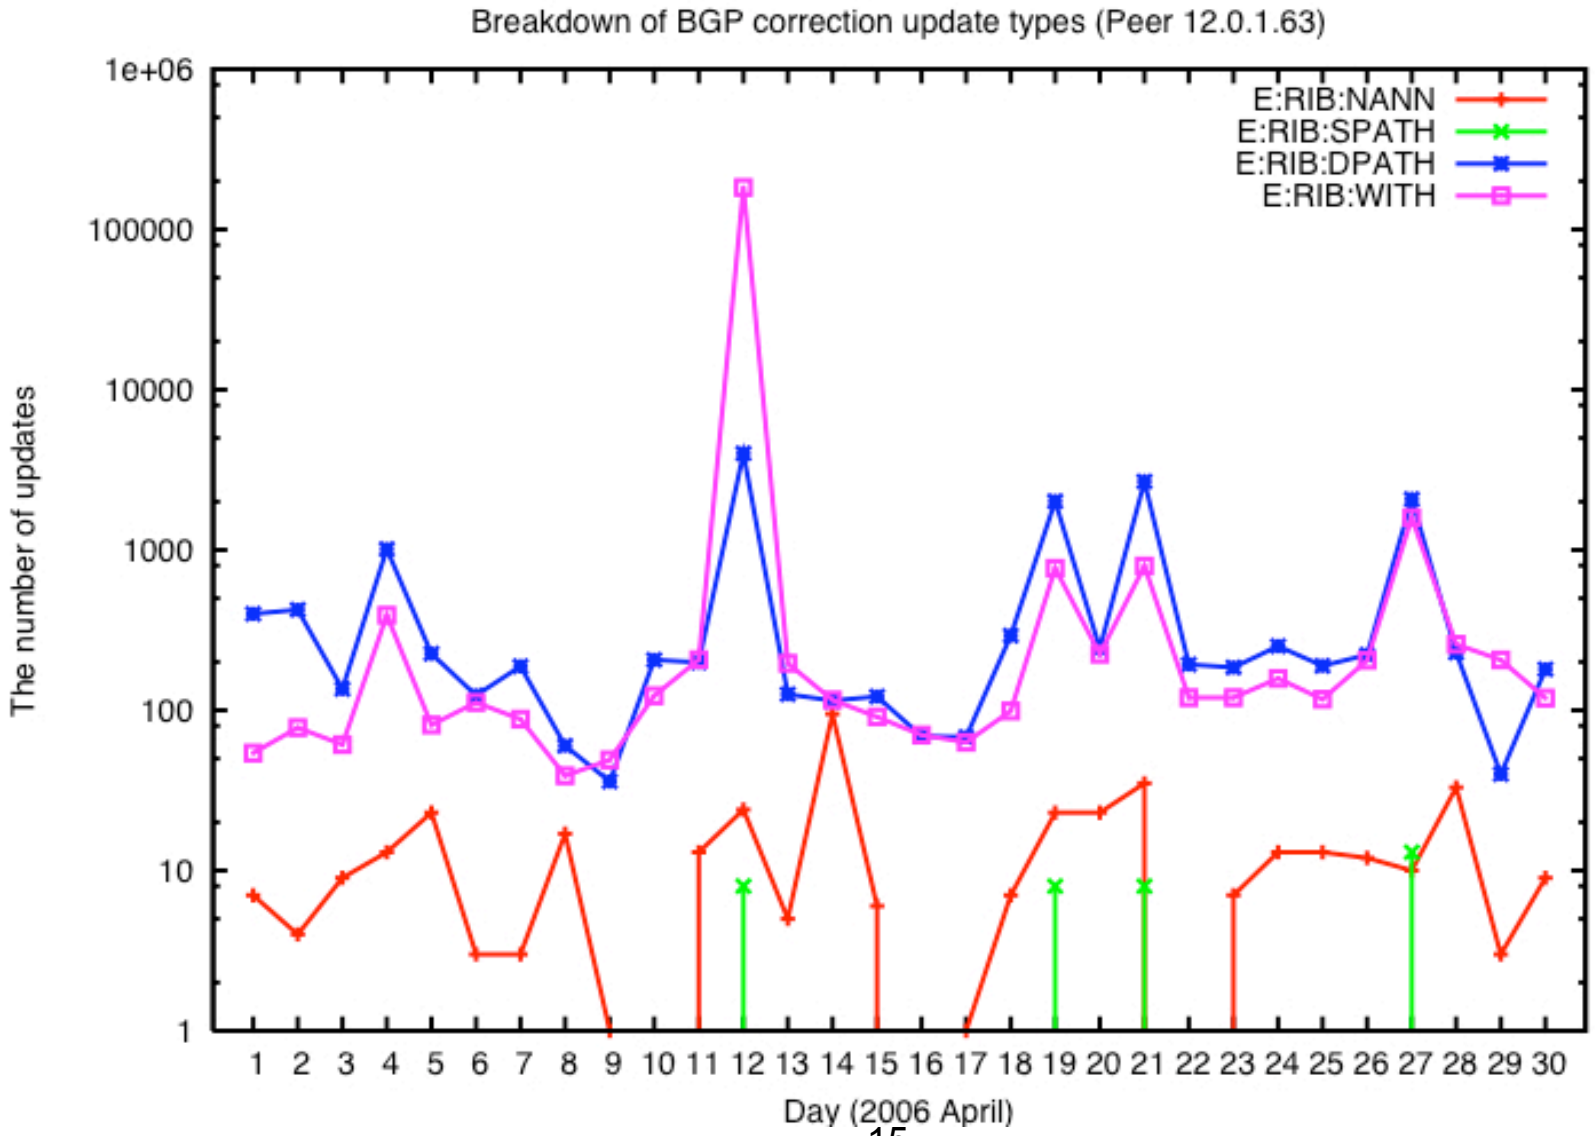

OBGP Summary

- Organizes data into a consistent format
- Adds labels to quickly find relevant data
- Adds additional state messages
  - When does a route table dump occur?
  - When does a table transfer occur?
- Identifies and corrects update error messages
  - May or may not be significant
  - Easy to evaluate use or ignore

If you are using RIPE/RouteViews/etc. data, consider OBGP as pre-processing tool

http://netsec.cs.colostate.edu/tools.html

Q&A

http://netsec.cs.colosate.edu/tools.html