

# WIT: A Watchdog System For Internet Routing

Dan Massey  
Colorado State University

*Beichuan Zhang*    *Dave Meyer*

University of Arizona

RouteViews/Univ of Oregon

*Lixia Zhang*

UCLA

## Pakistan Cuts Access to YouTube Worldwide

### RIPE NCC:

YouTube reacted **about 80 minutes after** the Pakistan Telecom announcements, and all the major events finished after **about two hours**.

### Renesys:

This story is almost as old as BGP. Old hands **will recognize this as, fundamentally, the same problem** as the infamous AS 7007 from 1997, a more recent ConEd mistake of early 2006 and even TTNNet's Christmas Eve gift 2004.

### Ars Technia:

This vulnerability has been known for a long time, and smaller-scale accidents of this nature happen at regular intervals. But so far, efforts **haven't produced any results yet**.

# Are You Being Hijacked Right Now?

- DNS root servers (and huge chunks of the Internet) were hijacked twice over the summer by the same AS!!
- And many related problems....
  - Are you unreachable from parts of the Internet?
  - Are your routes stable and reliable?
  - How did the last major cable cut impact you?

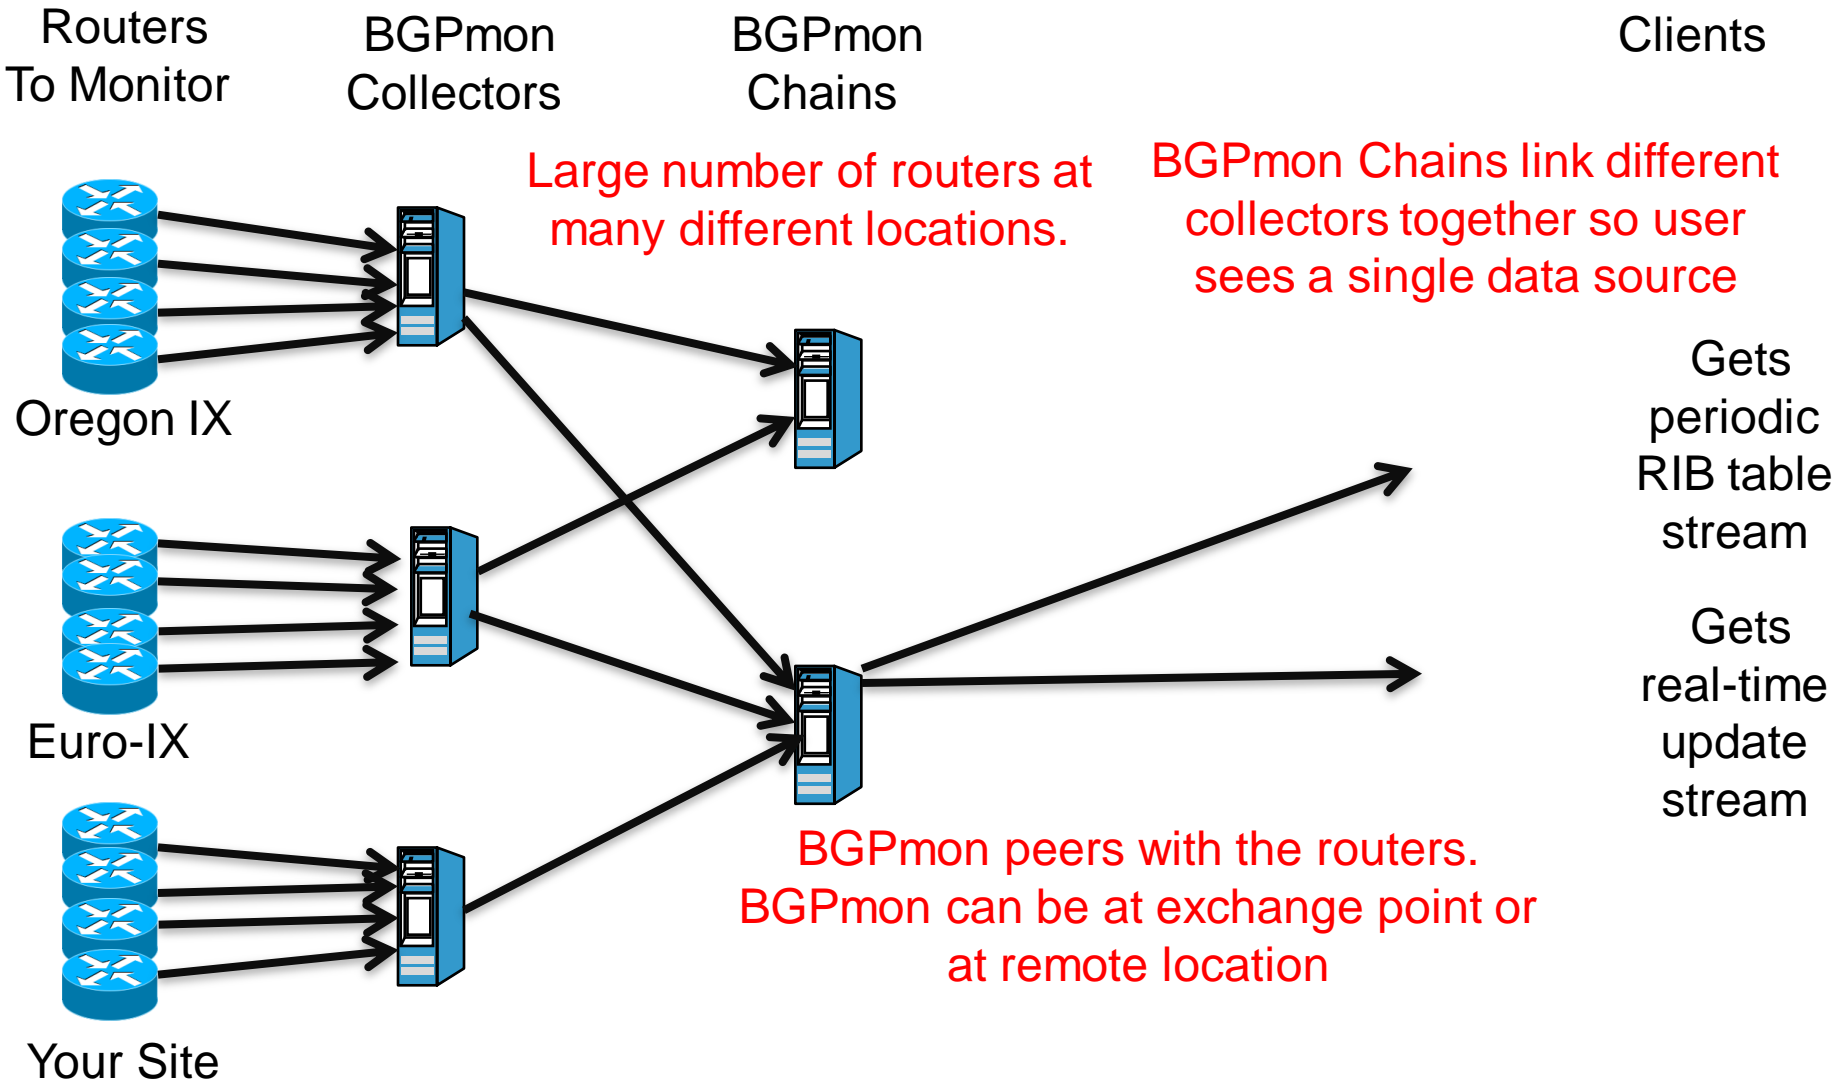
# Our Two Main Contributions

- **Routing Data Needed To Make Decisions**
  - Who is originating a route to your system?
  - Which routes changed during some major event?
  - Data from around the globe provided in real-time
- **A Prefix Hijack Alert System**
  - Digest the data and provide targeted alerts
  - Alarms indicating when you are being hijacked
  - Alert System Available Now

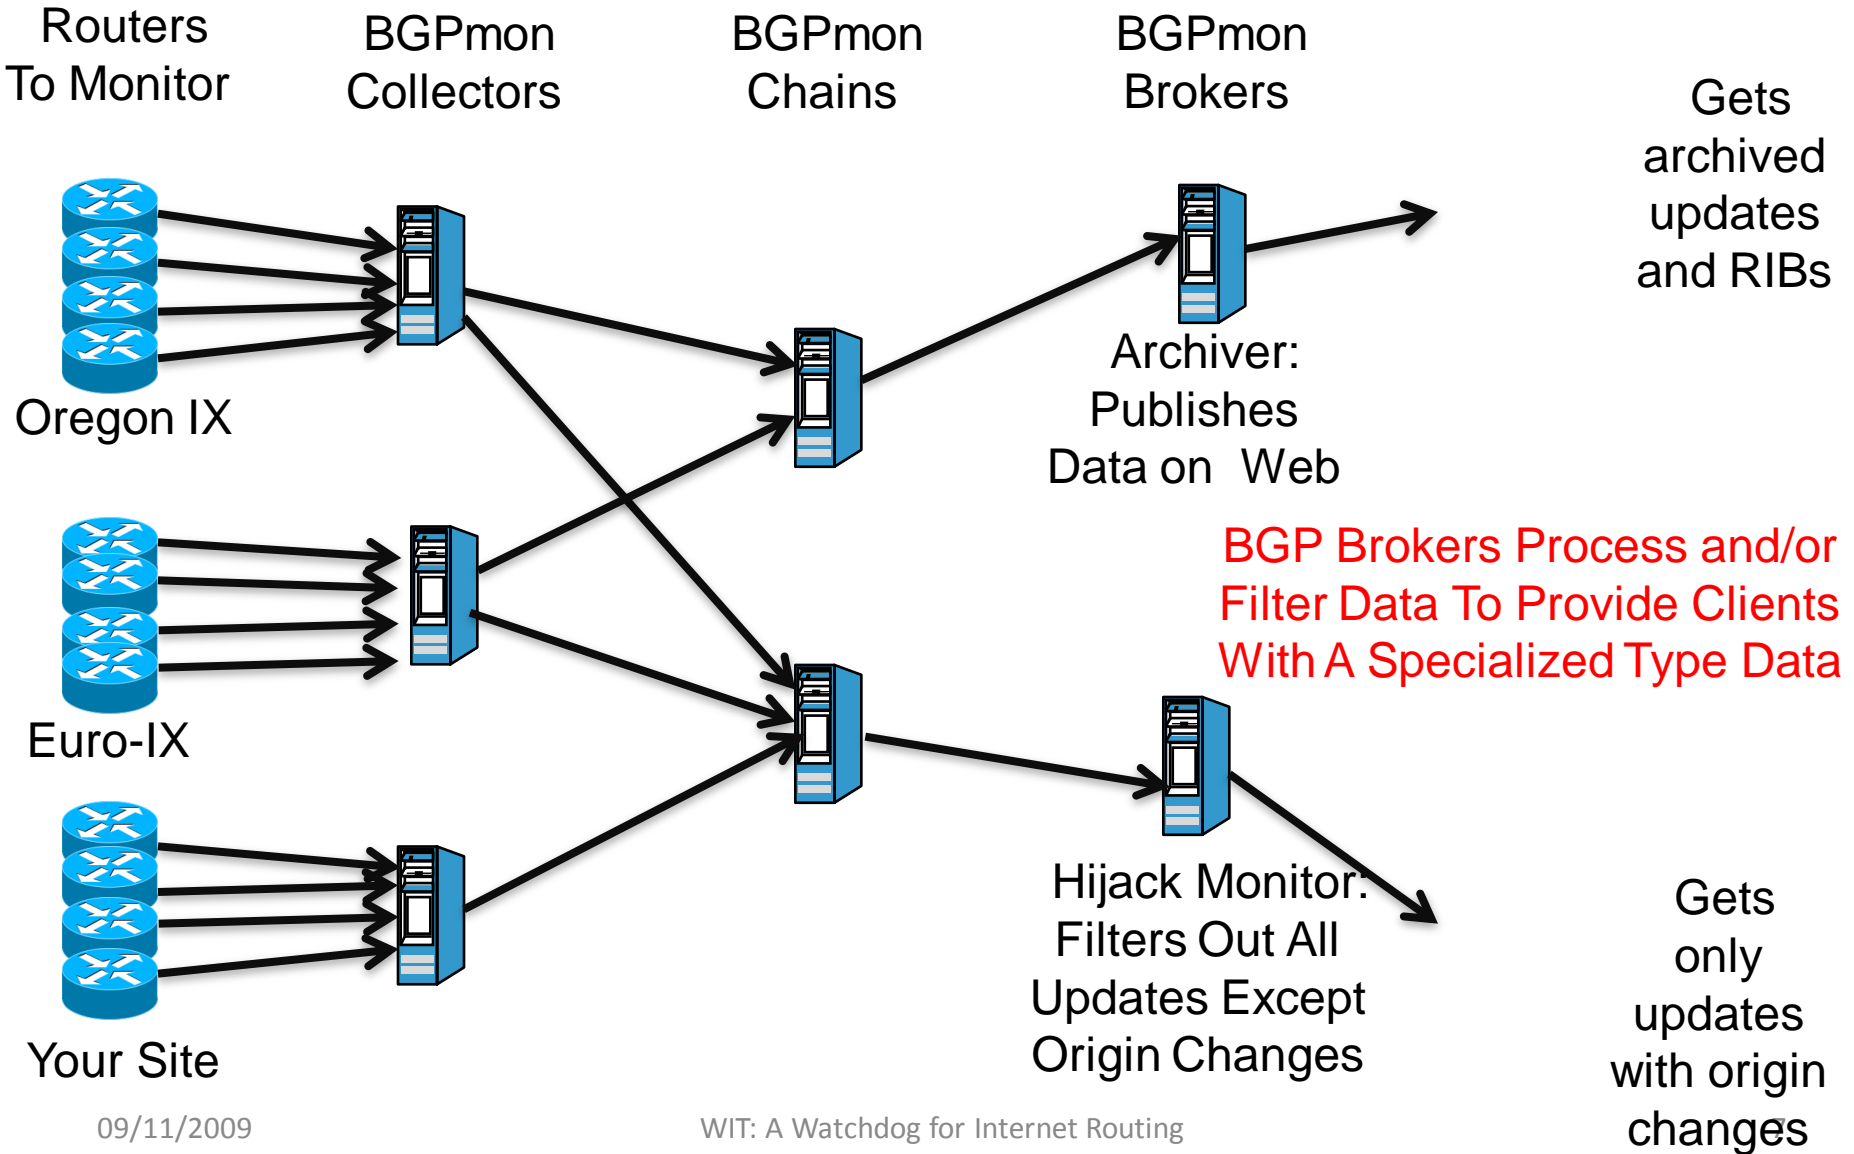
# BGPmon: Project Objectives

- Support Large Scale BGP Monitoring
  - Collect data from a large number of peers
  - Data collectors run at multiple locations/multiple exchange points
  - Users see a single coherent monitoring infrastructure
- Provide BGP Data in Real-Time
  - BGP update messages delivered to users in seconds
  - BGP RIB tables reported at regular intervals
  - Data archives also available for non-real-time users
- Create An Extensible and Easy To Use Data Format
  - Data available as bits off the wire and human readable format
  - Users can annotate data (flag customer prefixes, potential hijacks, etc.)
  - Users can easily ignore annotations they don't understand

# BGPmon Base Infrastructure



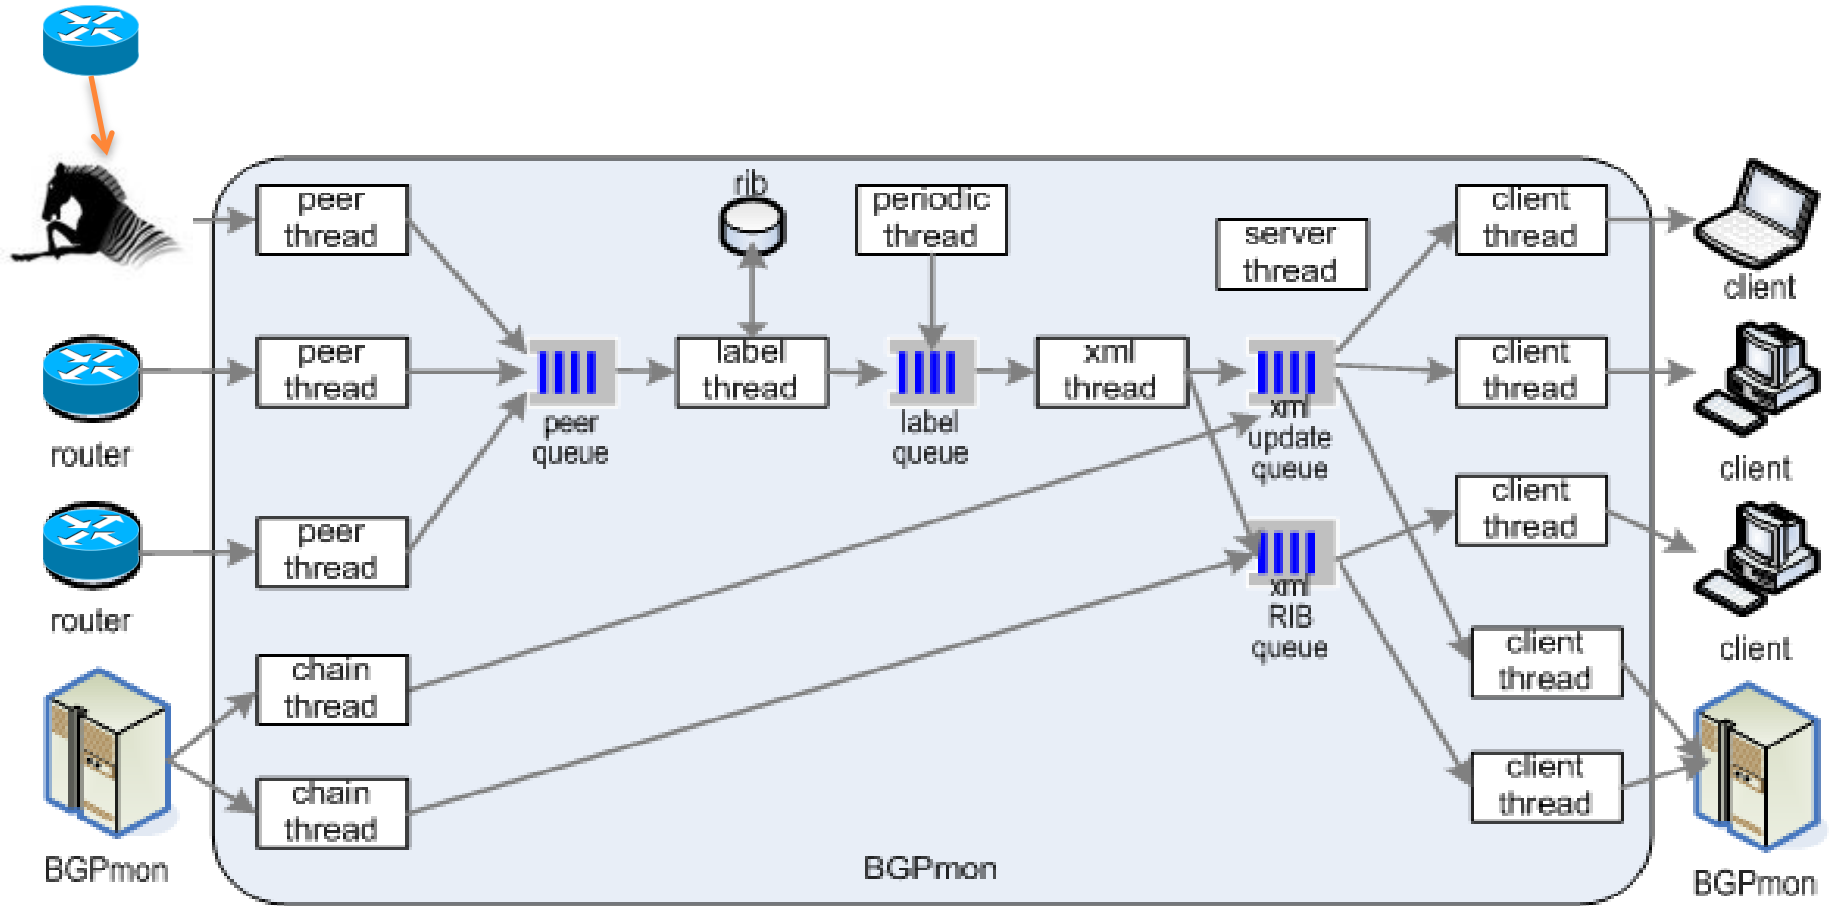
# Enhanced Monitoring With Brokers



# Large-Scale and Real-Time

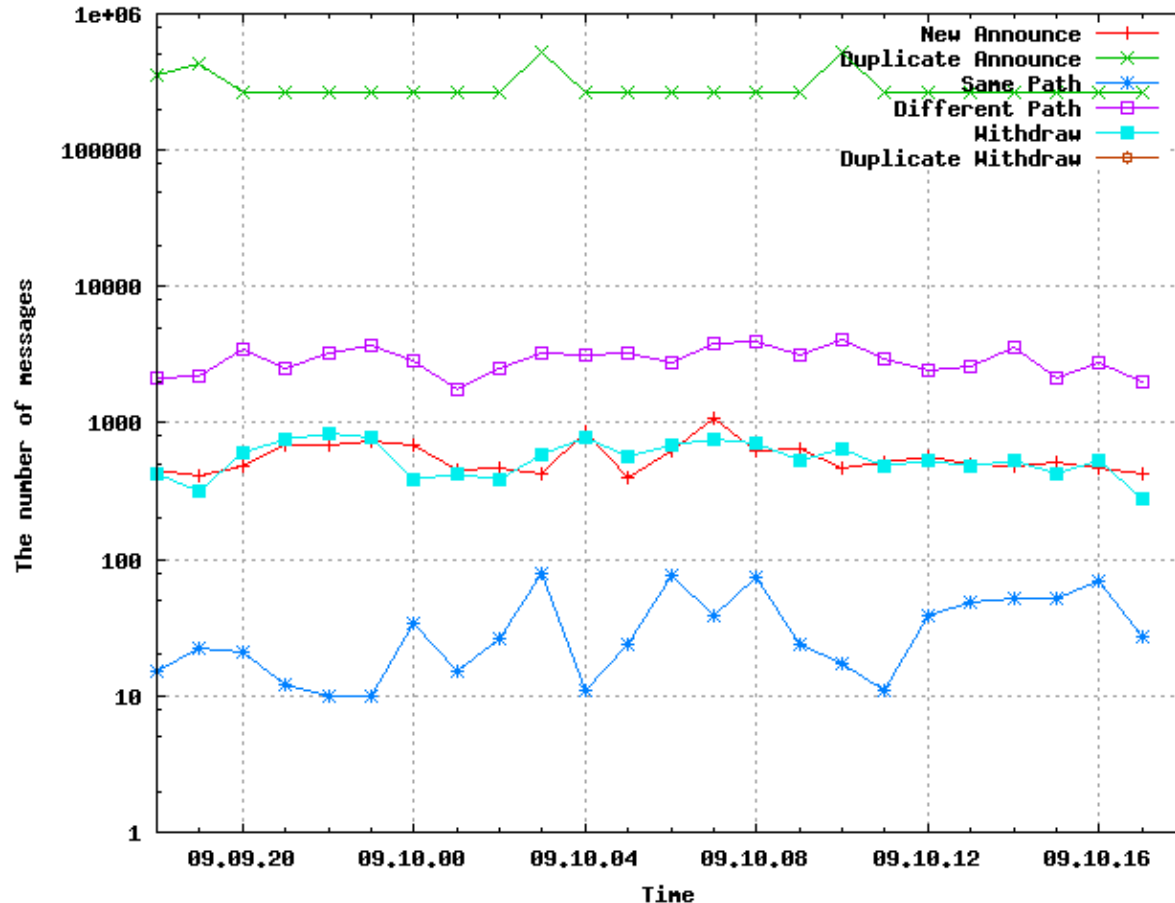
- Need Fundamental Changes In Monitoring Infrastructure
  - Provide real-time access to both route tables and incremental updates
  - Manage table transfers and update bursts from routers
  - Scale to large numbers of BGP peers and vast numbers of clients
  - Protect monitoring system from slow or misconfigured clients
- Requires Software Dedicated to Monitoring
  - BGPmon: dedicated software for monitoring and real-time delivery

# BGPmon Collectors



Internal Dynamic Flow-Control Algorithm  
"Push" Data To Clients In Real-Time  
**Version 7 In Alpha Out Now!**

# Message Count From One ISP



# Backward Compatibility

RouteViews and RIPE traditionally based on Quagga software and produce MRT output.

This is changing, but difficult to replace known infrastructure....

Small change to Quagga allows Quagga to BGPmon interaction.

Final BGPmon feature completed in August... Accepting Quagga/MRT data



# MRT to BMF Transformation

All MRT Table Dump v2 messages are converted to BGPmon internal format (BMF):

BMF header:

- Timestamp (32 b field)
- Precision Time (32 b)
- SessionID (16 b)
- Type (16 b)
- Length of BGP message (32 b)

BGP header:

- Marker (16 B)
- Length of Message (16 b)
- Type (8 b)

BGP message:

- Withdrawn Routes Length (32 b)
- Path Attribute Length (32 b)
- Path Attributes (variable)
- NLRI (variable)

# BGPmon Output Format

- Data Delivered in Real-Time Via TCP Connections
  - Format needs to capture exact bits off the wire
  - Format should be readable by humans
- Clients/Brokers should be able to annotate the data
  - Add locally useful tags to the data
  - Be able to pass the data and the annotations to other users
  - Clients who don't understand the added information should be able to easily ignore
- Achieve all of the above with an XML BGP data Format
  - Early version Internet Draft discussed at IETF GROW

# Sample Real-Time Data

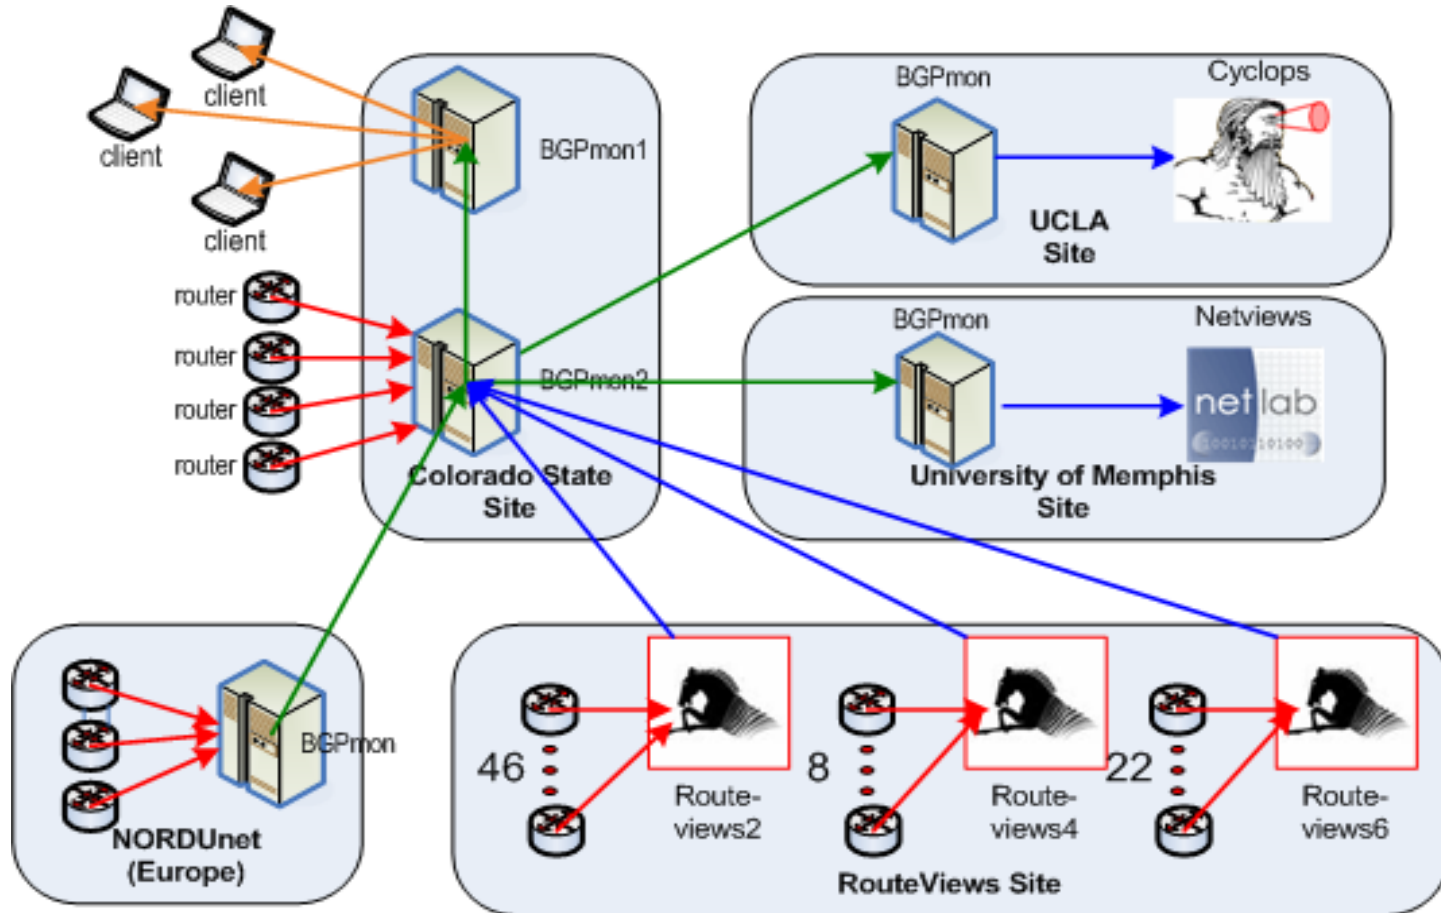
```
<BGP_MESSAGE length="00001630" version="0.2"
  xmlns="urn:ietf:params:xml:ns:xfb-0.2" type_value="4" type="TABLE">
<TIME>
<TIMESTAMP>1252041252</TIMESTAMP>
<DATETIME>2009-09-04T05:14:12Z</DATETIME>
<PRECISION_TIME>690</PRECISION_TIME>
</TIME>
<PEERING><SRC_ADDR afi="IPv4" afi_value="1" if_index="0">95.140.80.254</SRC_ADDR>
<SRC_PORT>179</SRC_PORT>
<SRC_AS>31500</SRC_AS>
<DST_ADDR afi="IPv4" afi_value="1" if_index="0">128.223.51.15</DST_ADDR>
<DST_PORT>179</DST_PORT>
<DST_AS>6447</DST_AS>
<BGPID>0.0.0.0</BGPID></PEERING>
<ASCII_MSG><MARKER length="16">FFFFFFFFFFFFFFFFFFFFFFFFFFFFFFFF</MARKER>
<LENGTH>71</LENGTH>
<TYPE value="2">UPDATE</TYPE>
<UPDATE>
<PATH_ATTRIBUTES>
.....
</PATH_ATTRIBUTES>
<NLRI count="2"><PREFIX afi="IPv4" afi_value="1" safi="UNICAST" safi_value="1">64.76.88/24</PREFIX>
<PREFIX afi="IPv4" afi_value="1" safi="UNICAST" safi_value="1">64.76.81/24</PREFIX></NLRI>
</UPDATE>
</ASCII_MSG>
</BGP_MESSAGE>
```

# Archived Data Storage Size

Format	Uncompressed (Bytes)	/MRT size	Compressed	/MRT size
<b>MRT</b>	26711666	1.00	5614650	1.00
<b>Bgpdump</b>	74551628	2.79	5645044	1.01
<b>XML</b>	264824363	9.91	13445451	2.39

**XML data requires more space to store,  
But compresses to nearly match binary format**

# BGPmon Deployment Today



# Questions