

Revisiting ASPATH Analysis

Determining 'Tier 1' through AS Path Analysis of Full Routing Tables

Vijay K. Adhikari, KCM

Gaurab Raj Upadhaya, PCH

@

APNIC 20, Hanoi






“Tier 1 is not a binary property”

Quoting Bill Woodcock




Flame-Shield

- ‘THE’ single Internet routing table does not exist. Our work is based on multiple views from multiple sources.
- There is no such thing as ‘Tier 1’ across the entire Internet.
- Content providers are not proportionately represented in this analysis.



Hypothesis

- More than two 'Tier 1' ASN should not appear in any ASPATH
- Our algorithm tests the above hypothesis on available routing data (oix-route-views and pfs-japan) against a list of possible 'Tier 1' ASN



Building a starting point list


- Routing table archives from Route-Views (oix-route-views) and Tokyo transit router data provided by Philip Smith (pfs-japan).
 - We found the 20 most frequently-occurring ASNs in all observed AS paths after removing the immediate next-hop transit providers
 - Data for the 03rd of each month in 2005 was used from oix-route-views, and pfs-japan.
 - We then found the intersection of the two lists
- The final list of top 10 gives a 'Tier 1' list against which the algorithm was tested



Our 'Tier 1' list

Created intersection of multiple data sets

- 701 - UUNET / MCI
- 3356 - Level 3
- 7018 - AT&T
- 209 - Qwest
- 3549 - Global Crossing
- 1239 - Sprint
- 2914 - NTT/ Verio
- 6461 - MFN
- 2828 - XO Communications
- 4637 - Reach




Looking at ASPATH

- A hypothesis we tested
 - There cannot be more than two 'Tier 1' ASN in a single ASPATH.
 - From pfs-japan data, AS 174 doesn't seem to be an Tier 1 in the AP region.

2516	1239	2914	174	812	i	
2497	2914	174	i			
2516	1239	2914	174	812	800	i

Where 1239 = Sprint, 2914 = NTT / Verio, 174 = Cogent



Point of Interest

- AOL does not appear on our ‘Tier 1’ list, because they do not provide transit to down-stream networks.
- When we add AS1668 (AOL) to our “Tier 1” list and run the algorithm, we get expected result.
 - Whenever AS1668 was on the ASPATH, at most one other ASN from our top 10 list were visible on the ASPATH.



AS-PATH peculiarities


- Further tests revealed certain interesting anomalies.
 - Customer ASN between two ASN on the 'Tier 1' list
 - Routing leaks ?
- Some examples ..




Example - Anomalies

- **Network: 200.124.192.0**
- **Data Source: 06-03-0300 pfs-japan**

```
• *> 200.124.192.0      202.249.2.169          0 2497 701 14551 11664  
3549 1239 6057 14234 i
```



- This is also the best selected path, where Telemax - an operator in Argentina is seen between two global networks.



Example - Anomalies

- **Data Source : 08-03-0800 oix-route-views**
- Network: 204.13.189.0
- ASPATH - 11608 **2914** 27506 10910 **701 1239** 6395 16852 35967 i

- Network: 204.13.189.0
- ASPATH - 5459 **2914** 27506 10910 **701 1239** 6395 16852 35967 i

- Network: 204.13.189.0
- ASPATH - 293 **2914** 27506 10910 **701 1239** 6395 16852 35967 i

- In this case, the same network is seen traversing customer path between NTT/Verio and UUNET
- 2914 is NTT/Verio and 701 is UUNet and 1239 is Sprint



Recent Anomalies



- **oix-2005-09-03-800**
- Network: 65.172.220.0
- * 207.246.129.6 0 11608 **6461** 16524 **3356**
1239 32500 i
- Network: 155.212.207.0
- * 207.246.129.6 0 11608 **6461** 16524 **701**
1239 14751 26684 i
- Network: 205.162.1.0
- * 207.246.129.6 0 11608 **2914** 16467
16467 16467 **3356 2828** 14815 i
- Network: 205.162.1.0
- * 213.140.32.146 0 12956 **2914** 16467
16467 16467 **3356 2828** 14815 i
- Network: 217.67.128.0
- * 207.246.129.14 0 11608 **6461** 5089 **3356**
3549 34245 15952 i





Where 6461 is MFN, 3356 is Level 3, 1239 is Sprint, 2914 is NTT/Verio, 701 is UUNet and 2828 is XO

Numbers of anomalies


- In the pfs-japan dataset, only one anomaly in June, referenced earlier
- In oix dataset, the numbers are:

Month	Against Top 10 List	Without AS 4637 (Reach)
January	114	2
February	105	2
March	128	2
April	155	2
May	129	37
June	125	2
July	85	5
August	141	6
September	83	5



Treating exceptions

- Removing AS4637(REACH) from the “tier 1” list yields significantly fewer anomalies.
- This tends to confirm the hypothesis that Reach may not be very ‘tier 1’ on the global scale
- But if only looking at the pfs-japan data, Reach doesn’t yield any exceptions making it more ‘tier 1’ in the Asia Pacific Region.



What we think ?

- 'Tier 1' is a socio-economic concept, but this technique can be used to find inter-AS relationships and determine the value of particular ASNs.
- We can look deeply at available data and see if we can determine regional 'tier 1', or ASNs that are regionally on the top.




What next ?

- We will need to run more tests over more data from other regions also.
- We'd like to run the full routes intersection with peering routes at each IX we are present.
- We also plan to review earlier research done on these topics and presented at SIGCOMM and at the recent GROW WG meeting at IETF.



Acknowledgement

- Bill Woodcock, PCH
 - Philip Smith, Cisco Systems
 - Tom Vest, PCH
 - Route-views Folks
-
- Few participants at the SANOG 6,
where this presentation was first made



Thanks

- You can reach us at gaurab@lahai.com

