

Quaking Tables:

The Taiwan Earthquakes and the Internet Routing Table

APRICOT Bali 2007

Alin Popescu, Renesys Corp
Todd Underwood, Renesys Corp
Earl Zmijewski, Renesys Corp

Overview

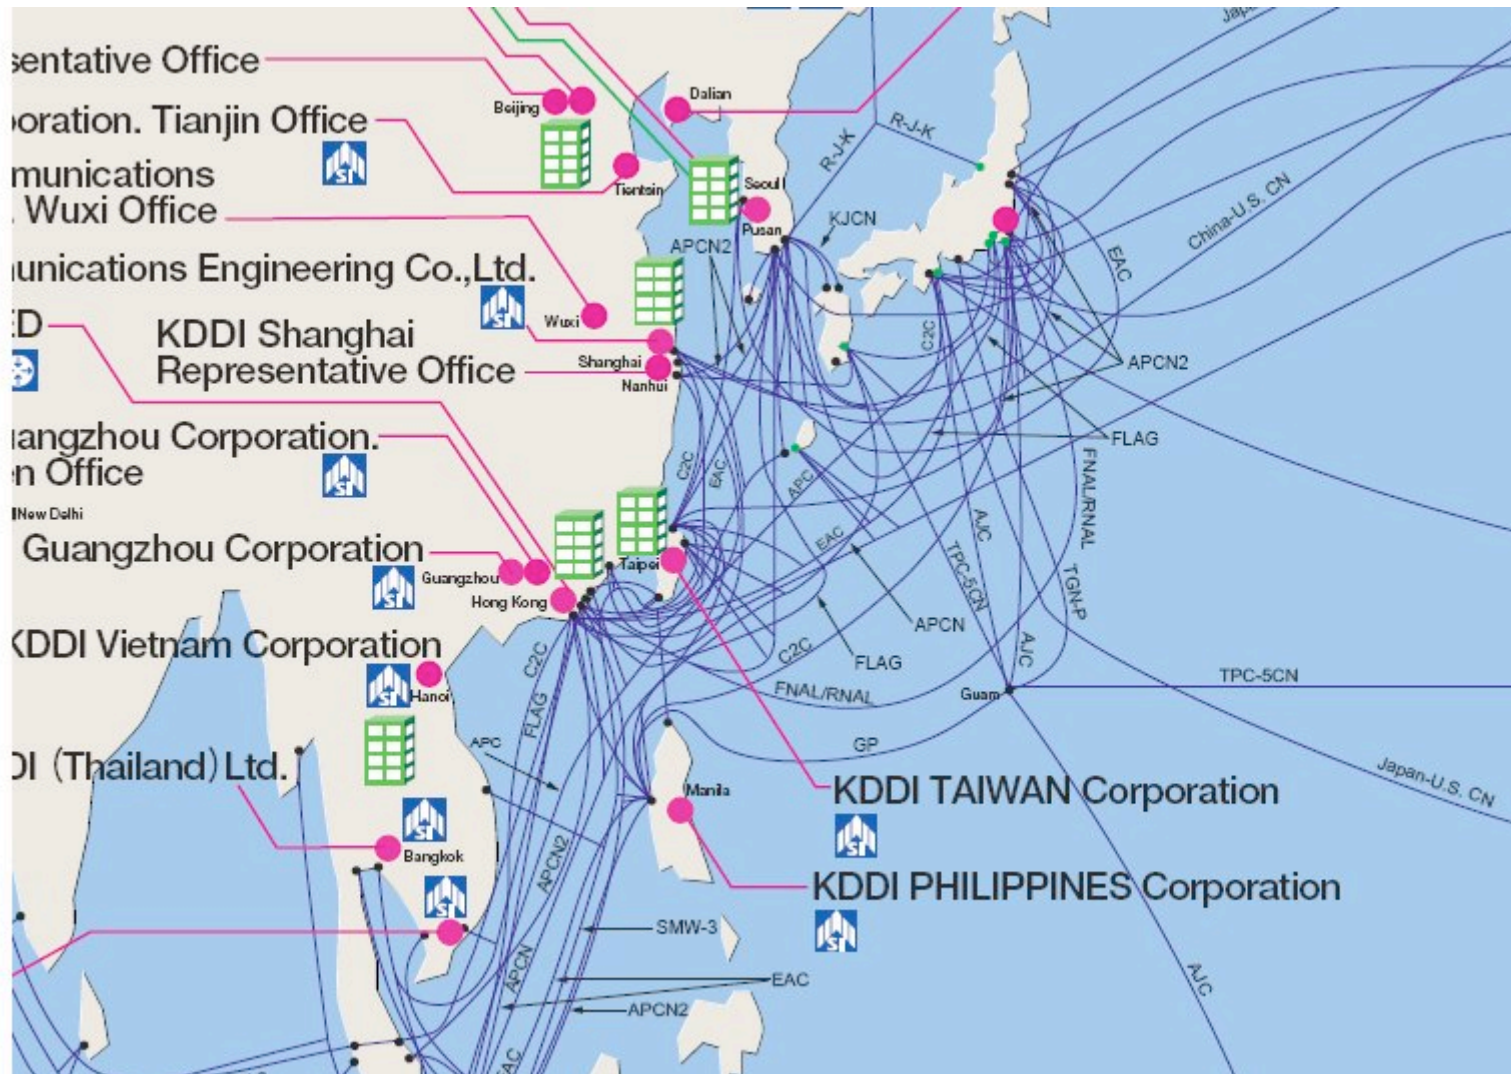
- Large earthquakes hit Luzon Strait, south of Taiwan on 26 December 2006
- Seven of nine cables passing through the straight were severed
- We review the event from a perspective of the Internet Routing tables
 - Routing outages occurred, significant congestion was reported, instability persisted
 - Recovery was delayed and uneven

Submarine cables in East Asia



- Two of nine cables **not** impacted:
 - Asia Netcom's EAC
 - Guam-Philippines
- All cables reported repaired as of February 14, 2007 (source: Office of the Telecommunications Authority of Hong Kong)

Submarine cables in East Asia (2)



Repairing submarine cables is difficult!




Image credit: Alcatel

Definitions: Outage, unreachable, unstable

- A **network outage** occurs when routes to the network are withdrawn by a large number of BGP routers worldwide.
- In this case if no less specific route is available, the network is **unreachable** and effectively disconnected from all or parts of the Internet.
- **Unstable networks** are not completely disconnected, but show frequent changes in network routing paths or alternating announcements and withdrawals (**route flapping**) – serious packet losses.

Timeline

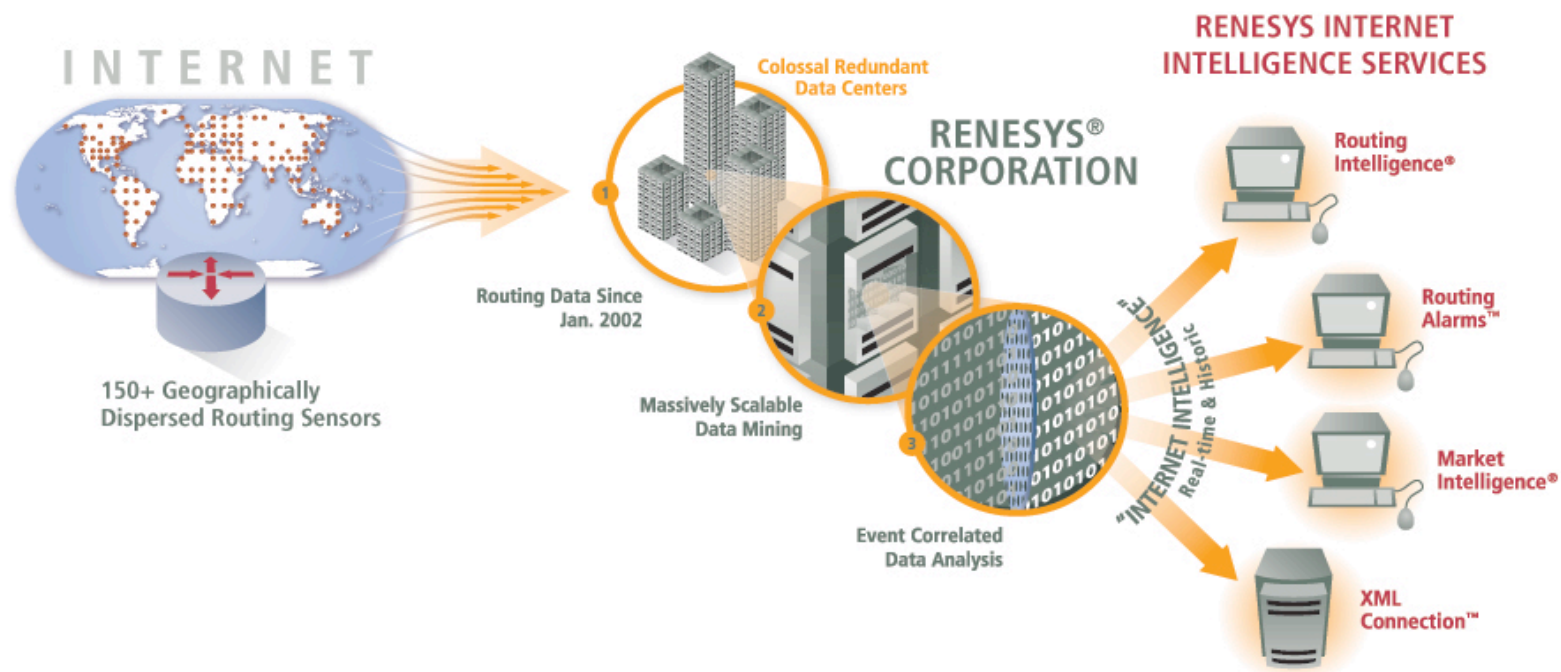
- Six earthquakes of magnitude 5.0 or higher hit the Taiwan region (all times UTC):
 - **12/26 12:26:21** **7.1** -- main quake
 - 12/26 12:34:14 6.9
 - 12/26 12:40:22 5.5
 - 12/26 15:41:44 5.4
 - 12/26 17:35:10 5.4
 - 12/27 02:30:39 5.6
 - 12/28 16:51:16 4.4
- Outaged prefixes ramp up from 400 to almost 1200 from the first quake through seventh

Timeline (2)

- 03:31 27 Dec 2006: 60 mins after the last quake, outaged network count spikes to 4k
- The “spike” is short-lived (< 2 hrs) but > 2k prefixes out for 6 hours.
- 31 Dec 2006 12:00: Outages return to pre-quake levels.
- Instability level remains high into January.

Data Collection Infrastructure


- 165+ peering sessions from 105+ different ASNs
- In this talk, we focus on East Asian prefixes only




•Disasters Have Signatures

- Sharp onset associated with some real-world event
- Slow return to baseline
 - Varies considerably
 - Power outages: fast
 - Major natural disasters, **much** slower
- Noise in the recovery (not in the onset)

•Power (Northeast US, 2003)




•Hurricane (Katrina, 2005)




The Pattern of the Taiwan Quakes

- Ramping up outages and spikes in instabilities
- Gradual increase in number of outages after major quake in Dec. 26
- Big spike in outages/unstables associated with smaller quake on Dec. 27
- Recovery typically noisy
- Pattern was probably affected by the number of different cable systems involved – this is not really one event but at least seven.


•Outages & Quakes – 3 Day



•Outages & Quakes – 10 Day



•Outages by Country – 3 Day



Why India?

Major subcontinent bandwidth heads East

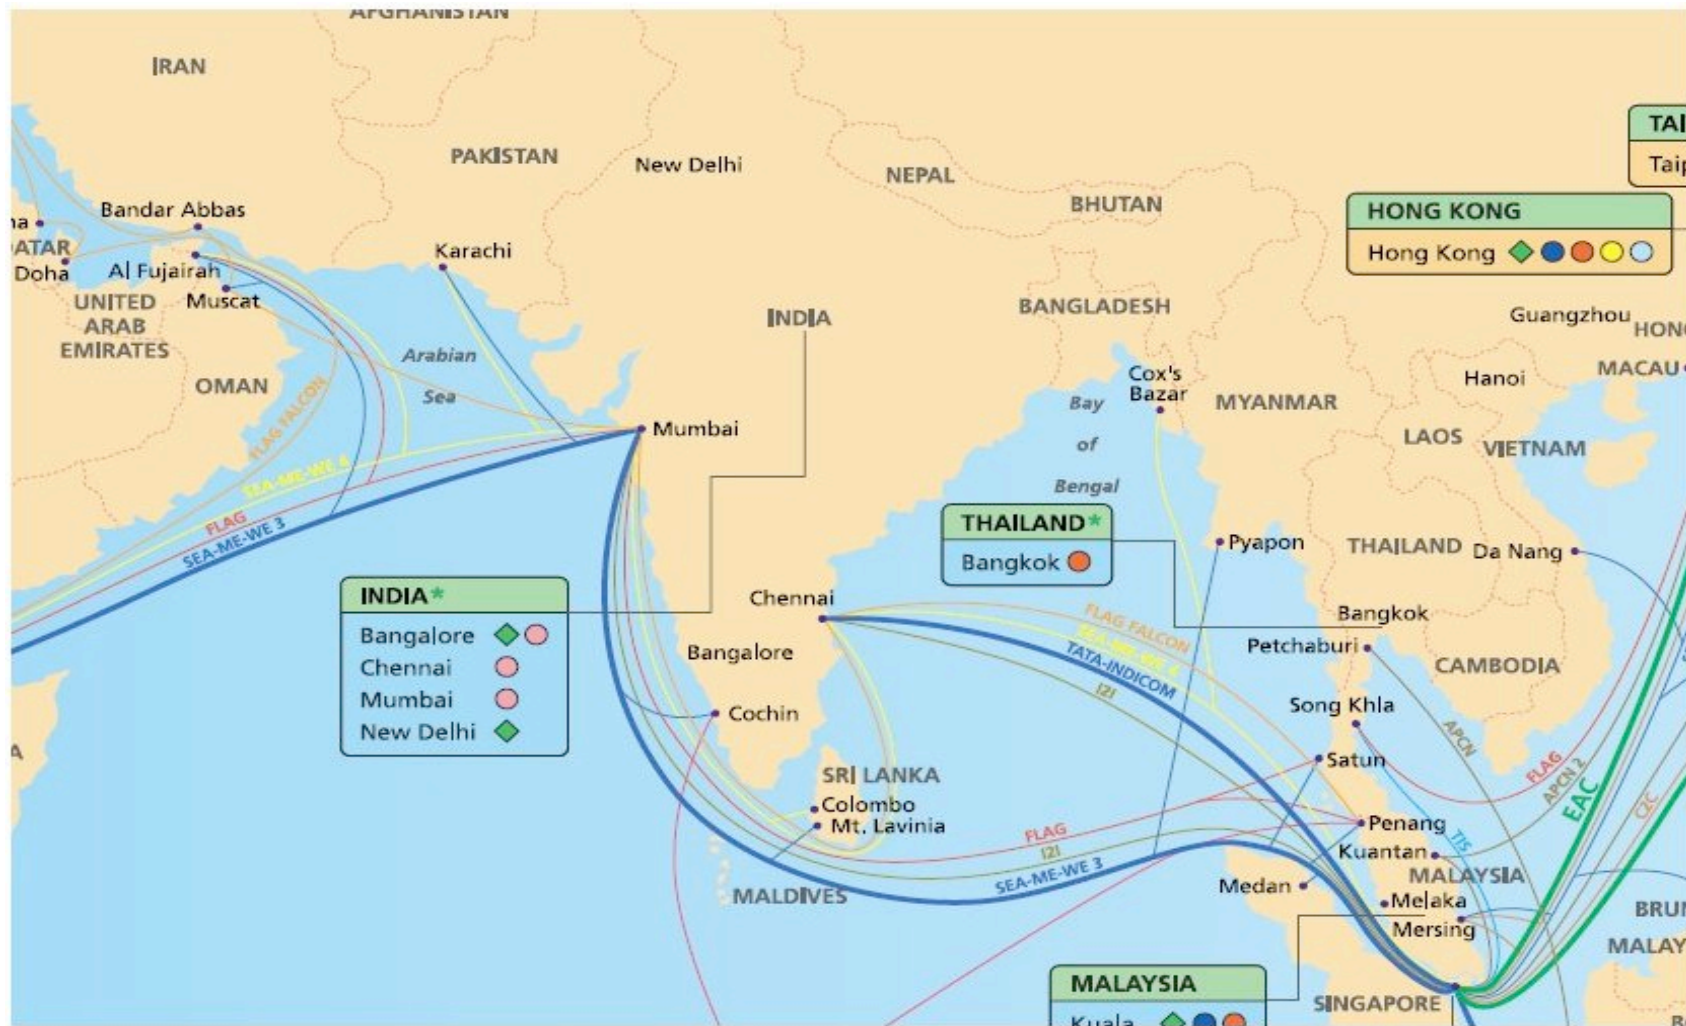




Image credit: Asia Netcom


•Outages by Country – 10 Day




•Outages by Country – Peak



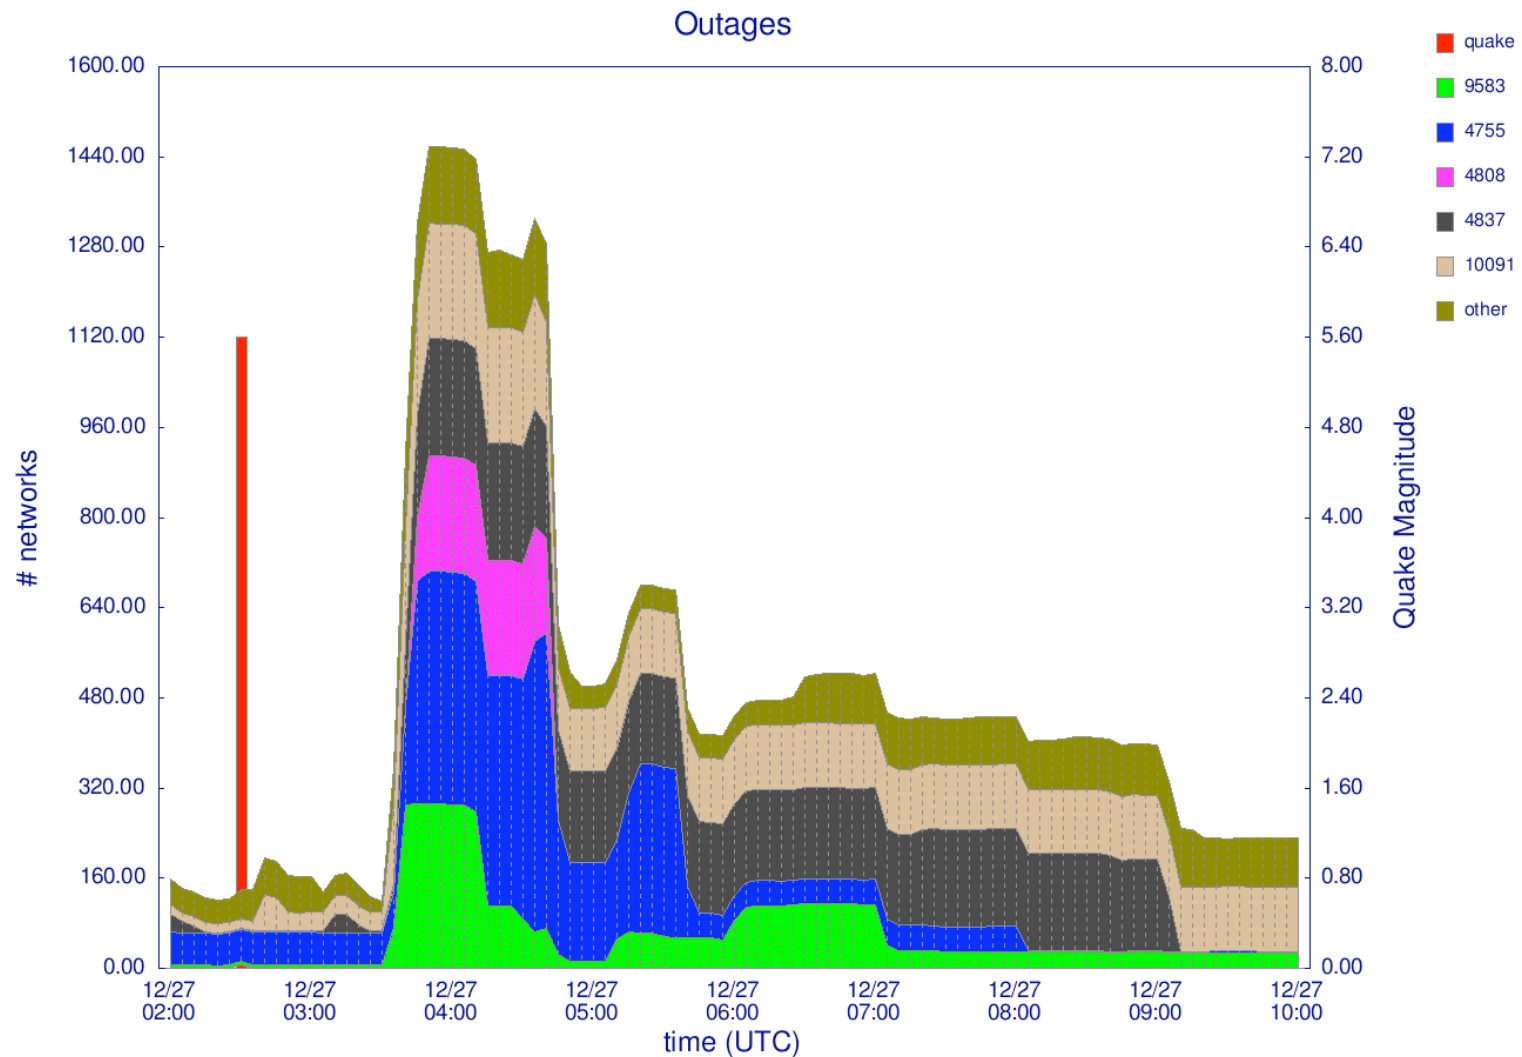
Outages by Origin ASN – 3 day



Outage by Origin ASN – 10 day




Outages by Origin ASN – Peak




•Instability Metrics

- Use algorithm described in <http://www.nanog.org/mtg-0402/ogielski.html>
- PenaltyBox(T,K,H,C): the number of globally routable prefixes at time T that have flap penalty K, using the classic flap dampening algorithm with half-life of H and ceiling of C.
- Reasonable: H=600s, C=15, K={0,...,C}.
- (Basically flap-dampening algorithm for scoring the “penalty” of a prefix).
- **Shown in graphs are pfxs w/ inst ≥ 3**


•Unstables & Quakes – 3 Day




•Unstables & Quakes – 10 Day




•Unstables by Country – 3 Day



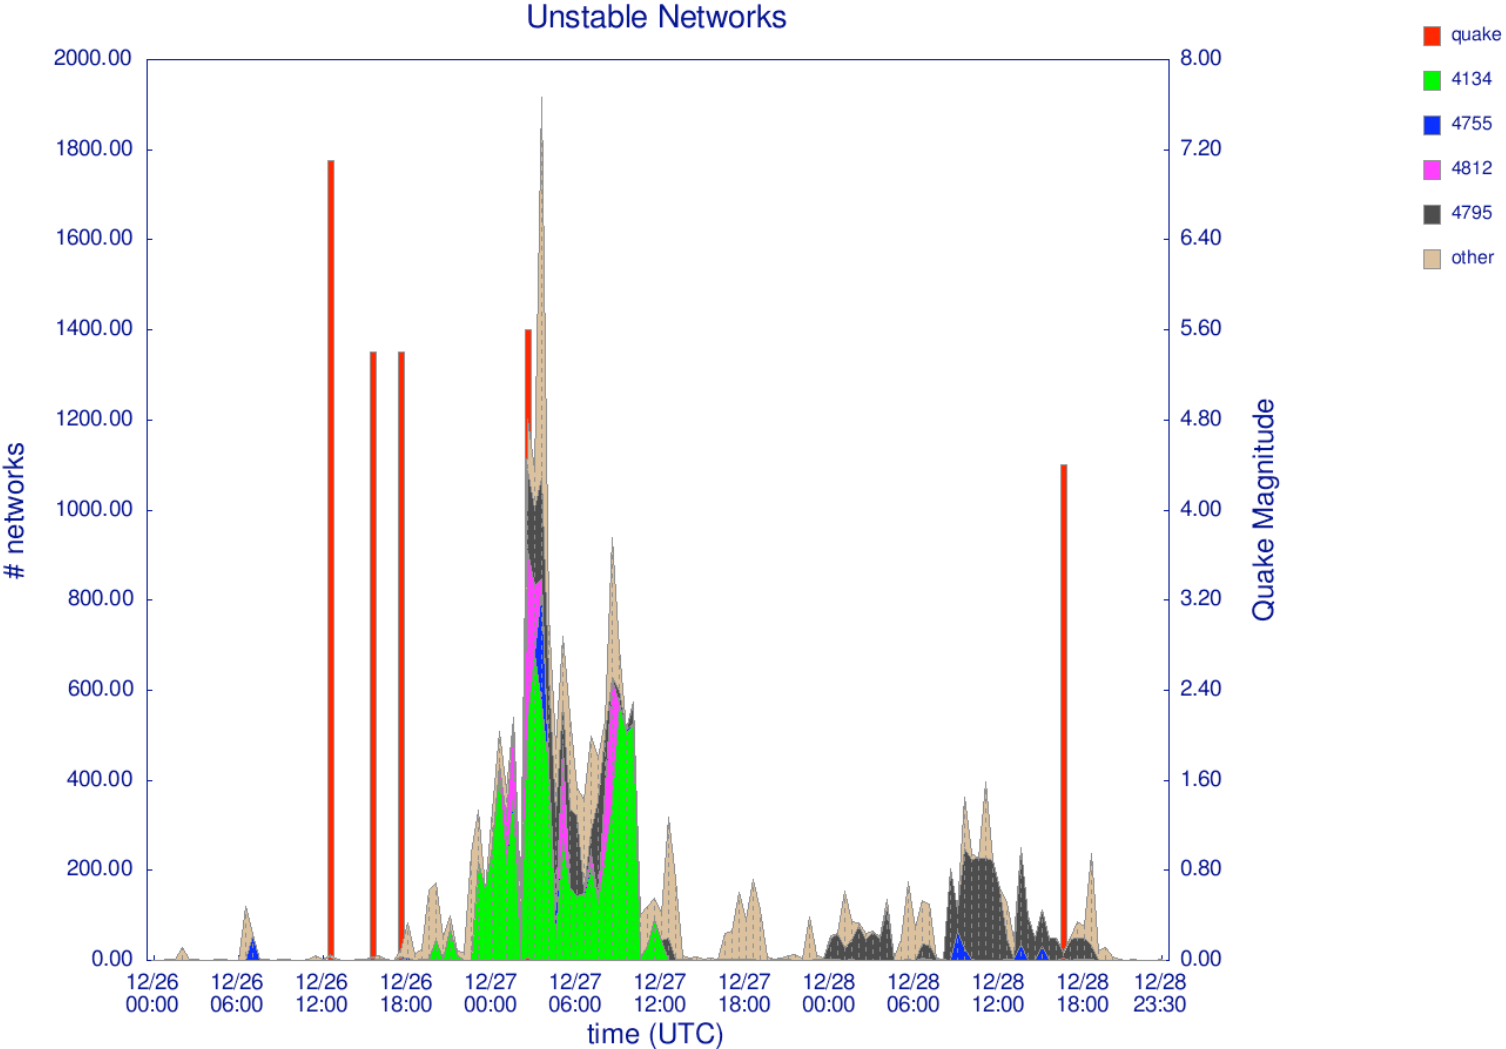
•Unstables by Country – 10 Day




•Unstables by Country – Peak



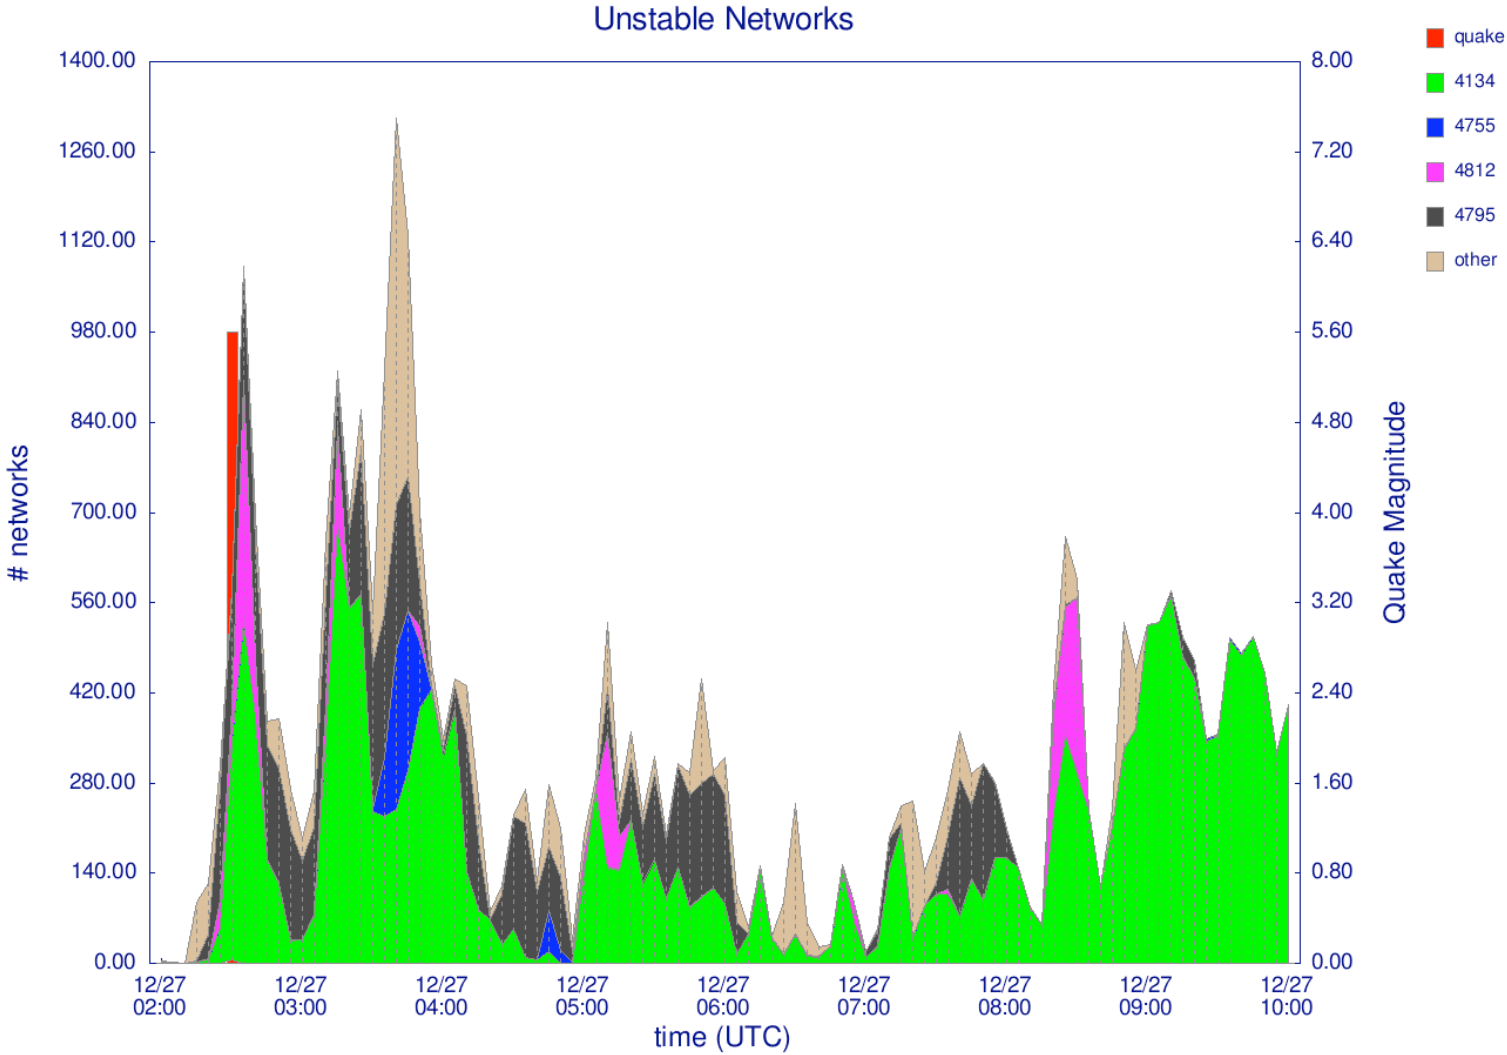
Unstables by Origin ASN – 3 day



Unstables by Origin ASN – 10 day



Unstables by Origin ASN – Peak




•Winners & Losers: By Country

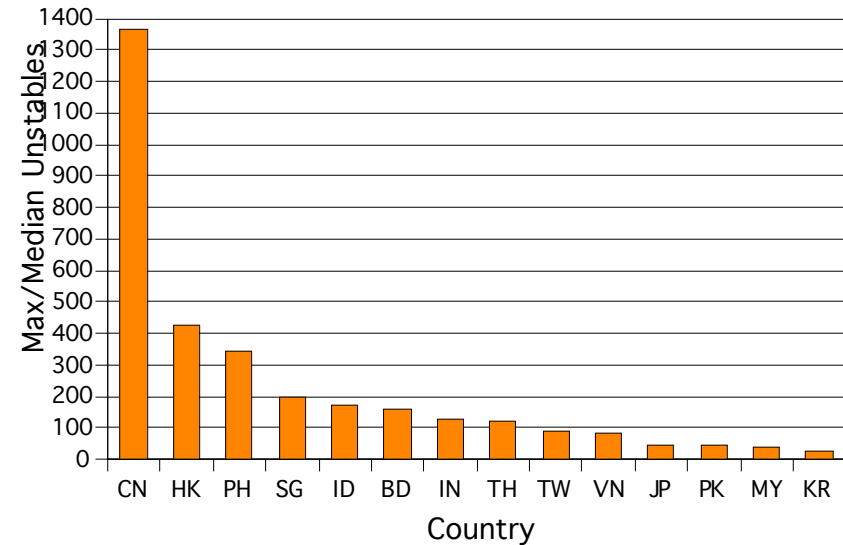
- Used maximum to median ratio of outages and unstable networks
- Worst Impacted:
 - China, Hong Kong
- Least Impacted:
 - Korea, Japan, Malaysia

•Winners & Losers: By Country (cont'd)

Max/Median Outages by Country




Max/Median Unstables by Country




Impacted ASNs

- Examined Asian prefixes outaged and/or unstable by origin AS – 1667 ASNs impacted:
 - China Telecom: AS4134, AS4812 (CN)
 - Sify: AS9583 (IN)
 - VSNL: AS4755 (IN)
 - Bharti BT Internet: AS9498 (IN)
 - PT Telekomunikasi: AS17974 (ID)
 - CNC Group (AS4808, AS4837) (CN)
 - Smart Broadband: AS10139 (PH)
 - INDOSAT: AS4795 (ID)

•Winners & Losers by ASN (cont'd)



of prefixes



% of total prefixes

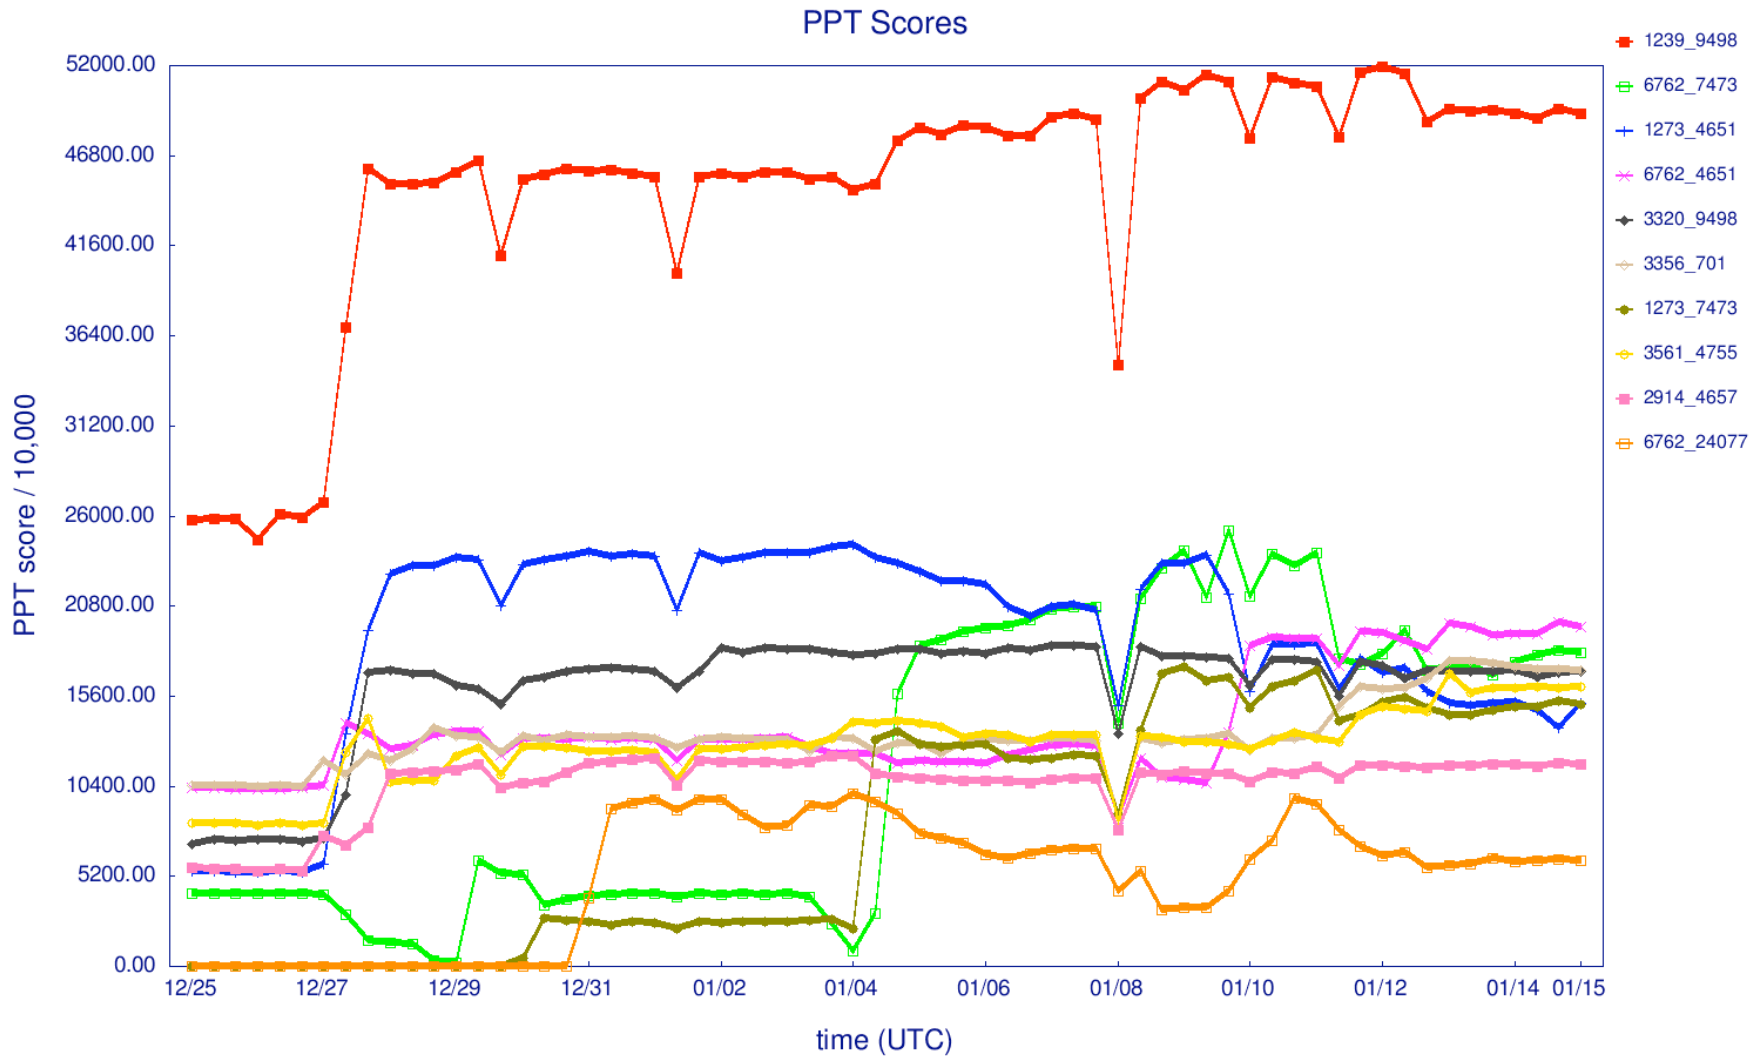
Edge Analysis

- PPT (Prefix, Peer, Time) score for each edge: for each prefix, for each peer, sum the amount of time the peer saw the prefix routed on the edge during a time interval
- Caveats:
 - All prefixes have the same weight
 - Cannot distinguish **between** an edge with a lot of prefixes seen by only few peers, **and** an edge with few prefixes seen by a lot of peers


Edge Analysis (cont'd)

- For each edge, generate time series with range 25 Dec 2006 to 15 Jan 2007 and 8 hours resolution
- Considered only prefixes from South-East Asia and the Indian subcontinent
- Filter out edges seen by less than 20 peers
- Evaluate the time series using the 2-day median difference between end and beginning

Top 10 Edge Winners




Top 10 Edge Losers



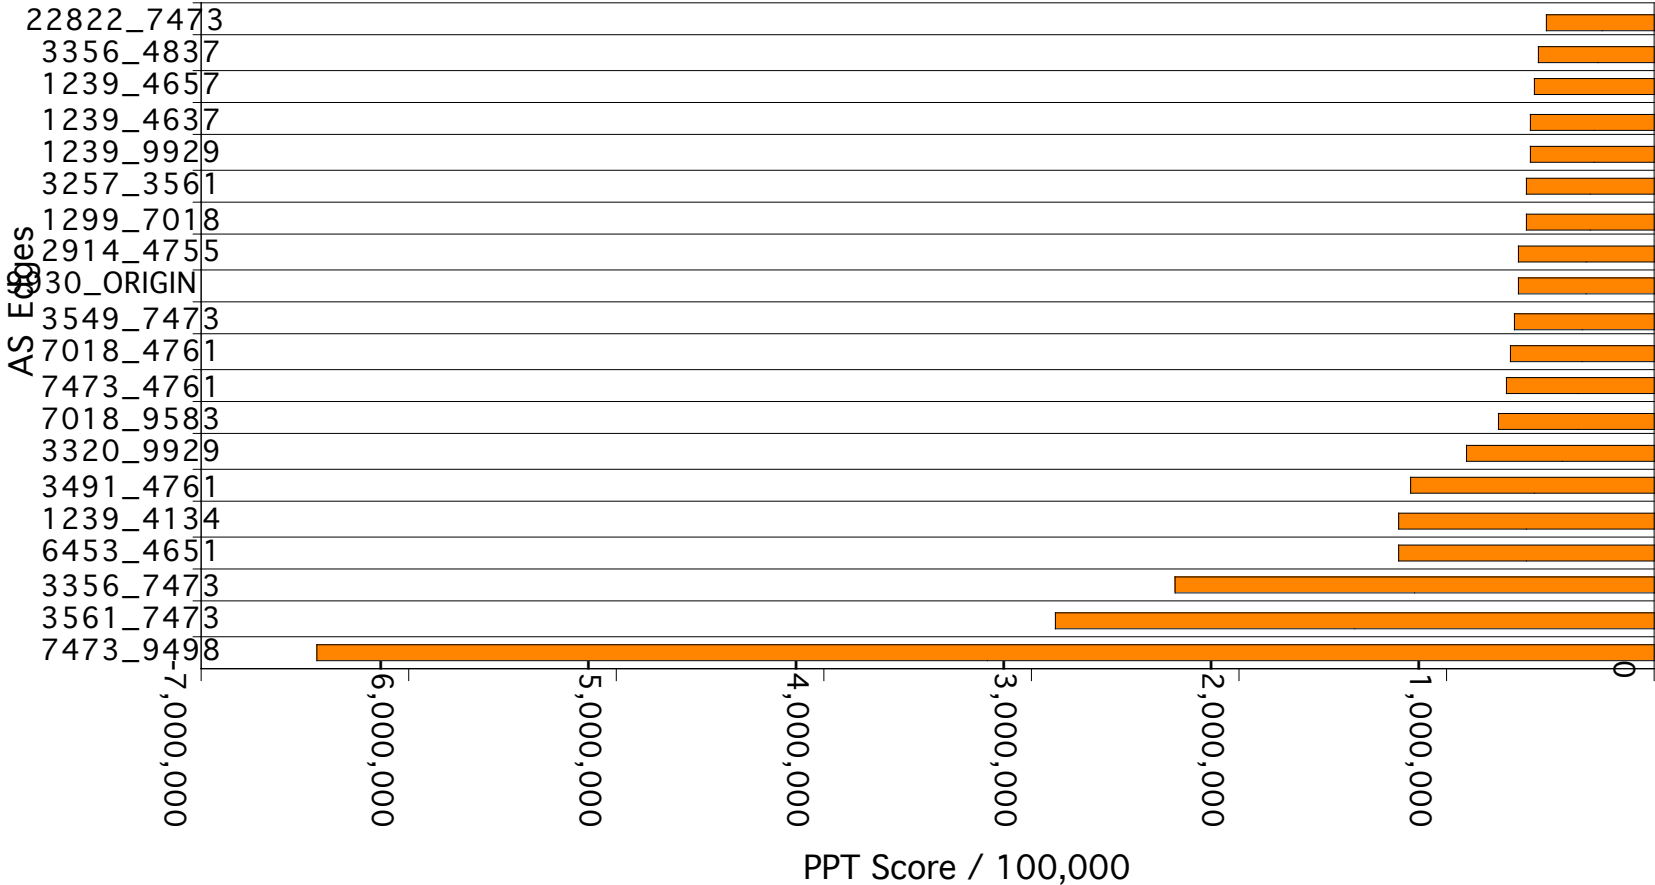
Winning Edges

Top 20 Winners



Losing Edges

Top 20 Losers




Regional Stories


- AS7473 Singapore Telecom (SG)
- AS4134 China Telecom (CN)
- AS9498 Bharti BT Internet (IN)
- AS4761 INDOSAT (ID)
- AS4651 Communication Authority of Thailand (TH)
- AS24077 TMHK Global Transit (HK)

Singapore Telecom (AS7473)



PPT Scores





Prefixes over Time



China Telecom (AS4134)




Bharti BT Internet (AS9498)




INDOSAT (AS4761)



PPT Scores





Prefixes over Time



Communication Authority of Thailand (AS4651)





TMHK Global Transit (AS24077)





Global Stories

- AS1239 Sprint
- AS3561 Savvis
- AS2914 NTT/Verio
- AS6762 Telecom Italia
- AS1273 Cable & Wireless



Sprint (AS1239)





Savvis (AS3561)





NTT/Verio (AS2914)



Telecom Italia (AS6762) – Winner



Cable & Wireless (AS1273) – Winner



Interesting Stories During Quake

- France Telecom (AS5511) provided temporary transit to Bharti (AS9498) from Dec 27 to Jan 5
- Indonesian routes move to INDOSAT (AS4761, AS4795) with transit mostly from DTAG (AS3320)
- China Netcom (AS9929) uses temporarily Sprint (AS1239) and DTAG (AS3320) as transits then drops them in favour of UUNet (AS701) and Savvis (AS3561)
- China Telecom (AS4134) routes move temporarily from Savvis to Sprint on Dec. 27

Interesting Stories After Quake

- Telecom Italia (AS6762) and Cable & Wireless (AS1273) are big winners adding Singapore Telecom (AS7473) and the Communication Authority of Thailand (AS4651) as customers
- Sprint (AS1239) gets to China Telecom (AS4134) through HiNet (AS9680) and Chunghwa Telecom (AS3462), i.e., 1239 9680 3462 4134

Conclusions

- Quake illustrates fragility of the global Internet
 - “Local” events can have broad impact
 - Physical failures can be difficult to remedy
- Asia is particularly vulnerable
- Impact will be felt long after the repairs are complete
 - New business relationships
 - New cable systems
 - Renewed interest in redundancy

Thank You

Todd Underwood todd@renesys.com

Alin Popescu alin@renesys.com

Earl Zmijewski earl@renesys.com