renesys

Rankings, Damned Rankings & Statistics

”There are three kinds of lies: lies, damned lies and statistics” – Benjamin Disraeli

MENOG 2, Nov 2007

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

Presentation Overview

- Background
- Why compute rankings?
- Who does rankings?
- What is wrong with rankings?
- Are there any good uses for rankings?
- Case studies of using rankings infrastructure (but not rankings themselves)
- Summary

renesys

Background – The Internet is an enormous graph

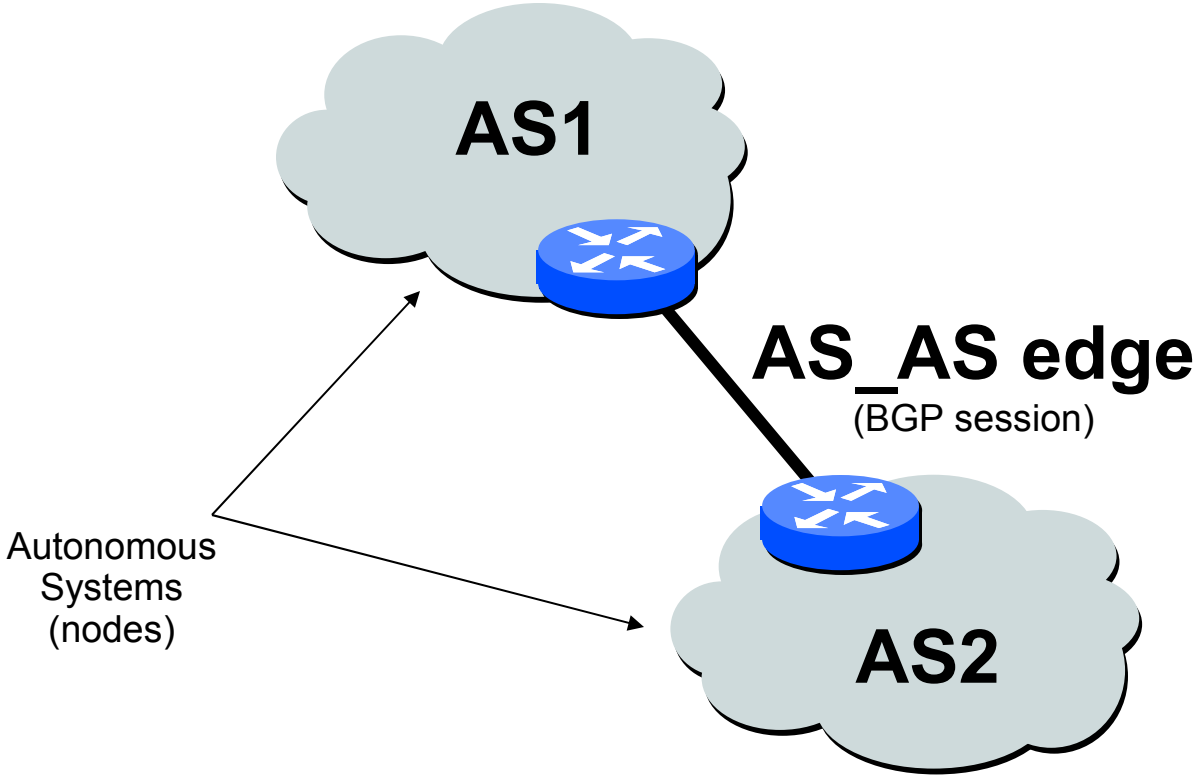

Background – Internet Routing depends on BGP

- **UPDATE** message: Announce new route or withdraw previously announced route.

UPDATE =
prefix + route attributes

- **Attributes:** Includes AS path to Prefix – e.g. 701 6543 8781
  Path data can be used to determine relationships between adjacent ASs, e.g., provider->customer, peer->peer, etc.
- **Prefixes:** Can be geo-located

renesys

Motivation for this talk …

- 2 years ago Renesys released Market Intelligence:
  - Focus on AS_AS edge relationship tagging and tracking
  - Geo-locate every prefix
  - Rankings (global and by geography)
  - Tried to pitch/target peering coordinators (too smart, too few, too broke)
  - Ended up used for sales/marketing/management
- But:
  90+% feedback (complaints) about rankings

renesys

Why Rankings?

Why do people care so much?

- Bragging rights. Mine is bigger ...
- Peering? Finding, evaluating, maintaining
- Marketing, Market selection
- Management oversight (uh-oh)
- Sales
- Engineering uses?

renesys

Rankings are everywhere

- Who has them:
  - CAIDA
  - Renesys
  - Netconfigs
  - Fixedorbit
  - Others
- What they're based on:
  - Fancy - relationship-tagged edges with scaled in-cone scores
  - Simple - # of prefixes, adjacencies

renesys

•Fixed Orbit “Knodes Index”

- “The Knodes Index is the best measurement of a network's connectivity to the Internet.”
- So we're done :-)
- “...average number of networks, or hops that must be traversed between any IP address on a given network to any other IP address on the Internet.”
- Big networks all the same.
- Lots of peering => High rank
- Useful?

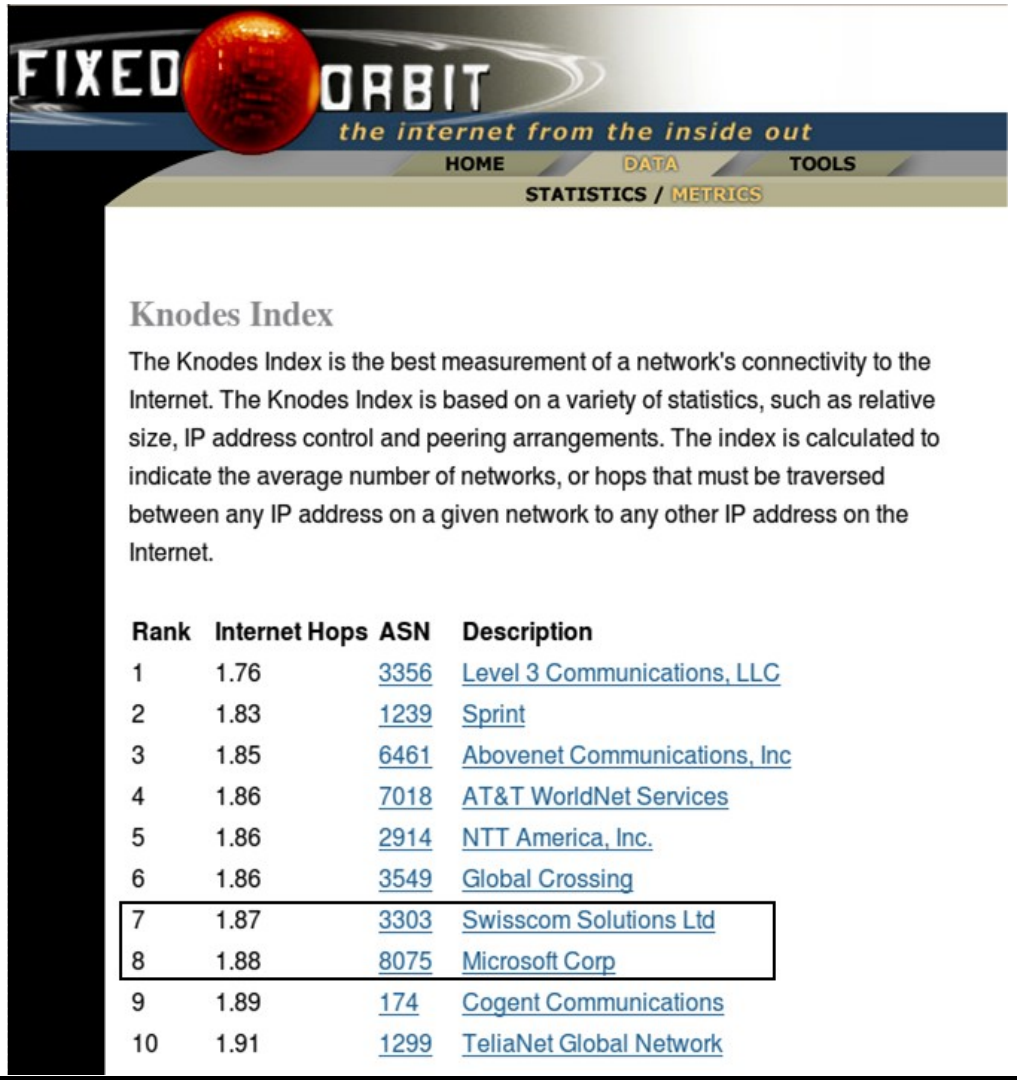

Knodes Index

The Knodes Index is the best measurement of a network's connectivity to the Internet. The Knodes Index is based on a variety of statistics, such as relative size, IP address control and peering arrangements. The index is calculated to indicate the average number of networks, or hops that must be traversed between any IP address on a given network to any other IP address on the Internet.

Rank	Internet Hops	ASN	Description
1	1.76	3356	Level 3 Communications, LLC
2	1.83	1239	Sprint
3	1.85	6461	Abovenet Communications, Inc
4	1.86	7018	AT&T WorldNet Services
5	1.86	2914	NTT America, Inc.
6	1.86	3549	Global Crossing
7	1.87	3303	Swisscom Solutions Ltd
8	1.88	8075	Microsoft Corp
9	1.89	174	Cogent Communications
10	1.91	1299	TeliaNet Global Network

renesys

•Netconfigs AS Rankings

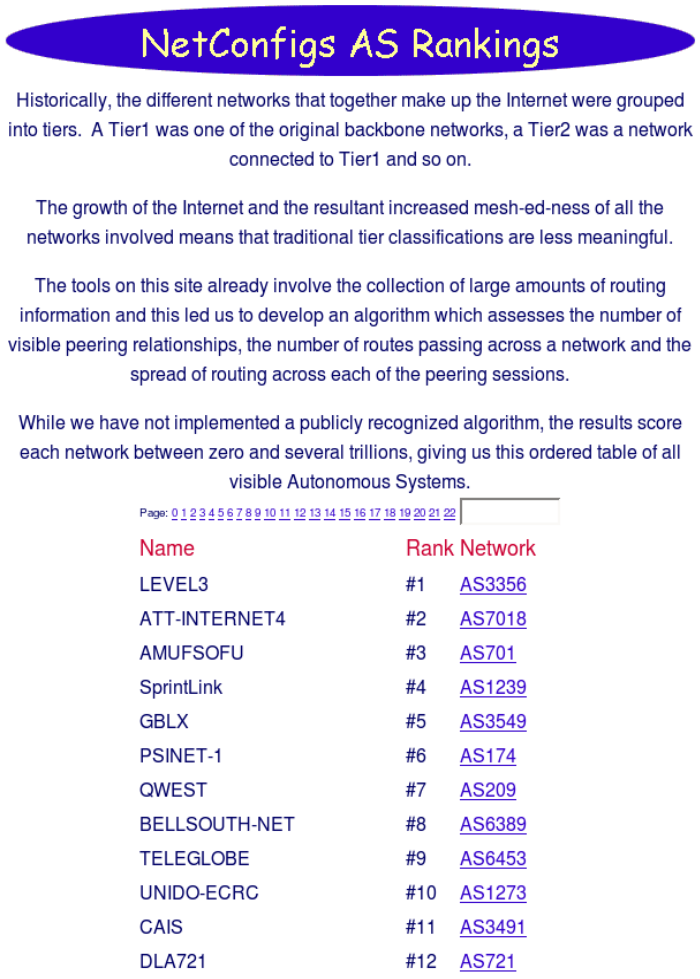

NetConfigs AS Rankings

Historically, the different networks that together make up the Internet were grouped into tiers. A Tier1 was one of the original backbone networks, a Tier2 was a network connected to Tier1 and so on.

The growth of the Internet and the resultant increased mesh-ed-ness of all the networks involved means that traditional tier classifications are less meaningful.

The tools on this site already involve the collection of large amounts of routing information and this led us to develop an algorithm which assesses the number of visible peering relationships, the number of routes passing across a network and the spread of routing across each of the peering sessions.

While we have not implemented a publicly recognized algorithm, the results score each network between zero and several trillions, giving us this ordered table of all visible Autonomous Systems.

Page: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22

Name	Rank	Network
LEVEL3	#1	AS3356
ATT-INTERNET4	#2	AS7018
AMUFSOFU	#3	AS701
SprintLink	#4	AS1239
GBLX	#5	AS3549
PSINET-1	#6	AS174
QWEST	#7	AS209
BELLSOUTH-NET	#8	AS6389
TELEGLOBE	#9	AS6453
UNIDO-ECRC	#10	AS1273
CAIS	#11	AS3491
DLA721	#12	AS721

- "Large amounts of routing information..."
- "...algorithm which assesses the number of visible peering relationships, the number of routes passing across a network and the spread of routing across each of the peering sessions."
- Strange names: Cogent is "PSINET-1", Verizon is "AMUFSOFU".
- Opaque
- Top-10 oddities …
  - Bell South (6389 at #8)
  - Cable & Wireless (1273 at #10)
  - Missing some big networks
    - NTT (2914 - #27)
    - Savvis (3561 - #42)

renesys

•CAIDA AS-Ranking

- Sophisticated, Best documented
- Actually tags AS_AS edges with **relationships.**
- Ranks size of the imputed "customer cone"
- Most useful of the free data.
- Quibbles:
  - Relationship tagging imperfect
  - Counting /24s doesn't mirror traffic well
  - Not run daily
  - Snapshots (not updates)

rank	AS number	AS information		customer cone			degree
		ISP's name	country	/24s	prefixes	ASes	
1	3356	Level 3 Communications, LLC	US	5,790,268	203,772	21,496	1,525
2	1239	Sprint	US	5,762,316	214,116	23,208	1,666
3	7018	AT&T WorldNet Services	US	5,598,886	202,927	21,762	2,068
4	701	UUNET Technologies, Inc.	US	5,380,772	199,773	21,722	2,626
5	174	Cogent Communications	US	5,248,608	193,814	21,040	1,723
6	3549	Global Crossing	US	5,125,837	190,913	20,309	983
7	3561	Savvis	US	5,104,900	188,405	19,946	495
8	7132	SBC Internet Services	US	5,094,856	184,525	19,942	700
9	702	MCI EMEA	NL	4,993,815	184,034	19,605	525
10	6939	Hurricane Electric	US	4,950,361	182,228	19,461	722

ranking mode: relationship based, pruning customer cone with inferred p2p links
alpha parameter: 0.01000
Whois: Feb 23, 2006 - AFRINIC, APNIC, ARIN, LACNIC, and RIPE
AS links: BGP RIBs from RouteViews (rv2)
AS links: 5 days starting on Oct 1, 2007 (15 snapshots at 8 hour intervals)
prefix-to-AS mappings: RouteViews BGP snapshot on Oct 3, 2007

renesys

Renesys Market Intelligence Rankings

- Based on relationship-tagged AS_AS edges and much more
- Multiple rankings for different purposes (more organizations get to be number one!)
- Customer base most popular: weighted size of on-net prefixes as routed

renesys

Renesys Ranking Algorithm

- Label the AS_AS edge set as one of:
  {customer-provider, provider-customer, peer-peer}
- Deal with anomalies: transit swaps, clustered ASNs
- Geolocate every prefix
- Reaggregate (“similarly routed” more specifics)
- Discount pre-CIDR allocations
- Ignore short-lived announcements (leaks)
- Score the remaining prefixes (nonlinearly) based on size (/8 - /24 only)
- Add up the scores in various ways; update daily
- Easy, no? Lots of possibly useful detail elided.

renesys

Route-based, Global Rankings are useless

- Global rankings by arbitrary metrics are not useful
  - Who cares who has the most aggregate on-net prefixes?
  - Who cares who has the largest number of AS adjacencies?
  - Notably incorrect for high-volume content (youtube, e.g.) and densely aggregated broadband consumer access networks.
  - Where is the **global traffic data?**
- Are there useful use-cases?

renesys

(Globally valid, representative)

Traffic Data is Non-Existant

- Routes are fundamentally *public*
- Traffic is fundamentally *private*
- Traffic data are unevenly collected, jealously guarded
- Routing data are global data from a local perspective
- Traffic data are local data. Unclear how to make representative.

Route-only Rankings Useful

- Who are the biggest providers of retail connectivity in Thailand (explore unfamiliar markets)
- Why did Qwest just lose market share to Cogent/Telia/Teleglobe ? (identify and track major changes)
- What new markets did my competitor just enter?

renesys

Renesys Market Intelligence Rankings

- RAYA dropped one place in the retail market in Egypt
  - AUCEGYPT reduced number of prefixes sent to RAYA (picked up Nile Online as a provider)
- TEDATA started transiting more prefixes (is a provider for Nile Online)

renesys

Forget About Rankings

- Rankings are one tool
- Much more interesting:
  - Edge tags in the AS_AS graph
  - Edge dynamics (the prefixes carried on various AS_AS edges)
  - Fabulous tools for analyzing large-scale changes

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

renesys

Uses of Edge Analysis

- Shifts in traffic
  - Natural disasters (e.g., Taiwan earthquakes)
  - Depeerings
  - Loss/gain of customers/providers
- Long term trends
- Geographic distribution of edges

renesys

Taiwan Quake Dec 2006 - Cable & Wireless (1273)

Gains more traffic from Communications Authority of Thailand (4651)
Gains Singapore Telecom (7473) as a new customer

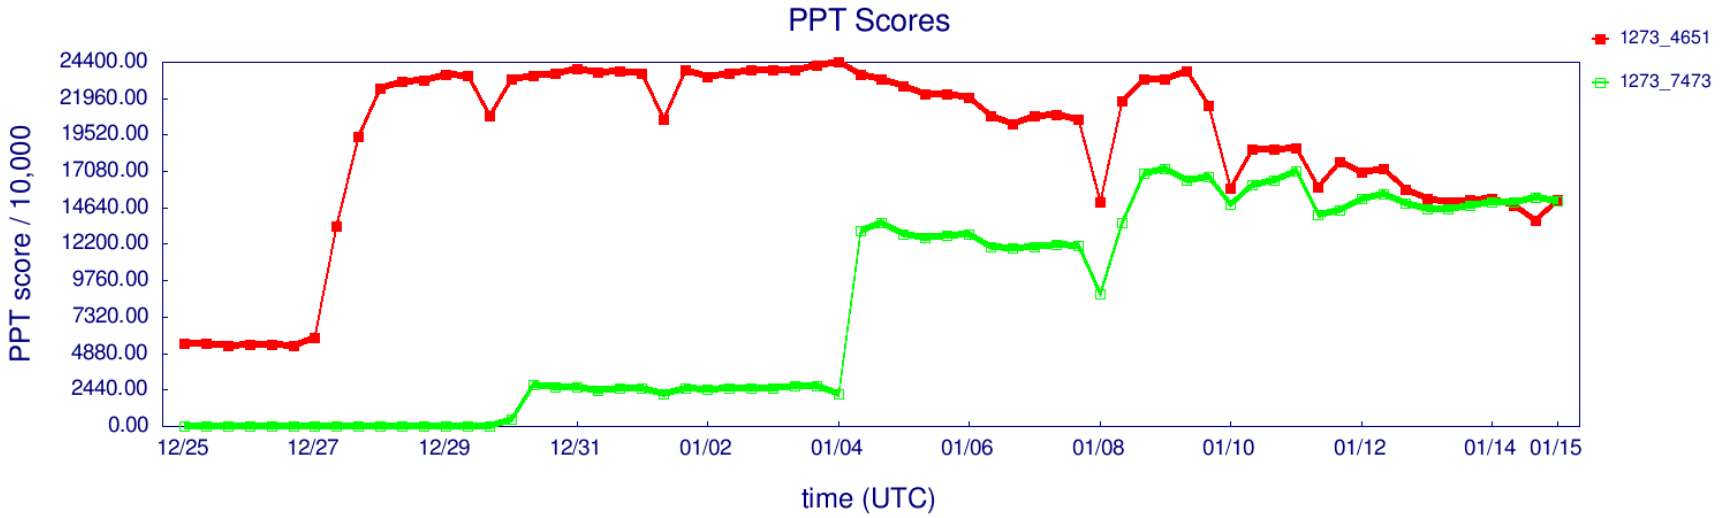

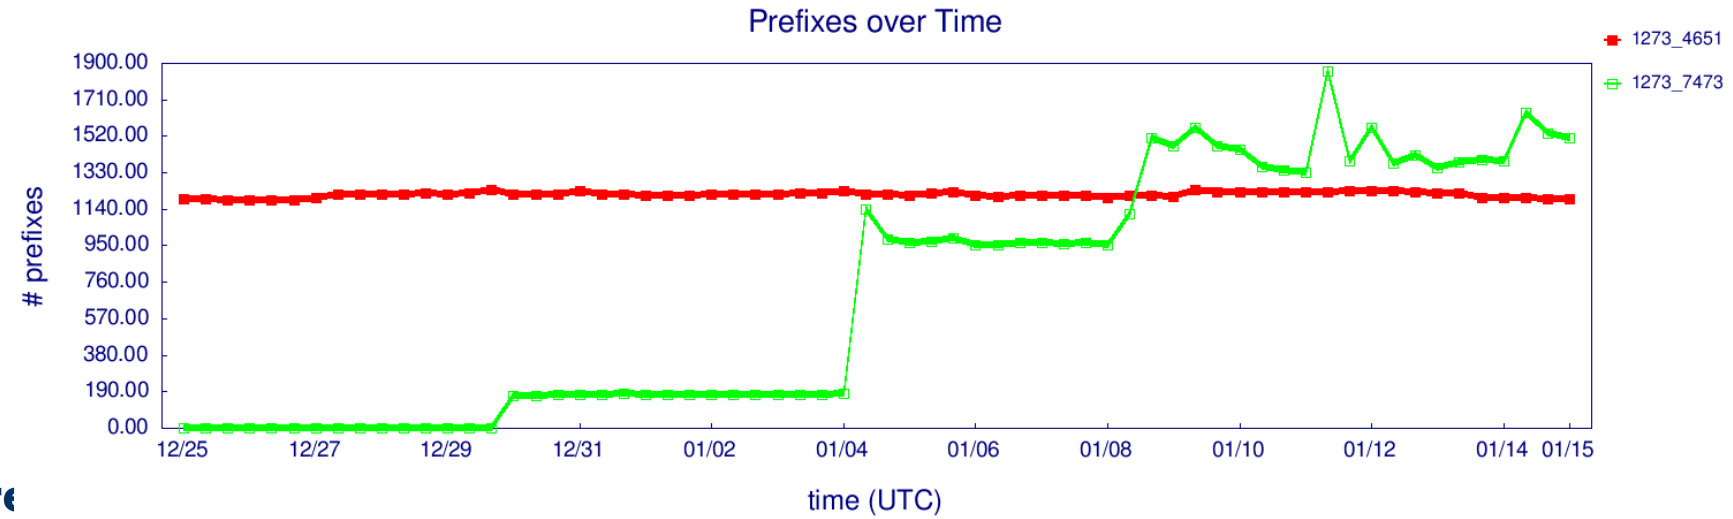

Cogent Depeerings: September 2007

- Cogent depeered more people
  - September 18
    - WVFiber (19151)
  - September 28
    - Peer1 (13768)
    - Limelight (22822)

renesys

Depeering - Cogent (174)

Cogent depeers Peer1 (13768), WVFiber (19151) and Limelight (22822)

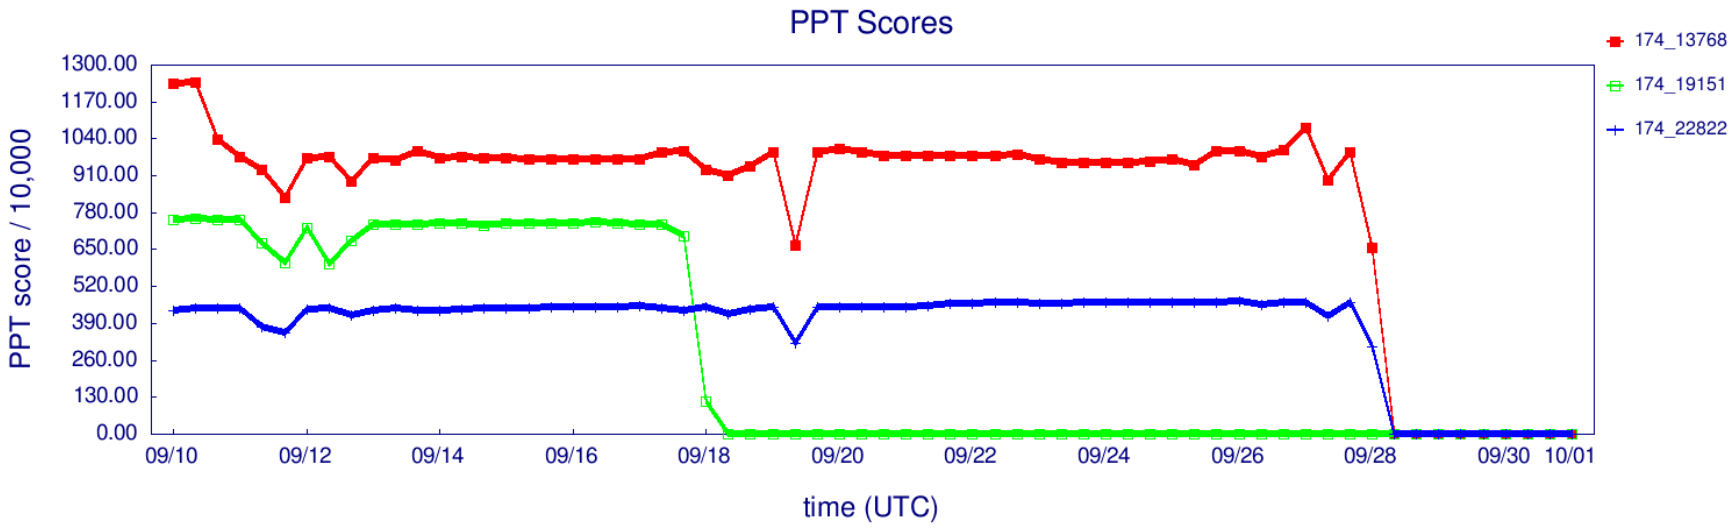

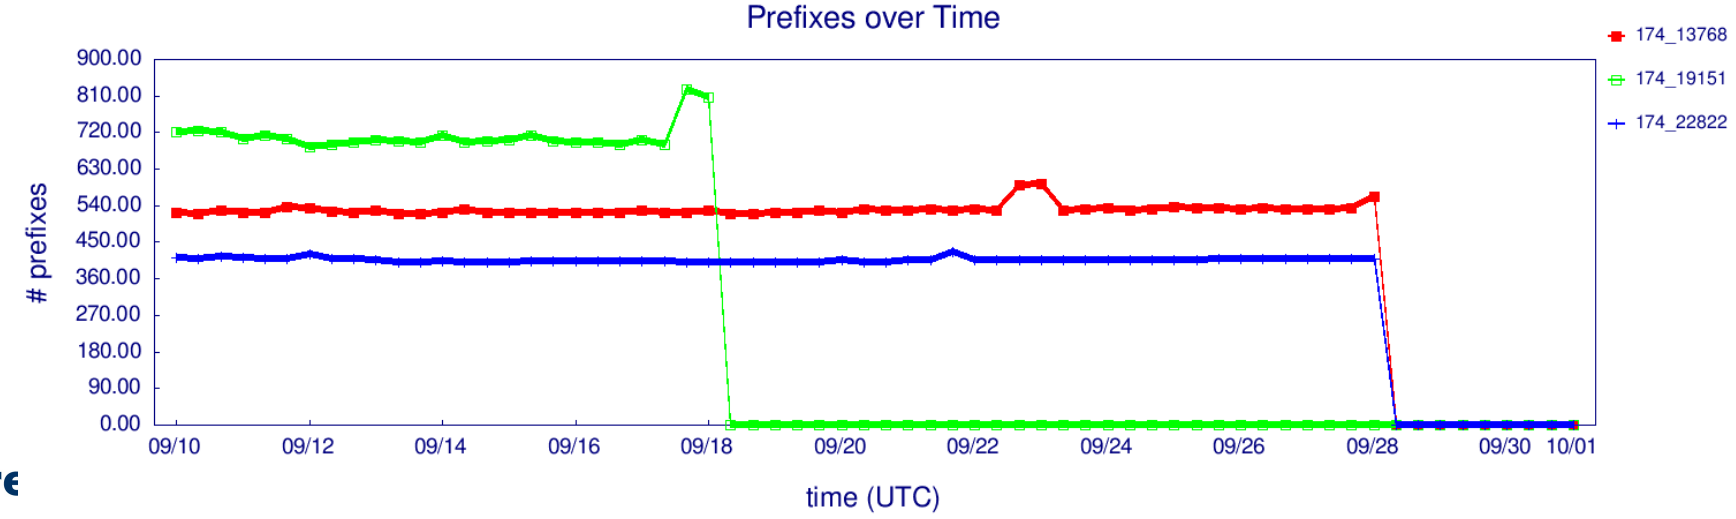

But edges are more complicated …

Where is the edge 701_1239 exactly?

- Edges are an abstract, not physical, concept
- Renesys sees 81,614 unique edges
- Prefixes seen on these edges can be geo-located
- Edges can “span” continents …

Continents	Edges
1	68,031
2	7,312
3	2,781
4	1,573
5	600
6	1,162
7	155

- Edges carry distinct sets of prefix “bundles”, each bundle geo-locating to a distinct region

renesys

Summary

- Worldwide global rankings are of little value.
- Specific rankings that make use of geo-location can be very useful.
- Routing changes and trends are immensely useful and can be indicative of …
  - Physical problems or human error
  - Customer wins and losses
  - Changes in traffic volume

renesys

Thank you

Alin Popescu alin@renesys.com

Todd Underwood todd@renesys.com

Earl Zmijewski earl@renesys.com