



Internet Captivity and the De-peering Menace

Peering Wars: Episode 1239.174

NANOG 45, Santo Domingo
January 2009

Martin A. Brown, Renesys Corp.
Clint Hepner, Renesys Corp.
Alin C. Popescu, Renesys Corp.

RENESYS ENTERTAINMENT COMPANY PRESENTS

PEERING WARS

EPISODE 1239.174

THE DE-PEERING MENACE



THE ADVENTURES CONTINUE...

Overview

- The Default-Free Zone (DFZ) and de-peering
- Clarify some terminology
 - (Transitively) single-homed, captives
- Brief history of de-peering events
- Who suffered from the most recent de-peering
 - Who were the captives
 - Geographic scope
- Who is likely to be affected by similar events in the future

The Default-Free Zone (DFZ)

- The Default-Free Zone (DFZ) is the set of ASes without a transit provider
- To ensure *global* connectivity, each AS in the DFZ must peer with all other ASes in the DFZ, i.e., form a clique
- We look at the DFZ from a routing, not a business perspective
 - We do not distinguish between paid or unpaid peering
- Spats between DFZ members will affect customers captive in their transit cones

The Default-Free Zone (DFZ)

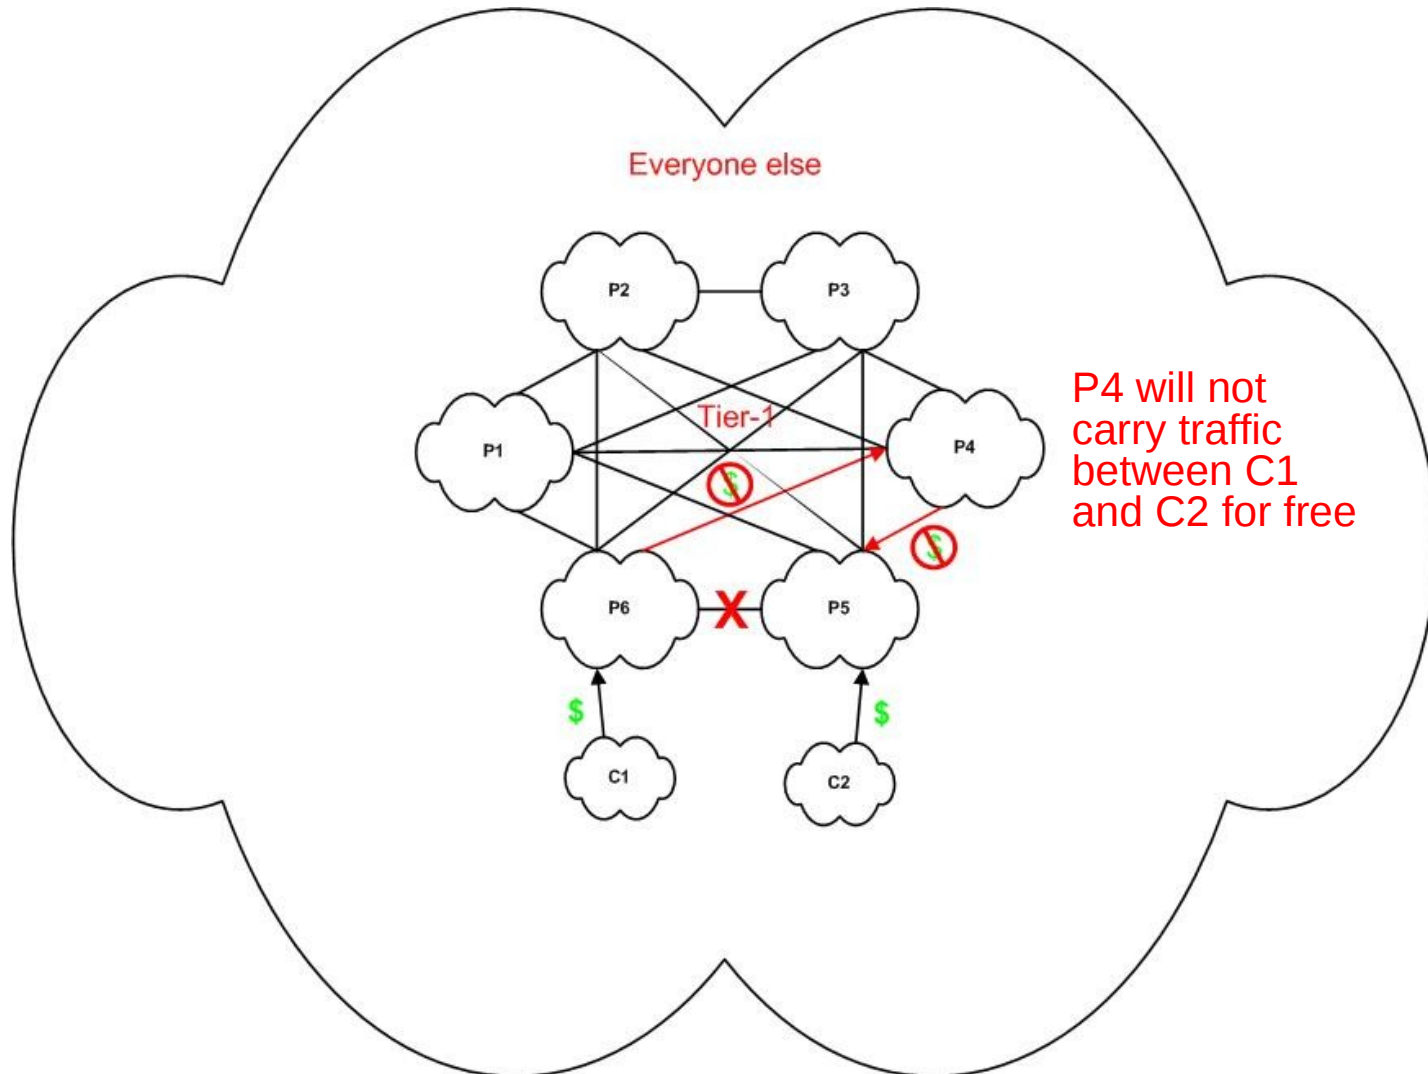
AS174	Cogent
AS209	Qwest
AS293*	Energy Sciences Network
AS701	Verizon
AS1239	Sprint
AS1299	Telia
AS2828	XO
AS2914	NTT
AS3356	Level 3
AS3549	Global Crossing
AS6453	Teleglobe
AS6461	Abovenet
AS7018	AT&T

* AS293 peers with all members of the DFZ, but is not a player in the commercial space


De-peering

- One peer in a peering relationship may be perceived to have disproportionate advantage
 - Traffic ratios are a common point of contention
- Aggrieved party may break the peering link
- Impact of broken peering link
 - Outside DFZ
 - Traffic follows transit links
 - Both parties pay more for transit
 - Internet remains whole, no partition
 - Inside DFZ
 - Traffic between ex-peers has nowhere to go, no transit links
 - The Internet is “broken” for innocent bystanders (captive customers)


De-peering Inside the DFZ



Being (Transitively) Single-homed



AS1 is
single-homed
behind AS2



AS1 is dual-homed but transitively
single-homed behind AS2

Definitions: Single-homed

- AS1 is **single-homed** behind AS2, if AS2 is the only provider of AS1
- AS1 is **transitively single-homed** behind AS2, if AS1's providers are *either AS2 or* transitively single-homed behind AS2 (note the recursion)
- Extension to prefixes: a prefix P is (transitively) single-homed behind AS1 if P is originated only by ASes that are (transitively) single-homed behind AS1

Another Definition: Captive

- A prefix is **suspected captive** to AS1 if all observed routes to that prefix contain AS1
- A prefix is **captive** to AS1 if all observed routes to that prefix contain AS1 and the prefix has no backup route
- Differentiating captives from suspected captives
 - Path analysis over time may reveal backup routes
 - A de-peering or outage

Observations on Captive Prefixes

- If an AS is captive behind AS1, all of its prefixes are also captive behind AS1
- Identifying captives:
 - De-peering events can provide proof of captivity
 - Maintenance windows expose backup routes
 - Routing advertisement history
- A captive prefix can be originated by a multi-homed AS
 - e.g., an AS with no internal connectivity and different providers at multiple locations

Real World Examples

- **AS11971** (Pfizer Inc.) is single-homed behind **AS7018** (AT&T)
- **AS40844** (Winn-Dixie Stores, Inc.) is transitively single-homed behind **AS7018** (AT&T), although it has two providers: **AS7018** (AT&T) and **AS6389** (BellSouth.net)

Real World Examples (continued)

- Prefix **116.66.128.0/24** (Cognizant Tech.) is originated by **AS17903** (Cognizant Tech.)
- **AS17903** has 5 providers:
 - **AS1239** (Sprint), **AS4755** (Tata), **AS7018** (AT&T), **AS9498** (Bharti), **AS18101** (Reliance)
- Most Renesys peers see routes to the prefix only via **AS1239** (Sprint), i.e., all AS paths follow this pattern: * 1239 17903
- Prefix **116.66.128.0/24** is **captive** behind **AS1239**
- Not all prefixes originated by **AS17903** are captive behind **AS1239**

High Profile De-peerings

- Oct 2005: Cogent (**AS174**) vs. Level3 (**AS3356**)
- Mar 2008: Cogent (**AS174**) vs. Telia (**AS1299**)
- Oct 2008: Cogent (**AS174**) vs. Sprint (**AS1239**)

Cogent – Level3 De-peering (Oct 2005)

- Partition lasted from 5 Oct 2005 to 7 Oct 2005
- Level3 had notified Cogent two months in advance of the de-peering
- Single homed customers
 - ~5100 prefixes for Level3 (~10% of transited prefixes)
 - ~2300 prefixes for Cogent (~5% of all transited prefixes)
- 4.3% of prefixes in the global routing table were partitioned as a result of the de-peering


Cogent – Telia De-peering (Mar 2008)

- Partition lasted from 13 Mar 2008 to 28 Mar 2008
- Most impacted geographic regions:
 - United States, served by Cogent
 - North-central Europe, served by Telia
- After the link was restored, evidence suggested that the dispute may have been about traffic ratios
 - Telia chose almost 3000 more prefixes via Cogent
 - Cogent chose 600 fewer routes via Telia
- 1.6% of prefixes in the global routing table were partitioned as a result of the de-peering

Cogent – Sprint De-peering (Oct 2008)

- Partition began on 30 Oct 2008 at 20:00 UTC and ended on 2 Nov 2008 at 21:00 UTC
- Prior to de-peering, Renesys peers saw
 - 8029 prefixes from Sprint to Cogent
 - 2875 prefixes from Cogent to Sprint
- After re-peering, Renesys peers saw
 - 7356 prefixes from Sprint to Cogent
 - 2791 prefixes from Cogent to Sprint
- 3.3% of prefixes in the global routing table were partitioned as a result of the de-peering

Prefixes (NOT) Carried



- Each point corresponds to the number of prefixes seen by Renesys peers on the 174_1239 (Cogent/Sprint) and 1239_174 (Sprint/Cogent) edges during an 8-hour interval

Captives Behind Sprint


- 214 ASes were single-homed behind Sprint
- 6603 prefixes were captive behind Sprint, out of which 857 of them were registered to Sprint
- Interesting captives
 - 246 prefixes from Sprint PCS
 - US Dept. of the Interior, US Dept. of Justice, US National Park Service, US Bureau of Reclamation
 - Commonwealth of Massachusetts
 - 65 educational institutions (e.g., Brandeis University)
 - Northrop Grumman, Pfizer, Merck

Captives Behind Cogent

- 289 ASNs were single-homed in Cogent's cone
- 2349 prefixes were captive behind Cogent
- Interesting captive prefixes
 - NASA
 - Maryland Dept. of Transportation, NY Court System
 - 63 educational institutions (e.g., Rider University, York University)
 - GMAC Mortgage, ING Canada

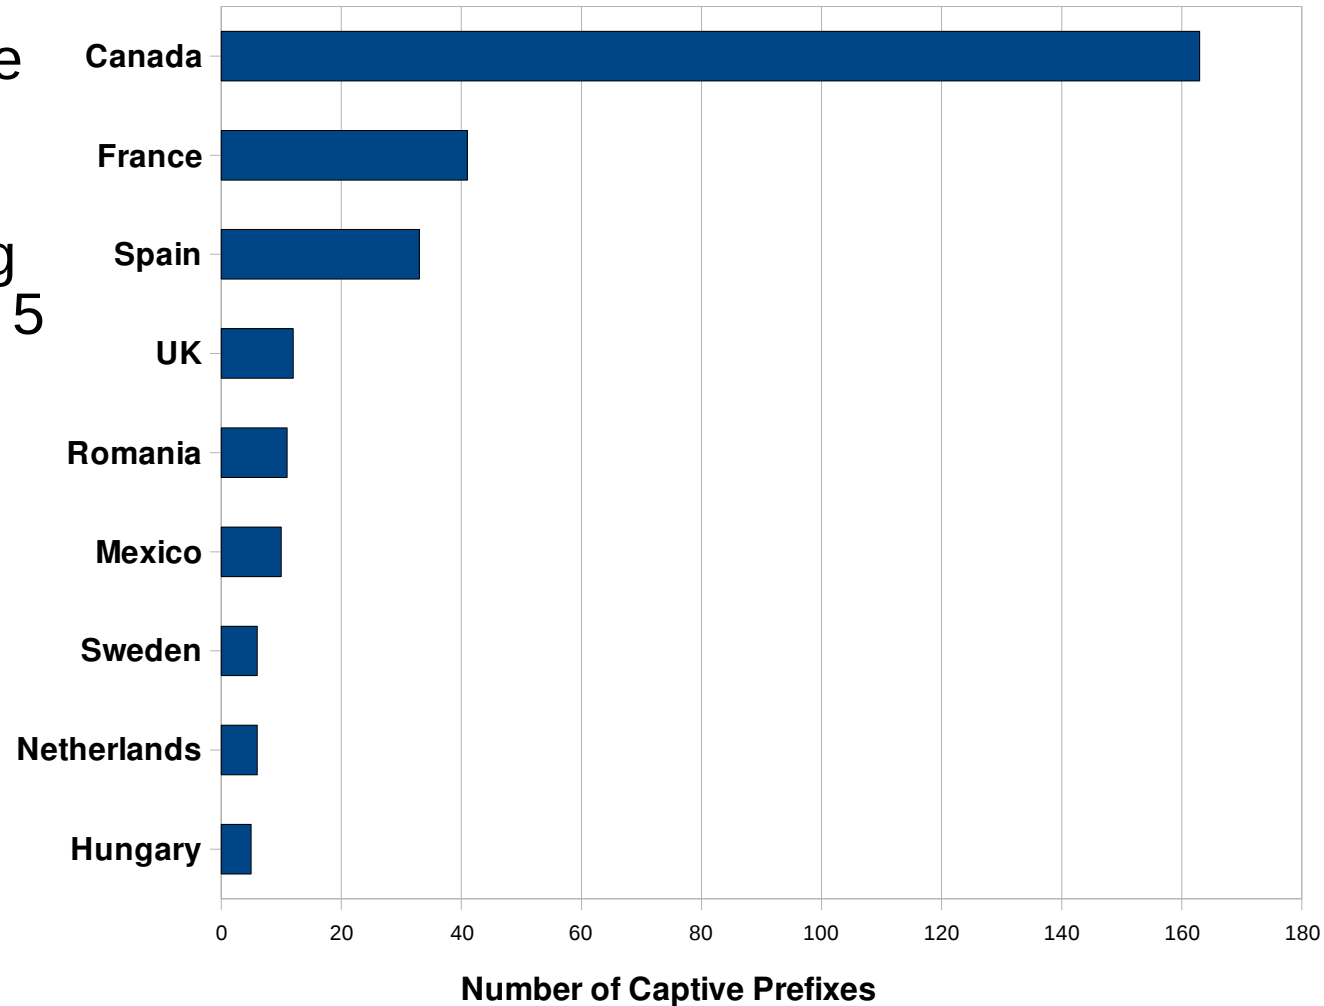
Sprint's Captives by Country

- Heavy US presence (5035 prefixes)
- Shown in chart are countries (excluding US) with more than 10 captive prefixes



Cogent's Captives by Country

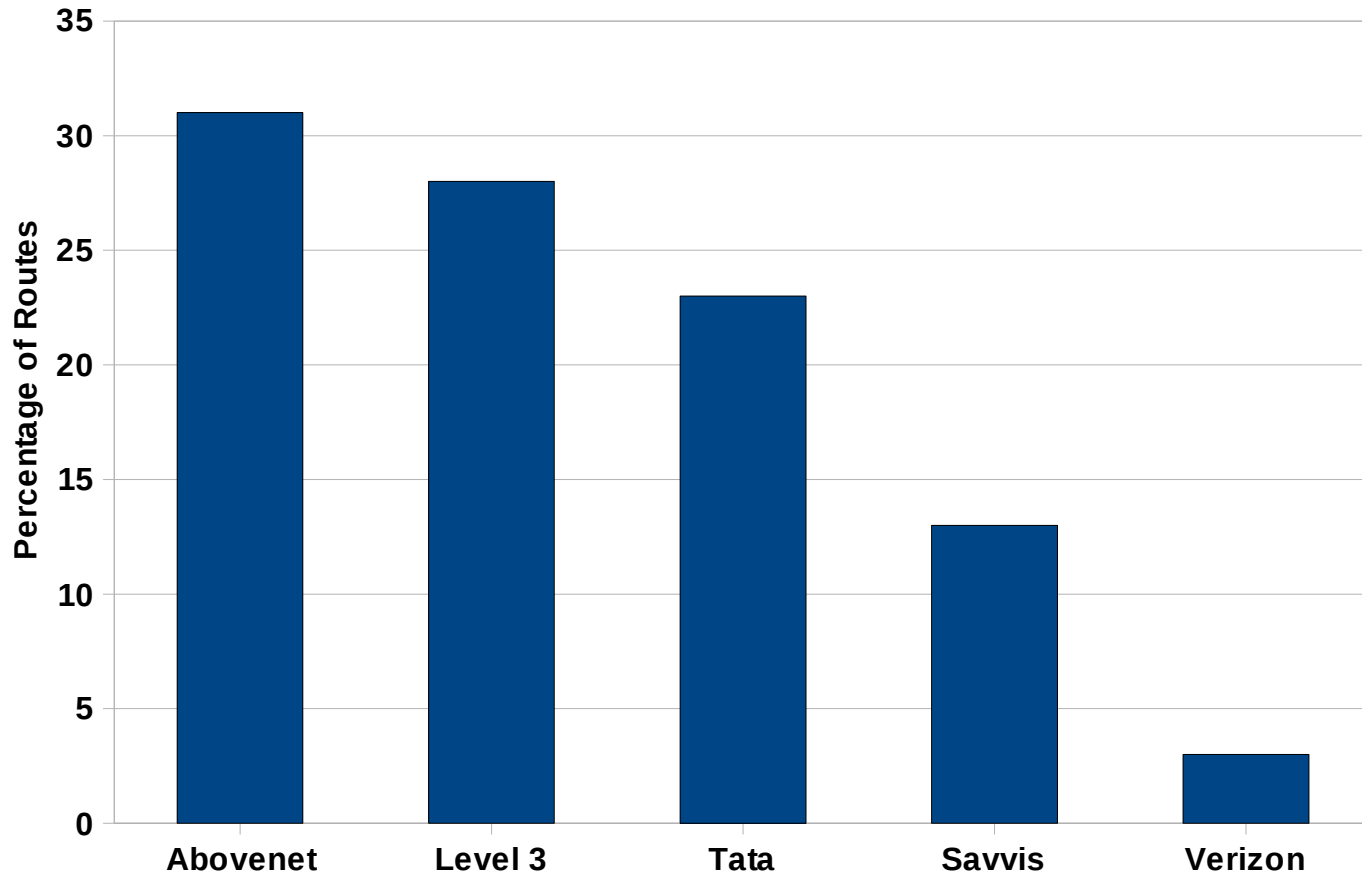
- Strong US presence (1766 prefixes)
- Shown in chart are countries (excluding US) with more than 5 captive prefixes



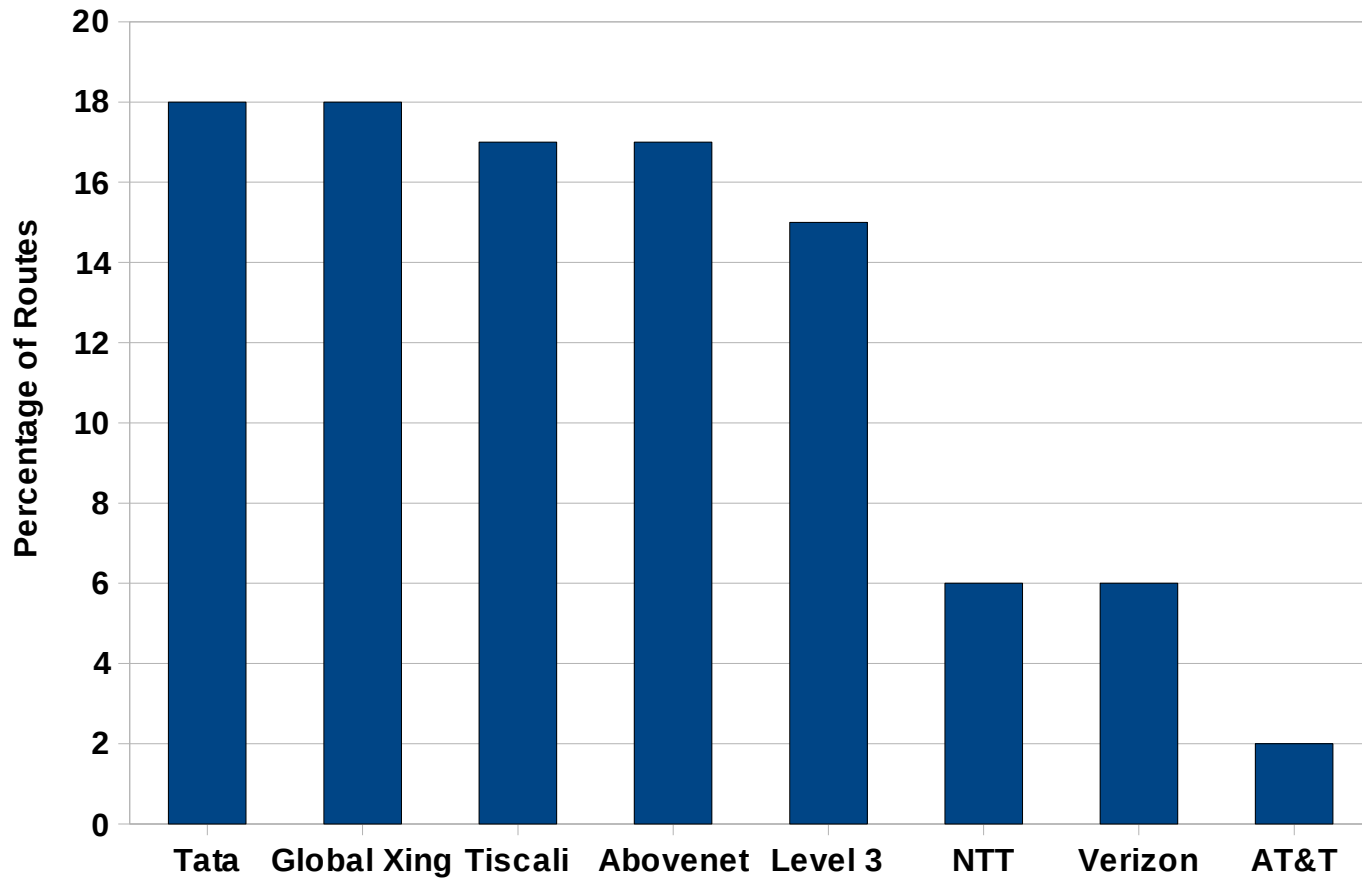
Route Winners (for non-captives)

- 1426 of Sprint's prefixes and 526 of Cogent's prefixes on the peering link were reachable via alternate paths
- Winners are the providers replacing one combatant to reach the other combatant
- A single prefix may be won by multiple providers, depending on the peer reporting the new path

Who won Cogent's prefixes?





Who won Sprint's prefixes?



The Aftermath

- Peering link restored on 2 Nov 2008 at 21:00 UTC
- Sprint's menace
 - The restoration was only a “temporary reconnection”, and *re-de-peering* would follow unless the issue was resolved
- On 22 Dec 2008 peace ensued
 - “Sprint and Cogent announced that they have reached a multi-year interconnection agreement for the purposes of exchanging Internet traffic. This agreement will benefit the customers of both Sprint and Cogent and resolves the earlier dispute to the satisfaction of both parties. The agreement is in accordance with both parties' previous and long standing interconnection policies and agreements. The specifics of this agreement are confidential.” (source: <https://www.sprint.net/cogent.php>)

Suspected Captives in the DFZ



AT&T Suspected Captive Prefixes

- **Telecom**
 - Mediacom Communications, Cable One, Windstream Communications, Charter Communications, Northland Cable TV, Fairpoint Communications, Heartland Communications
- **Financial**
 - Citicorp, The Vanguard Group, Countrywide Home Loans, Deutsche Bank, ING
- **Universities**
 - St. Louis University, State University of New York (SUNY), Princeton, Arizona Tri-University Network (ASU, UA, NAU)
- **Other interesting captives**
 - UPS, Honeywell, Apple, Staples, Lockheed Martin, Alcoa, Pfizer, National Weather Service, Entergy

Conclusions

- Intelligent multi-homing is good, but increases global routing table size
- No single DFZ provider can guarantee global connectivity
- De-peering events in the DFZ can do significant damage to captives
- Risk for a large spectrum of organizations: small and large, government, commercial



Thank You

Martin A. Brown mabrown@renesys.com

Clint Hepner clint@renesys.com

Alin C. Popescu alin@renesys.com