

# *Enhancing the Internet's Administrative Look-up Service*

Mark Kosters  
VeriSign Labs  
ARIN 8, October 2001

- Universal Whois
- VeriSign has committed undertaking in agreement with ICANN
- Public consultations (3 formal, ongoing informal)
  - business, intellectual property holders (Aug 2001)
  - international input (Nov 2001 at end of ICANN meeting)
  - civil liberties, other ngo's (TBD)
- Informal Meetings
  - RIPE (Oct 2001)
  - NANOG (Oct 2001)
  - ARIN (Oct 2001)
  - APRICOT (Mar 2002)

- universal whois
- non-centralized
- not specific to particular tld registry
- non-proprietary, open standard as outcome

- coordinating
- commitment
- listening
  - Web page  
<http://uwho.verisignlabs.com>
  - Mailing list  
[Uwho@lists.verisignlabs.com](mailto:Uwho@lists.verisignlabs.com)  
Use [uwho-request@lists.verisignlabs.com](mailto:uwho-request@lists.verisignlabs.com) to subscribe
- would welcome input from the ARIN community

- **Perhaps the best methods are not on port 43.**
- **Points of exploration**
  - RWhois
  - Whois++
  - Referral LDAP Service - take advantage of the many LDAP standards.
  - XLDAP - looking at a layered XML approach similar to EPP.
- **This work actually started before VeriSign committed to Appendix W.**
- **Work located at <http://www.verisignlabs.com>**

- Re-use known technology.
- Don't invent too much
  - Does the world really need another application transport?
  - Does the world really need another schema language?
- Structured queries and structured results for better machine readability.
- Referrals to bridge domain registries and registrars and other needs for entity reference.
- Client authentication to address privacy concerns
  - Provide the mechanism, not the policy
- Toolkits should be readily available.

- May be one part/component of the answer
- RWhois was primarily developed to handle network reassignment
- No protocol enhancements for a couple of years
- What is the problem
  - Usability issues
  - Backend built when no freely available 3rd party databases
  - Documentation lacking on installation/operations/registration
- Question - What is the appropriate way to help ARIN to prioritize development resources to enhance RWhois

# RWhois Update

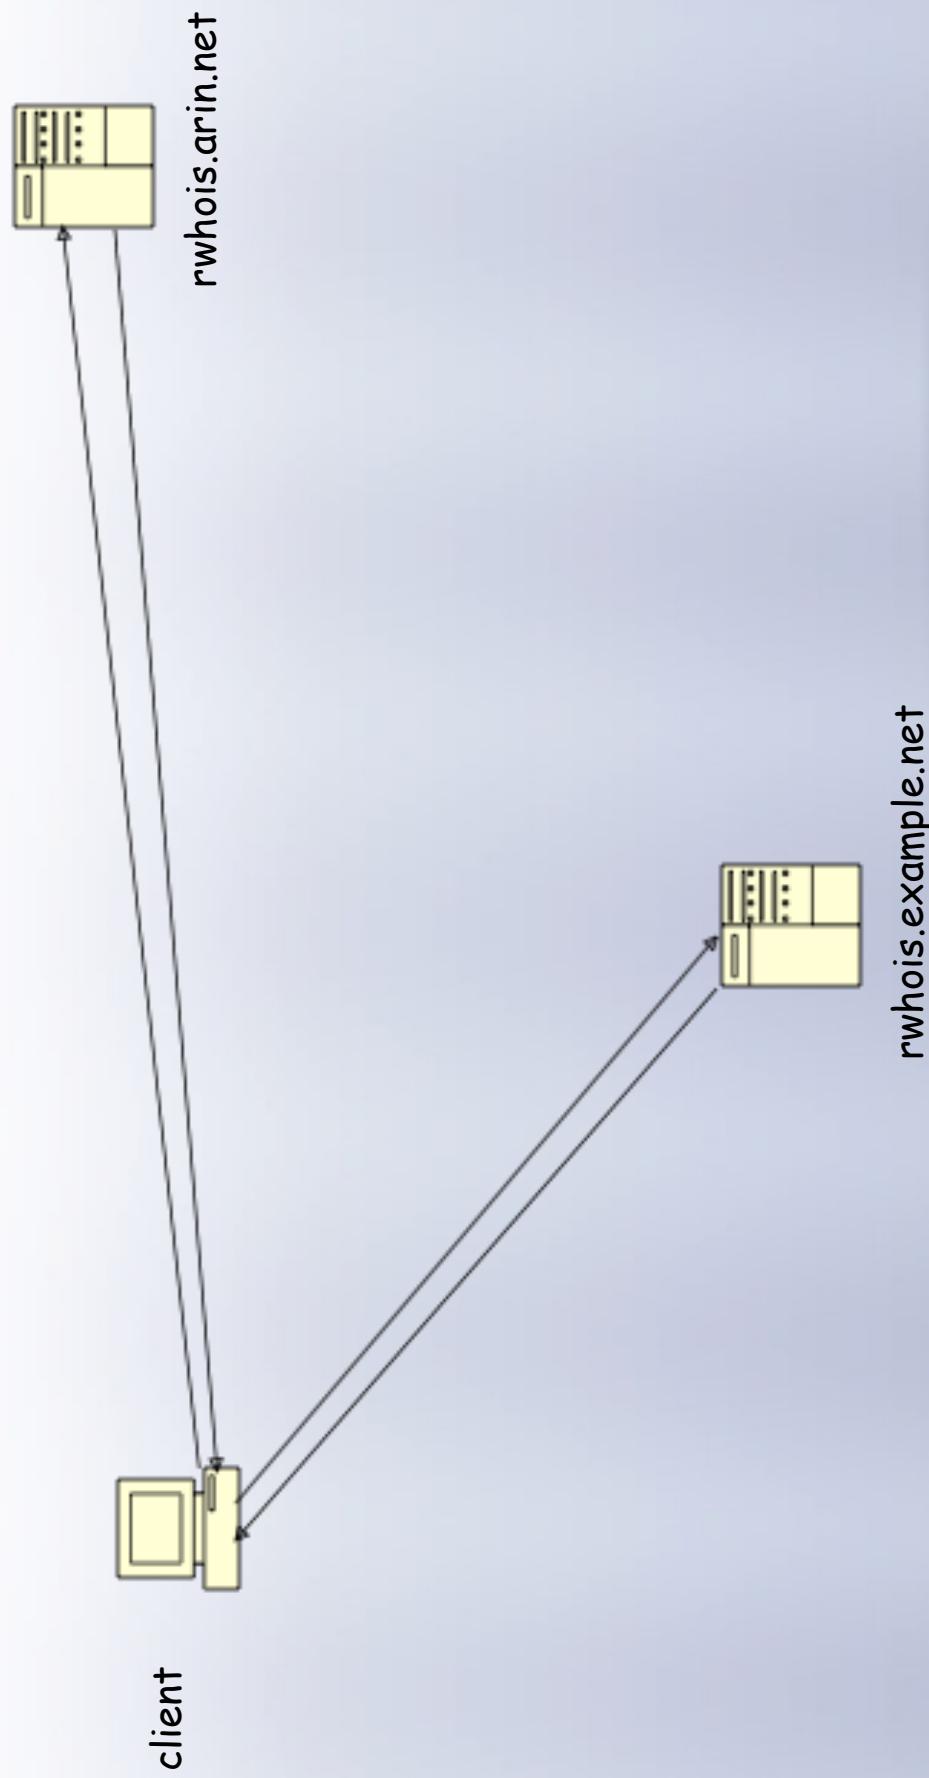
Cathy Murphy  
Principal Software Engineer  
[<cathym@arin.net>](mailto:cathym@arin.net)

# Overview

- What is RWhois?
- Why do we need it?
- Who is using it?
- What is wrong with it?
- How do we fix it?
- What is its status?
- What's next?

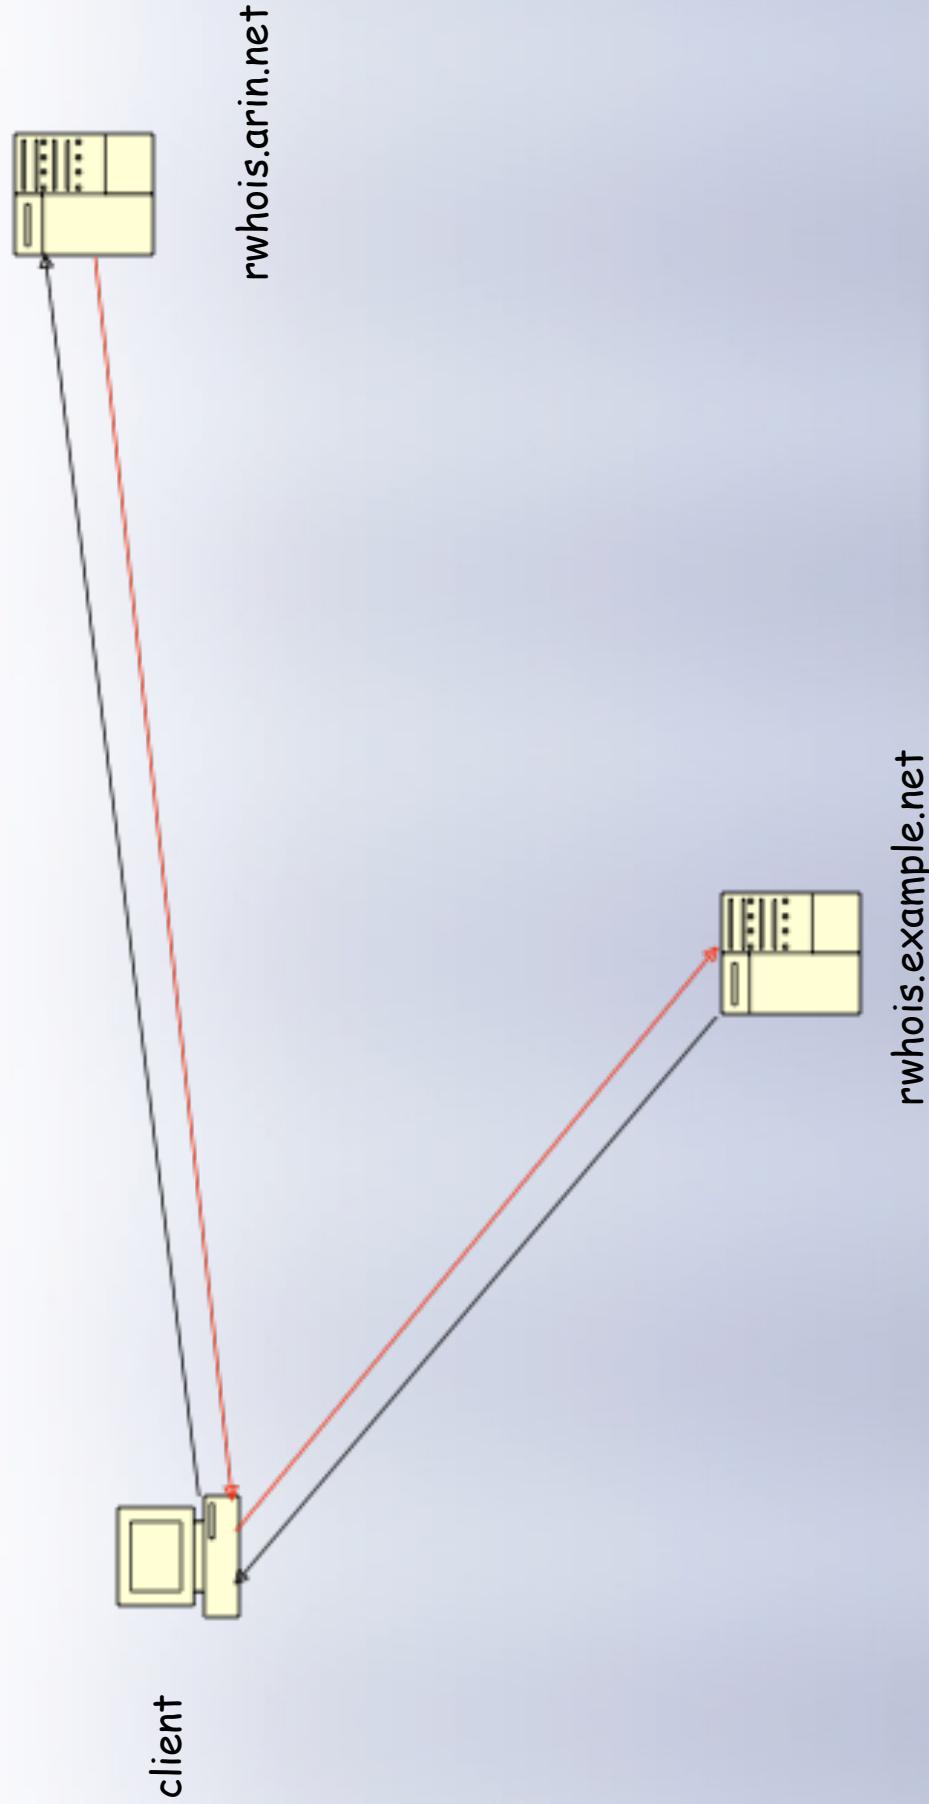
# Current RWhois

- What is RWhois? •



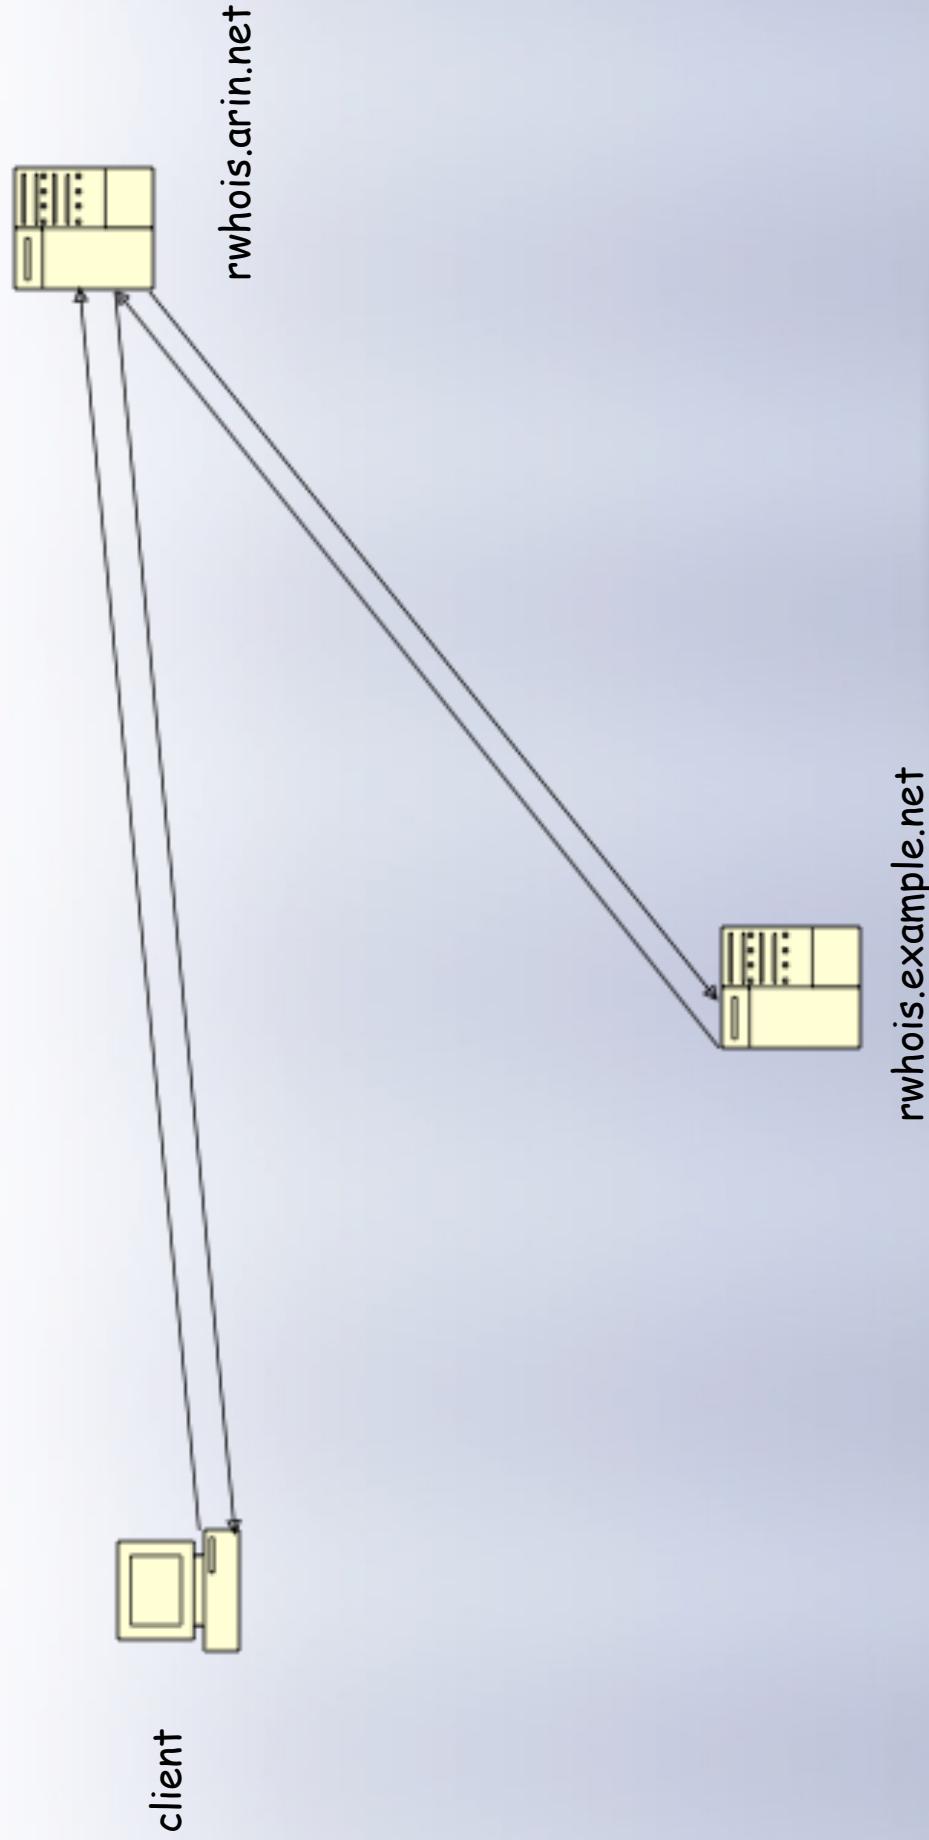
# Client Referral

- What is RWhois? •



# Query Routing

- What is RWhois? •



# Purpose of RWhois

- Why do we need it? •

- Technical Support
  - Public availability of technical contacts for networks
- Justification
  - To document 80% usage of existing address space as justification for future allocations
- Local control of reassignment data
  - More accurate information
  - More frequent update of information
  - Easier to maintain
- Good engineering practice to keep the data in just one place and as close to the source as possible.

## Current Installation Base

- Who is using it? •

- 263 Maintainers have told ARIN that they run RWhois servers
- 129 referrals in ARIN's RWhois
- 99 do not currently respond
  - 30 are active
- 134 do not have referrals

# Deficiencies As Identified by DBWG

- What is wrong with it?
- Query routing doesn't work
- Uses a non-standard database format
  - file structures != database
- Insufficient documentation
- Set-up and administration is just too difficult

## Additional Deficiencies

- What is wrong with it? •

- Code Architecture
  - Not conducive to updating
  - Not conducive to adding modular functionality
    - Query format
    - Database backend
    - Output format
- Code Audit
- Management Tools
  - Existing tools are of limited use
  - Not enough of them

# Goals

- How do we fix it? •

- Query routing doesn't work
  - Fixed per ARIN's beta release
  - Production implementation not yet scheduled
- Uses a non-standard database
  - Make the default backend an open source solution, e.g., MySQL
  - Make the backend a configurable component
- Insufficient documentation
  - Develop new FAQs and how-to's
  - Provide a more active support mechanism
- Setup and administration is just too difficult
  - Develop a canned set of working configuration and data files
  - Simplify, simplify, simplify!

# Goals

- How do we fix it? •

- Code Architecture not conducive to update
  - Develop new modular code base
- Code Audit
  - Develop new modular code base
- Management Tools
  - Develop new tools
  - .

# Current RWhois Situation

- **What is its status? •**

- ARIN hired a software engineer to work full time on RWhois
  - Adam Guyot [adamg@arin.net](mailto:adamg@arin.net)
  - Sept '01 – ARIN updates RWhois server software (rwhoisd-1.5.8 [beta])
    - allows RWhois server to perform referral queries
    - requires *libstr* regex library from Ralf Engelschall running internally at ARIN; outstanding issues approximately 30 downloads
  - Oct 4 '01 – VeriSign provided a bug fix release (rwhoisd-1.5.7.1 [beta])
  - Oct 26 '01 – ARIN provided bug fixes for recently discovered exploit
    - Oct 30 '01 – VeriSign incorporated bug fixes (rwhoisd-1.5.7.2)

# What's Next?

- New RWhois code base?
  - Client referral?
  - Server referral?
  - Both?
- Universal Whois?
- Other solution?
  - <http://www.freeipdb.org/>

# David Huberman

## Global Crossing

- 2001-1: Changing minimum SWIP size from /29 to /28
  - Proposal:  
It is currently required all /29 and shorter reassigned be reported to the ARIN WHOIS database via SWIP or RWHOIS. It is proposed this policy be modified to require reporting for /28 and shorter reassigned only.

# RTMA Agenda

---

RTMA Update

Presentation of Cengiz's Ptomaine Slides

Andrew Dul – Discussion of Policy Proposal 2001-6

Single organizations with multiple aggregation points

Lee Howard - Discussion of Policy Proposal 2001-2

Multi-homing justification for /24 reassigned

Bill Manning – Discussion of Policy Proposal 2001-5

Micro-assignments for multi-homed sites

Any Other Items??

# RIPE/RIS Project BGP Analysis

## CIDR at Work

CJ Wittbrodt  
Slides “Stolen” From: Cengiz Alaettinoglu

# BGP is not perfect!



Too many catastrophes

often due to misconfiguration and bugs

Known convergence problems

up to 30 minutes

Questionable scaling properties  
linear?, quadratic?, exponential?, *hyper-exponential?*

# Why care about growth?

Linear or quadratic growth

need to add more memory/cpu power once in a while

Exponential but slower than Moore's law

we can still cope with this, but gets expensive

Faster than Moore's law

next generation hardware costs more

need a new EGP, fast!

no time to understand all the issues...

# RIPE / RIS Project

Massive BGP data collection project

At 7 locations, oldest since September 1999

Every single BGP message, state change

in mrt format, using zebra software

our analysis uses mrtlib from Merit

We use rrc00 at RIPE NCC

peering 13 routers now

peered 22 routers over the course of its operation

We are extremely grateful to RIPE/RIS project  
for their efforts and making the data available.



# BGP Messages vs Snapshots

---

We want to understand BGP's operation, not just  
routing table growth  
routing table growth  
churn  
convergence  
volatility

Snapshots are not nearly as powerful even for  
routing table growth  
e.g. under-estimates multi-homing, AS topology

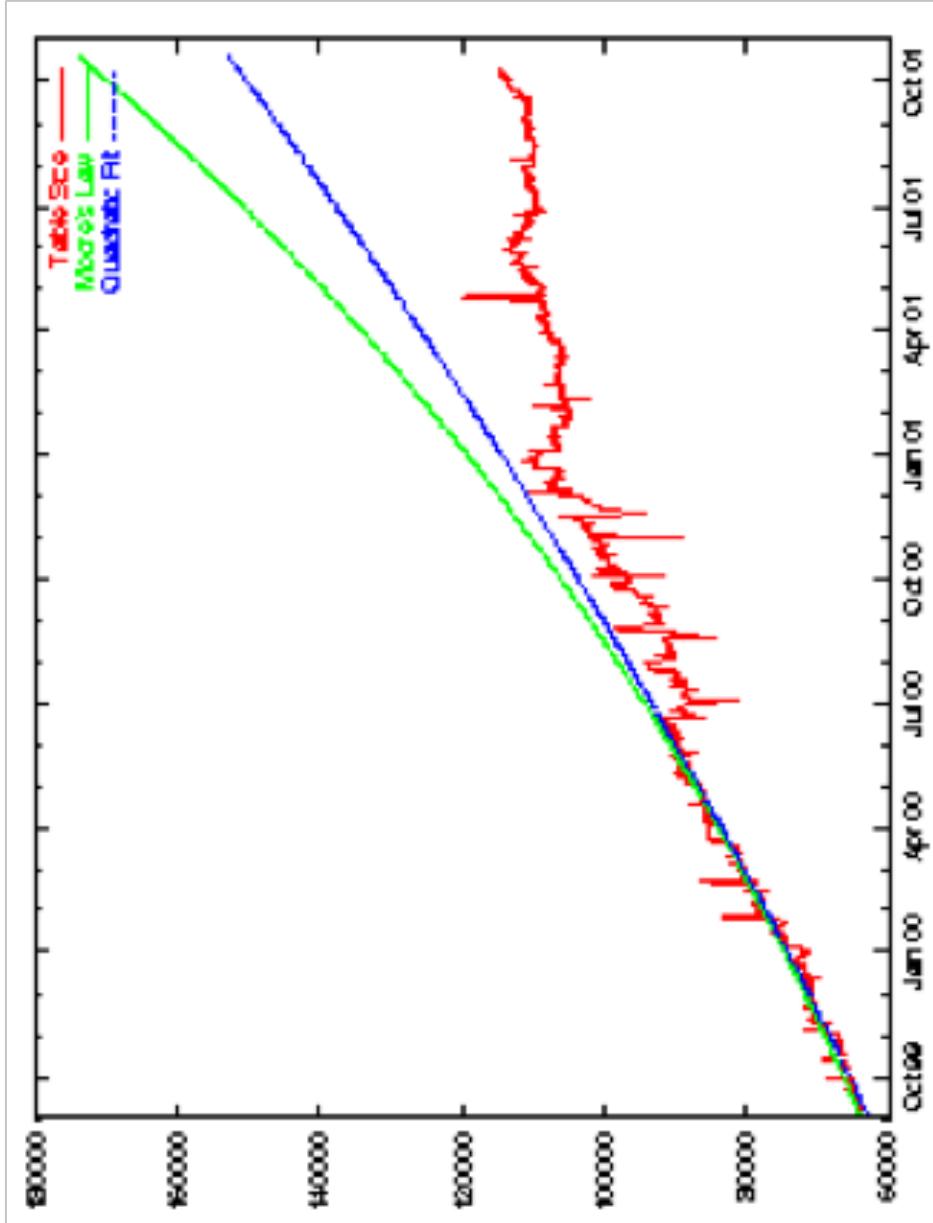
# Routing table size

---

Total number of prefixes over all active peers

not the average routing table size  
each routing table has prefixes the others don't

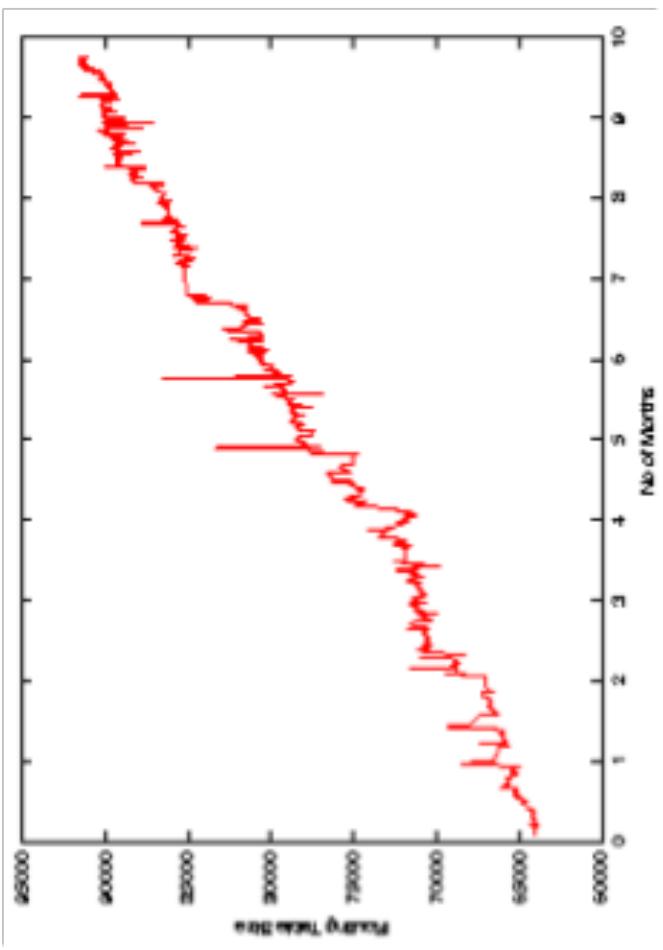
# Routing table growth



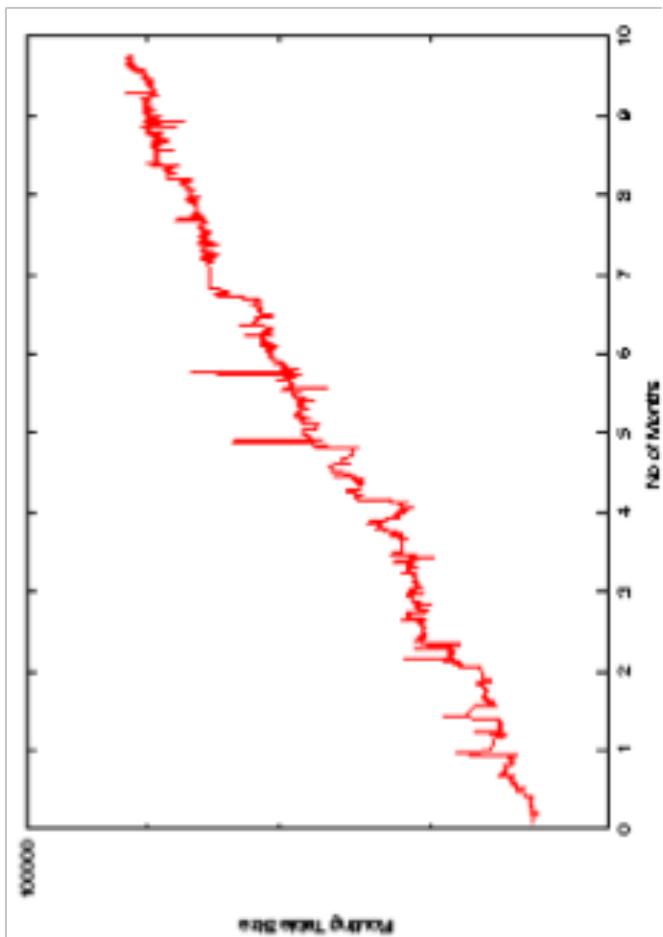
Smoothed w/ moving average of 5 data points



# First Ten Months



Linear y axis scale

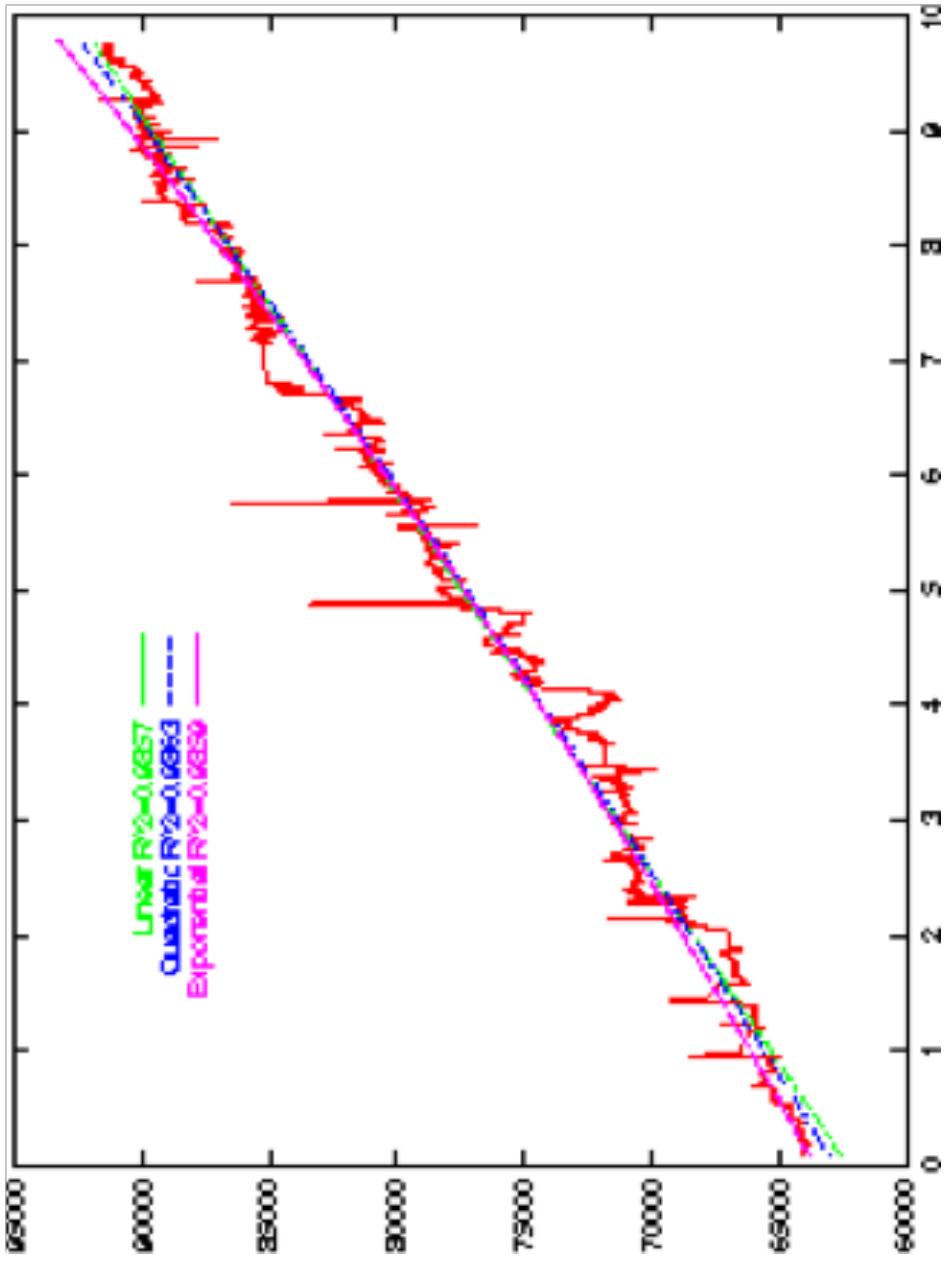


Log y axis scale

Polynomials and exponentials can look alike depending on coefficients, constants and the x axis range



# The Fit



Linear, quadratic, exponential all fit well

thru residual analysis methods

Packet  
Dacinn

# CIDR at work

How effective is CIDR in presence of  
multi-homing  
inter-domain traffic engineering

## Historic prefixes

Prefixes that were used to be advertised

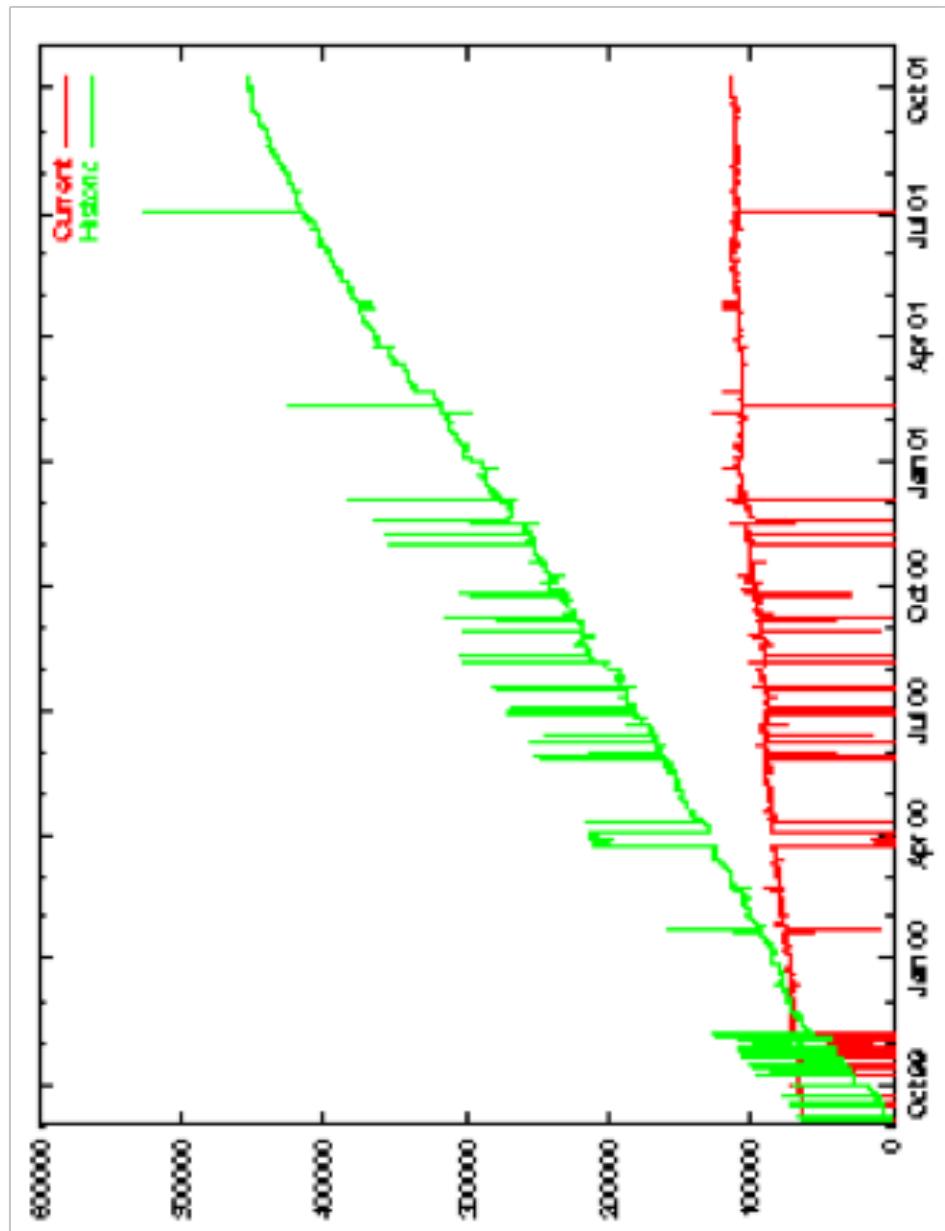
CIDR

ISPs leaking their more specifics => CIDR

Private address space (fixed number)

Un-assigned address space (spammers)

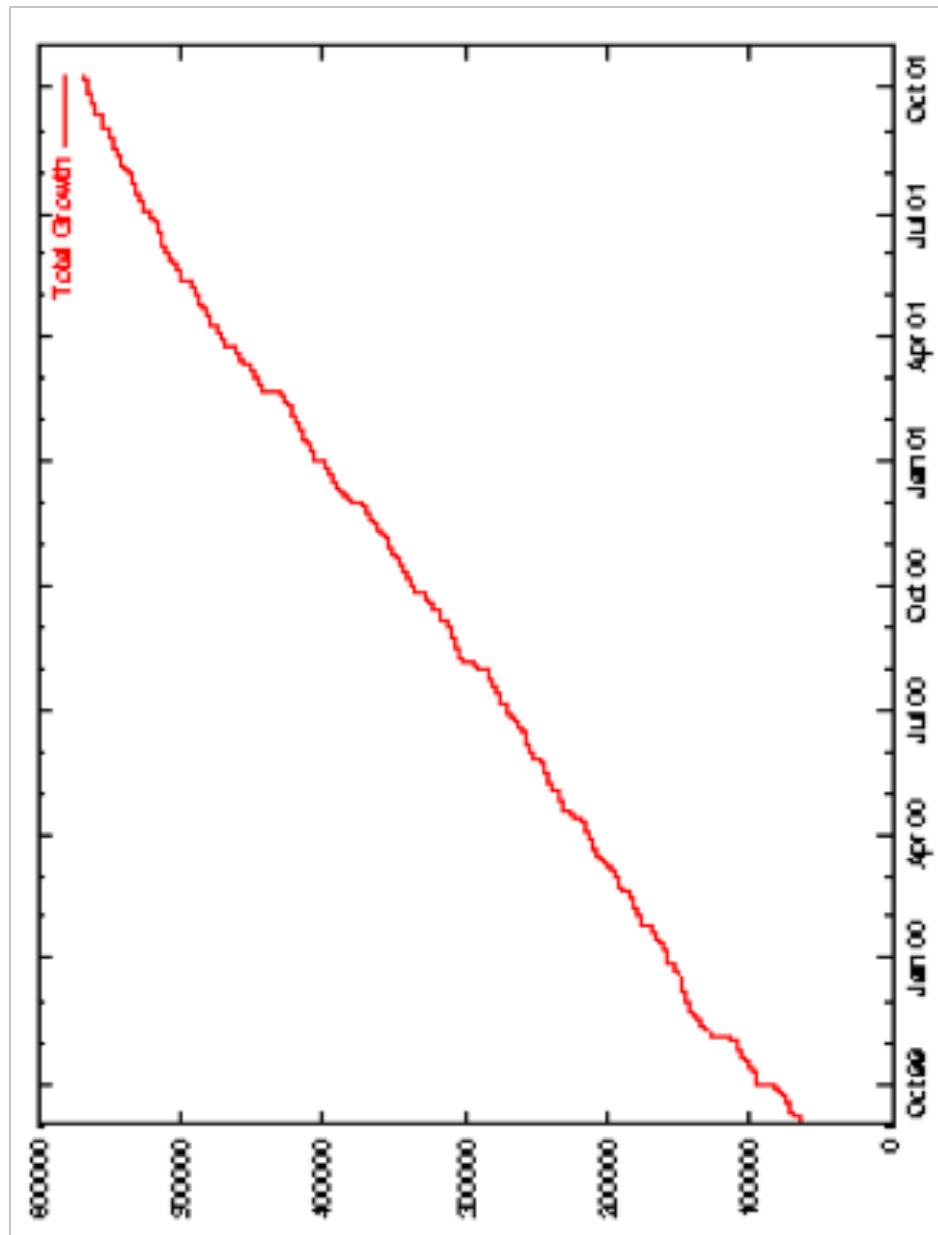
# CIDR is working very well



Without CIDR routing table would have been ~5 times larger



# The total growth



Exponential does not fit



# Where is the growth coming from?

---

Classify a prefix into **one** of the following groups:

Multi Homing

with origin AS doing BGP w/ multiple ISPs

with multiple origin ASes

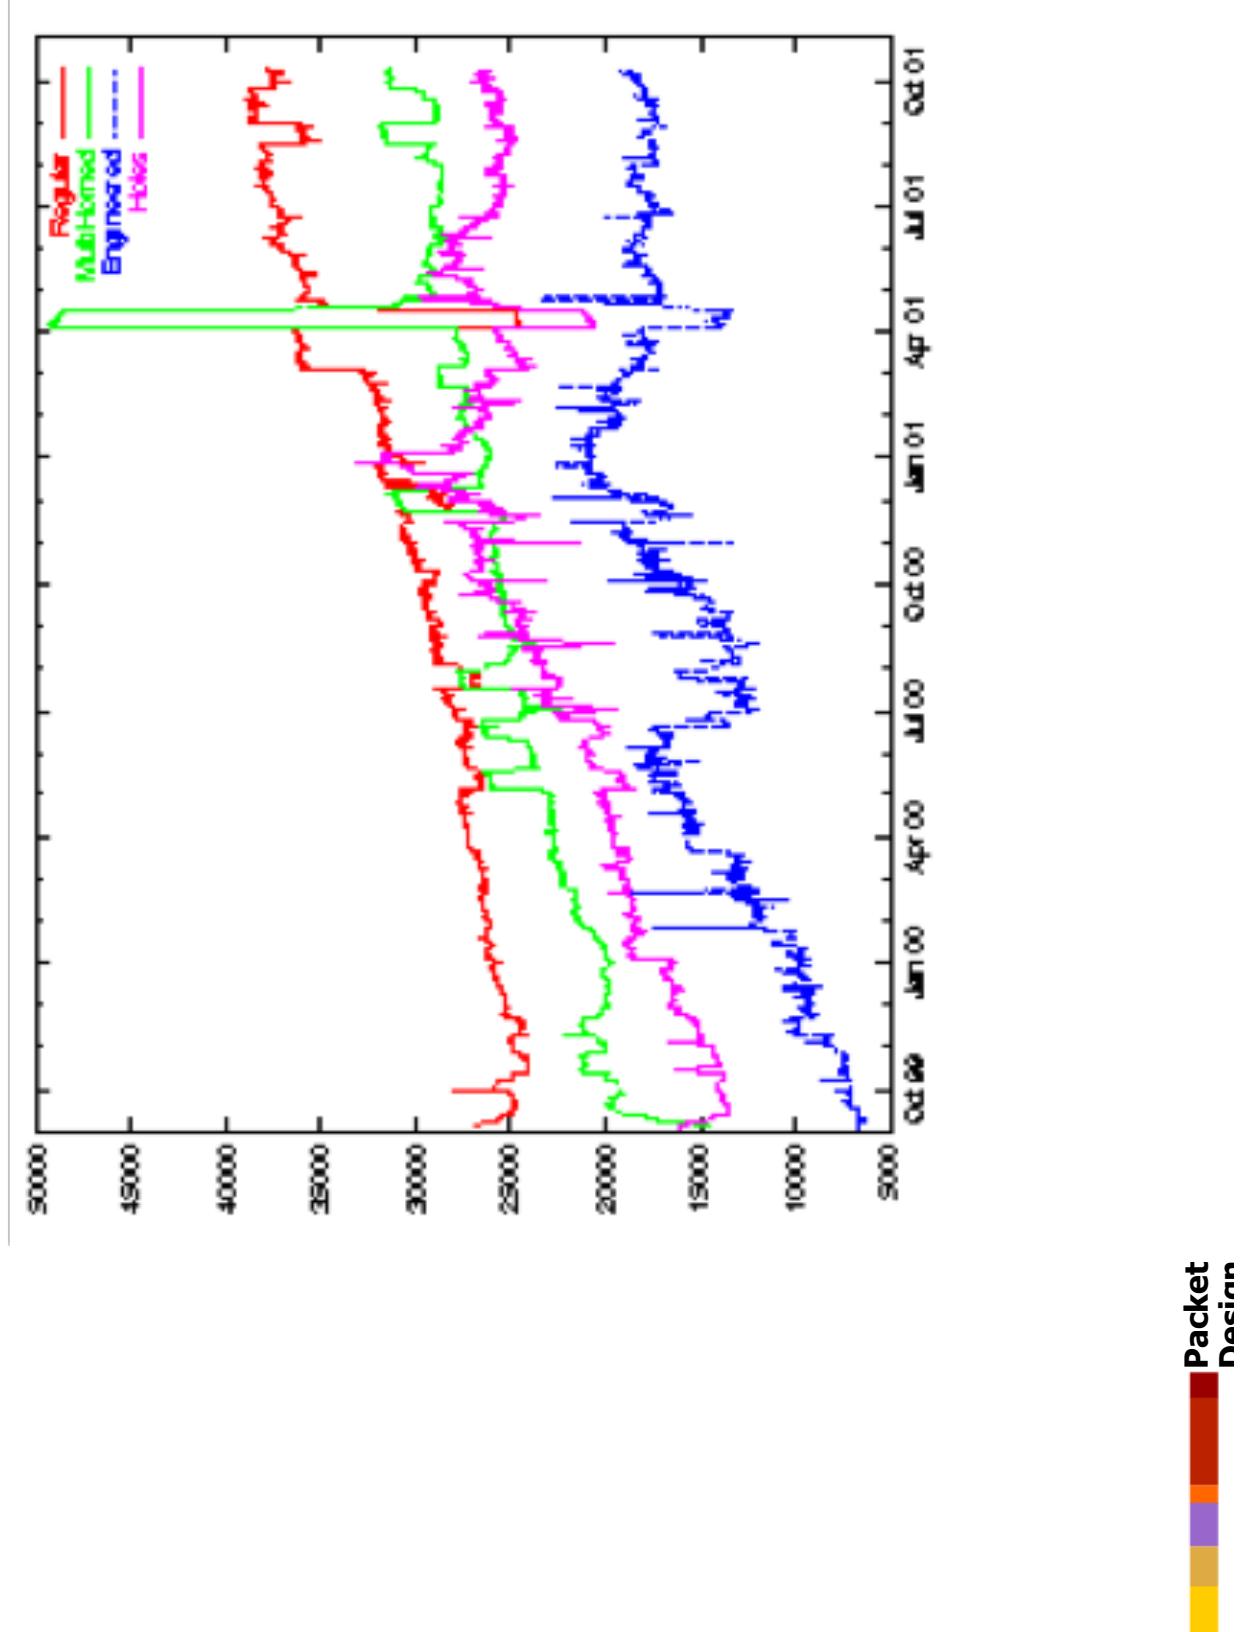
Engineered prefixes

prefixes and their components w/ the same origin AS

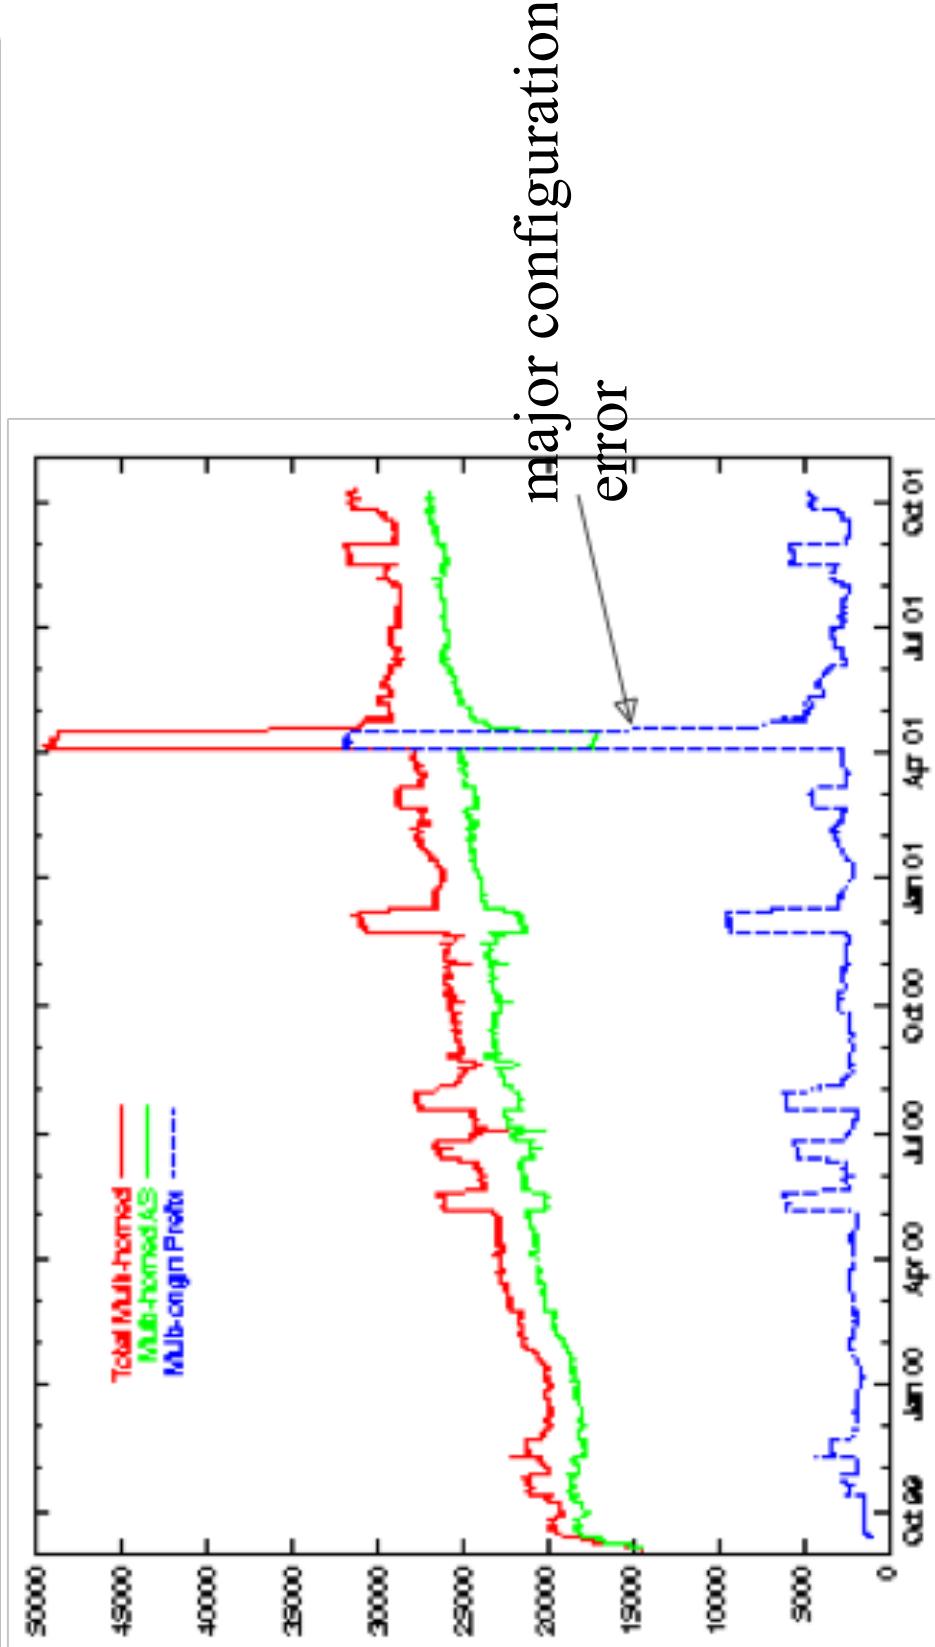
Punching holes

prefixes and their components w/ no common origin AS  
perhaps some are multi-homed

# Components of Growth



# How do People Multi-Home?

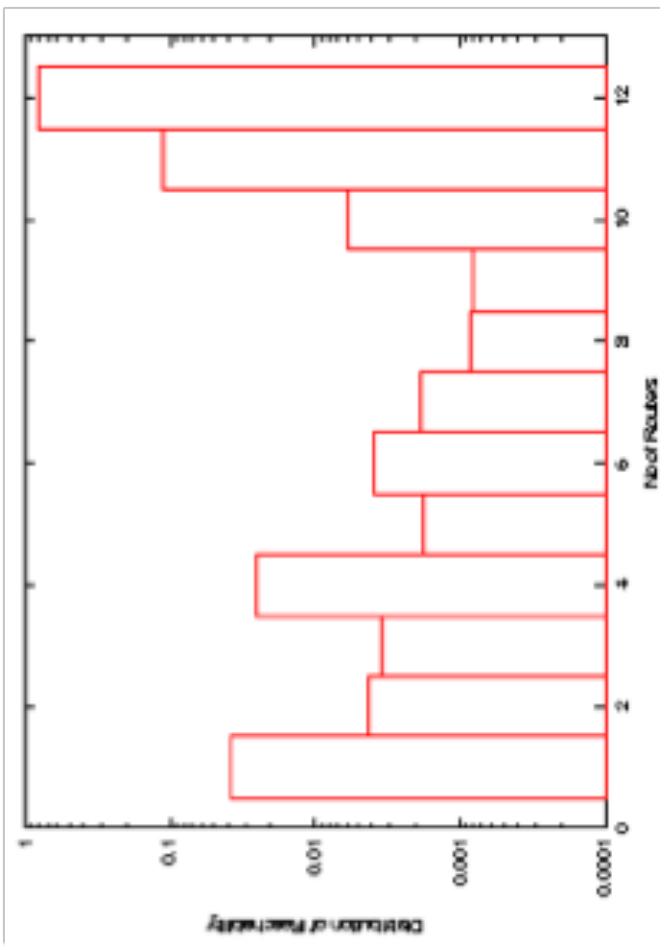
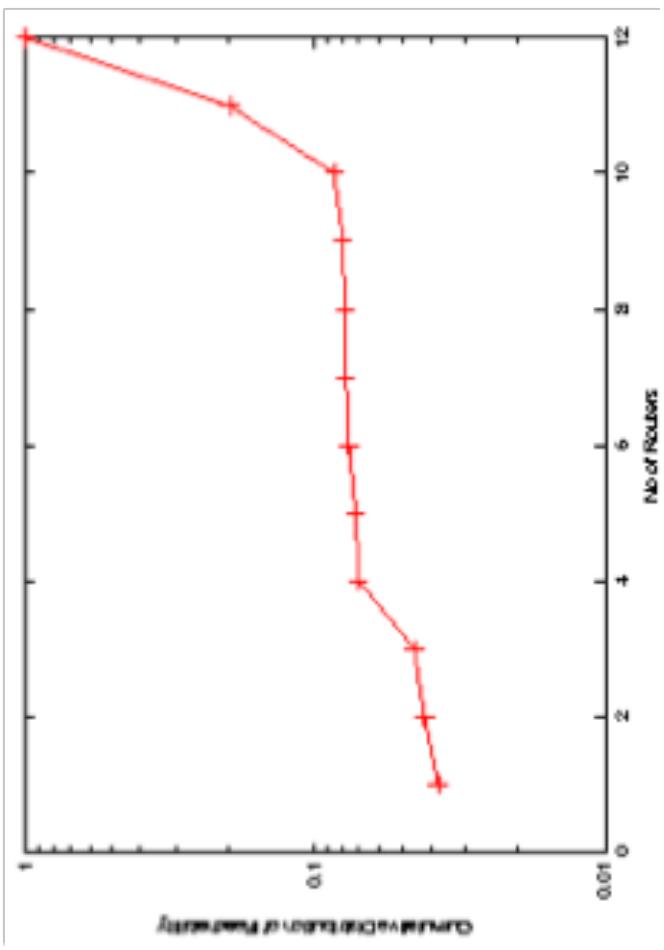


Dominated by running BGP at the site

Some do it through their ISP's BGP



# Prefix to Router Distribution



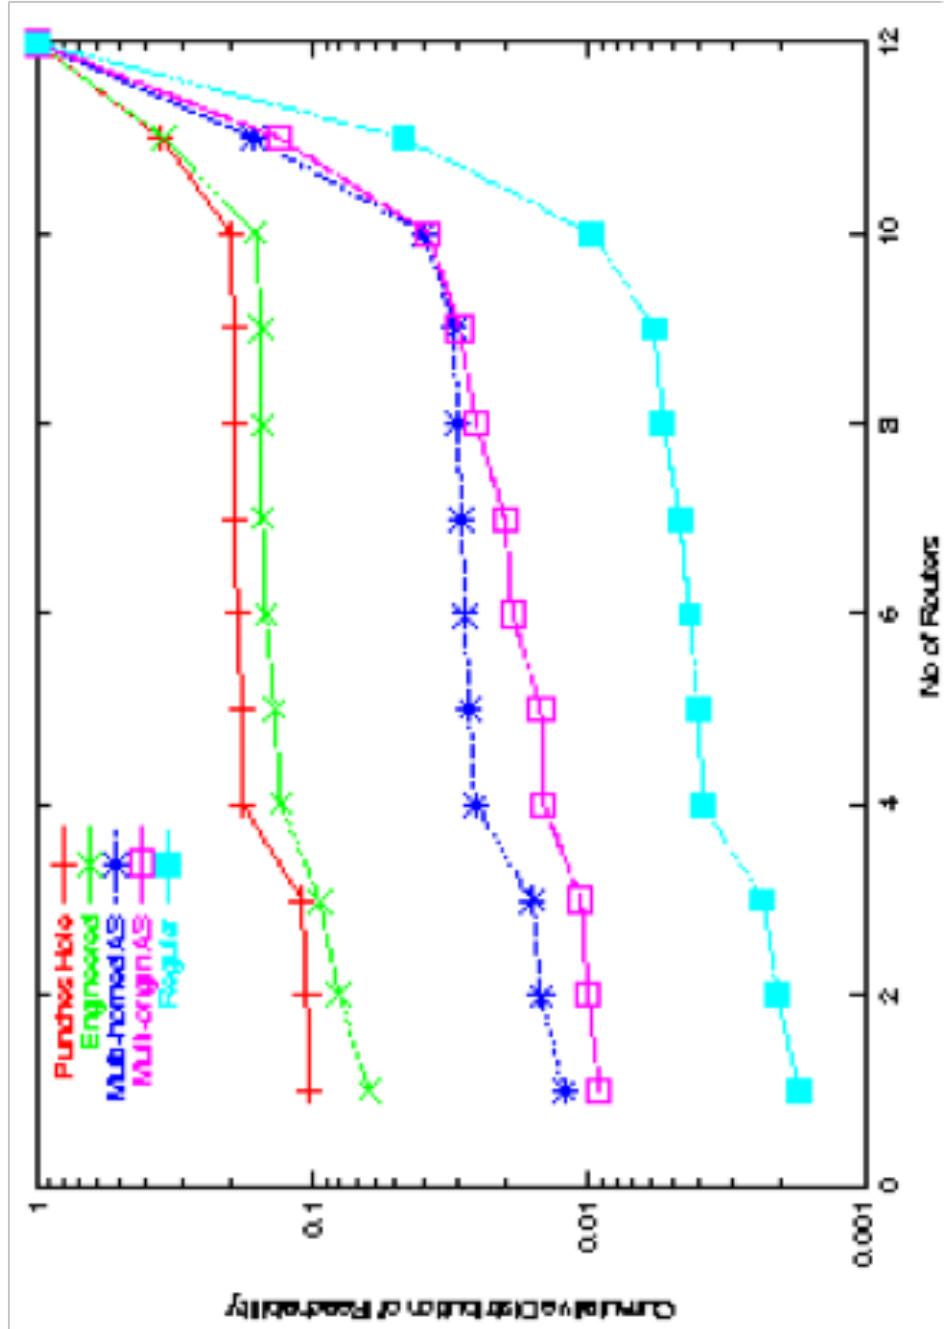
**Only 80% are seen by all 12 routers**

If you take 2 ISPs, the number is much higher

**Caveat: using a snapshot only**



# Filtering and Routing Distribution

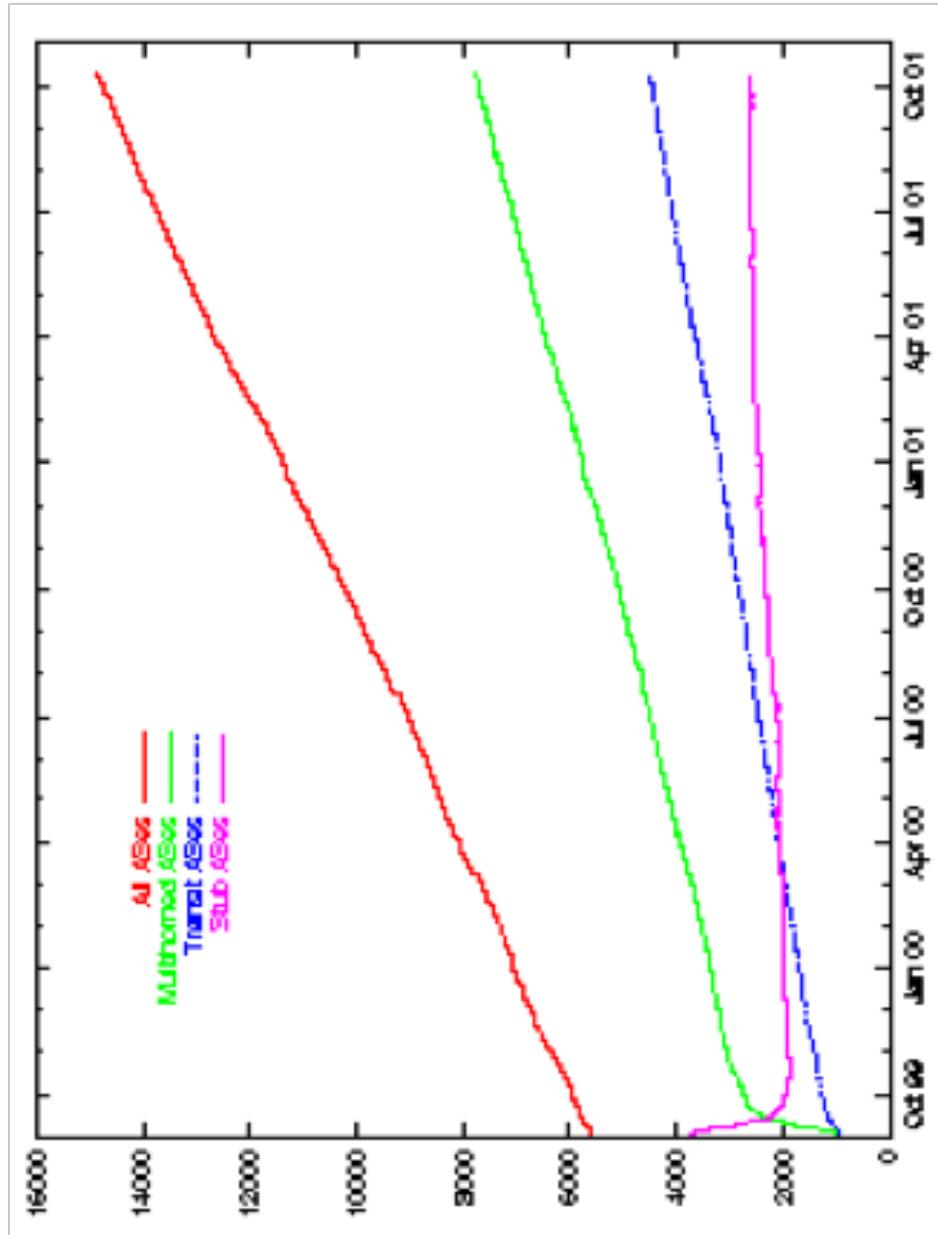


# Word on AS number growth

---

- We were told: exponentially growing
- If each of these ASes had 1 prefix, we would have seen exponential multi-homing growth...

# AS Growth



quadratic fits best:  $R^2=0.9996$  (also by Tim Griffin)  
exponential fits very badly (residuals are not stationary)



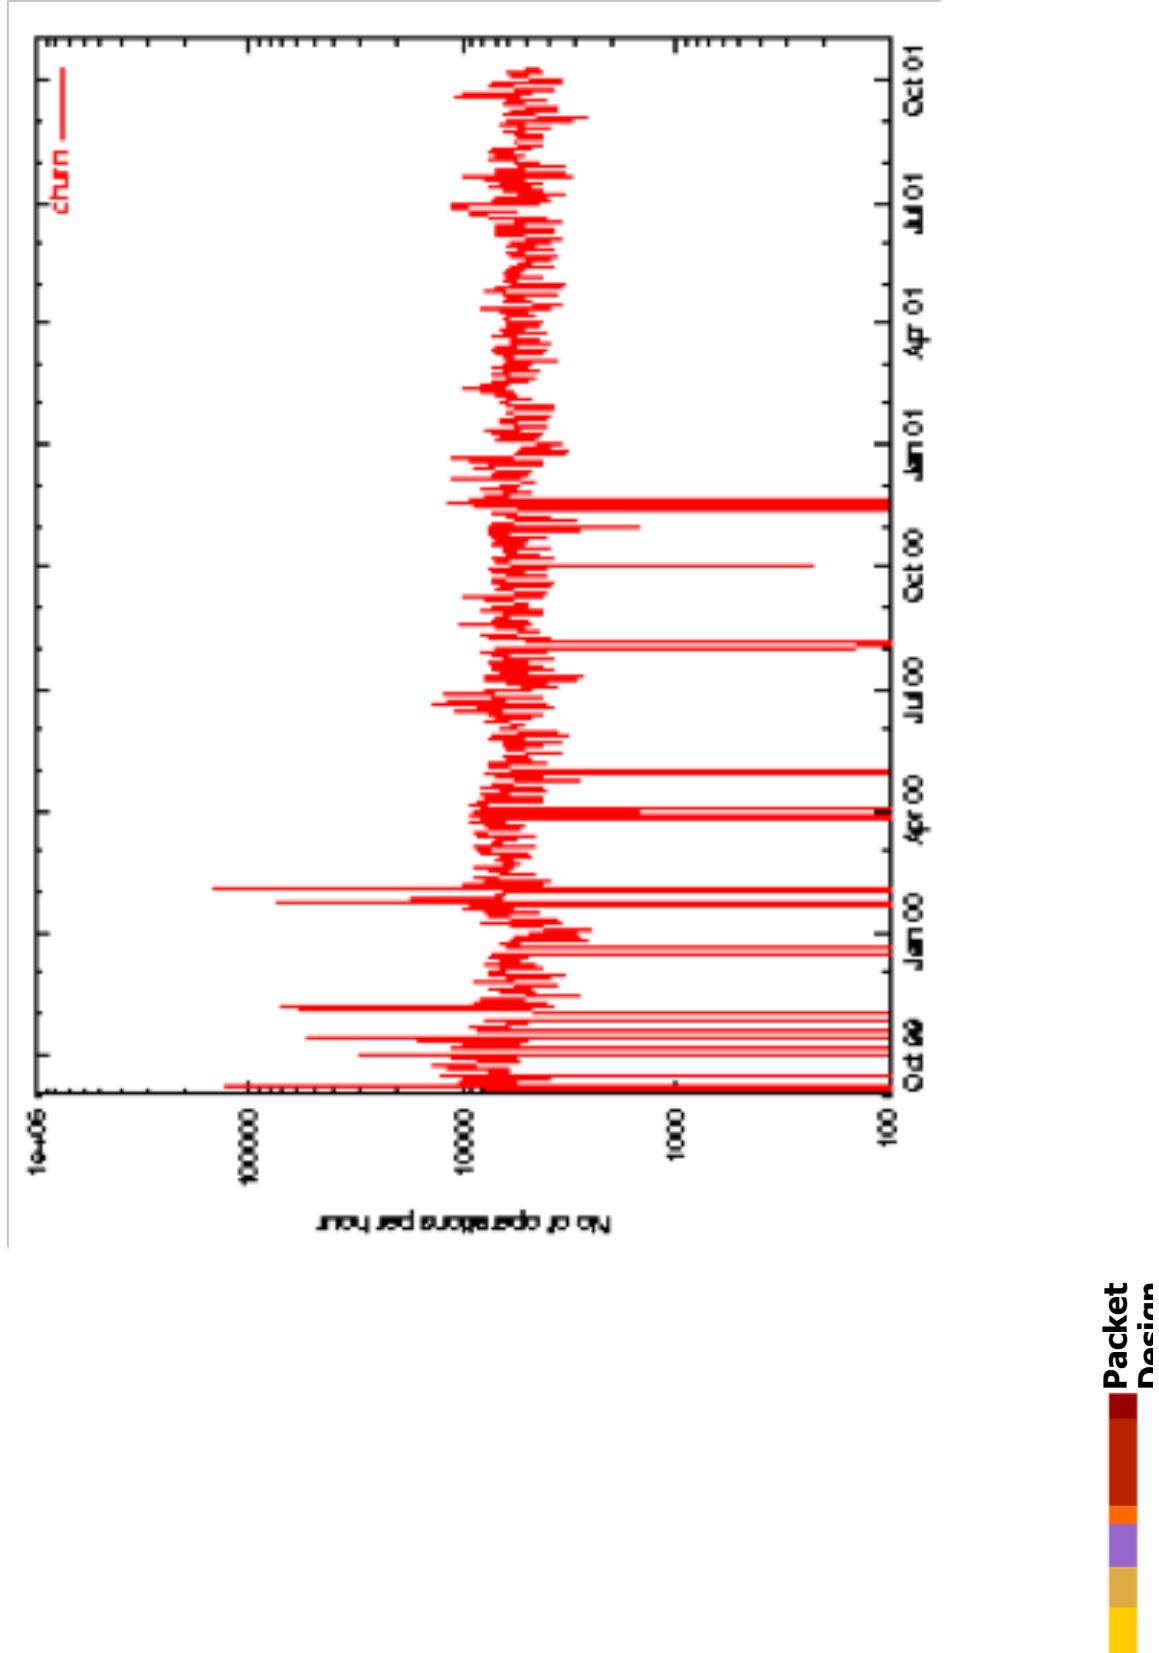
# The churn

---

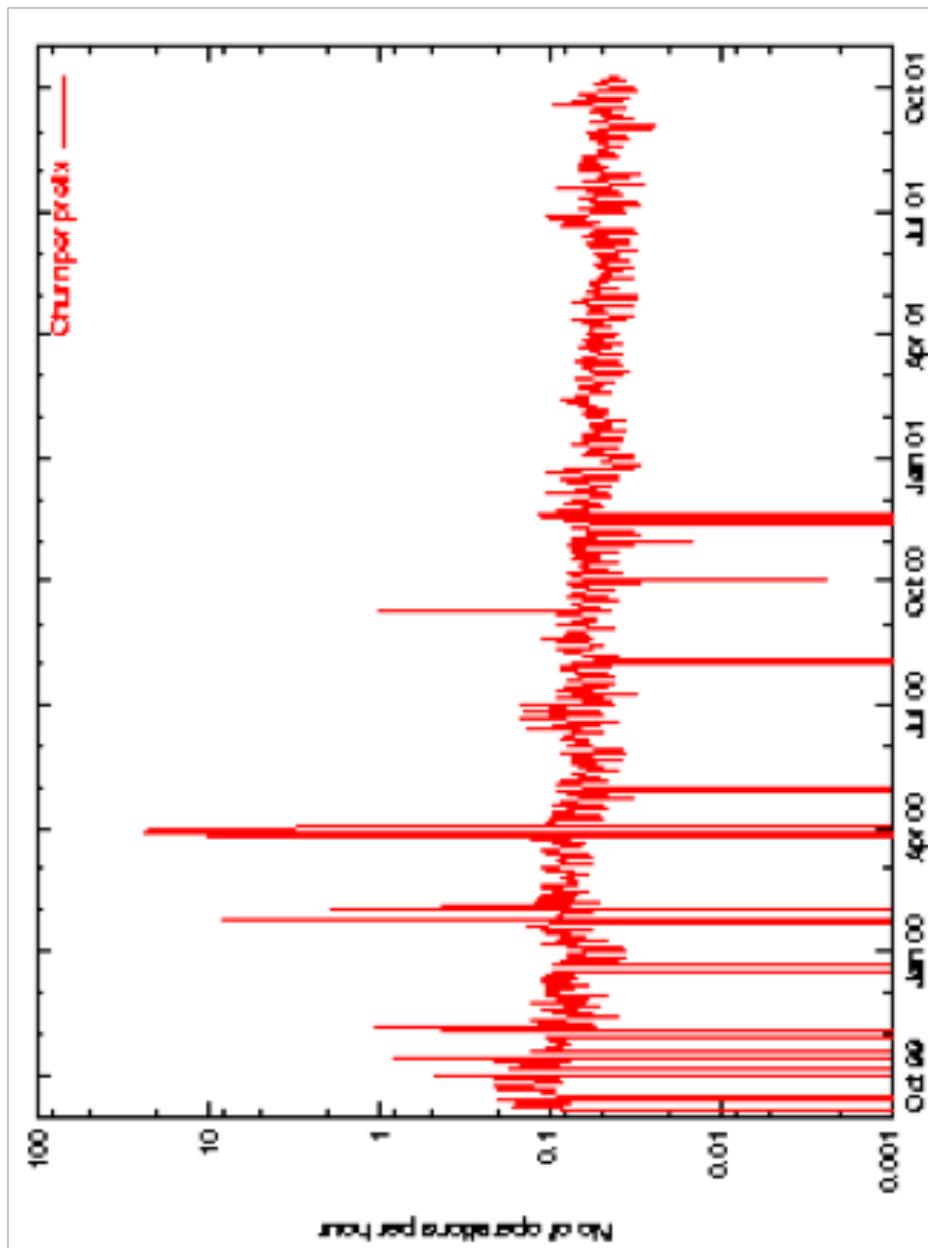
During BGP convergence factorial number of alternate paths are examined  
Churn may be growing faster than the space

We looked at churn by each router  
took median/min/max/ave hour to represent the day  
took the median router  
for each class of prefixes

Overall churn is decreasing

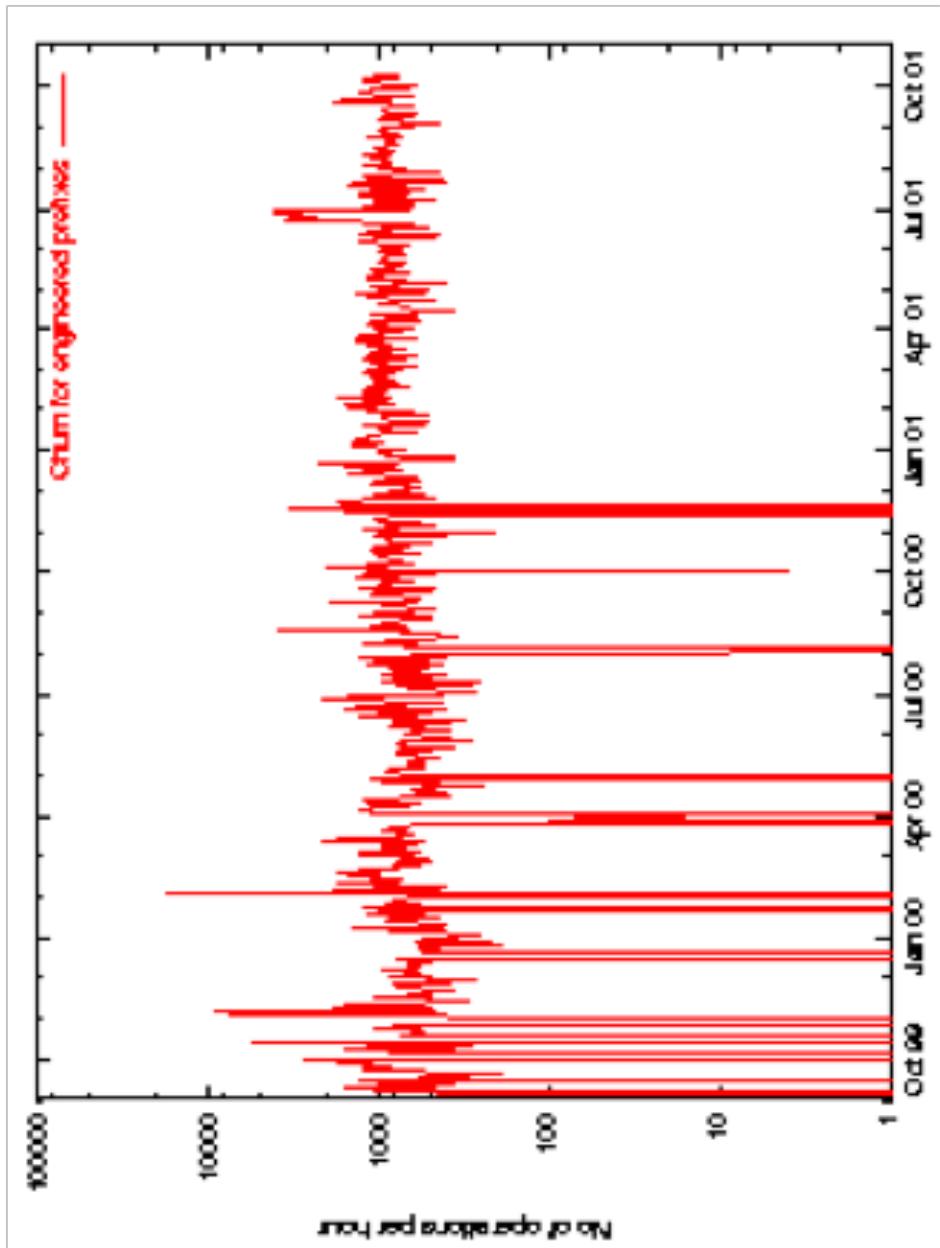


# Per prefix churn: stability



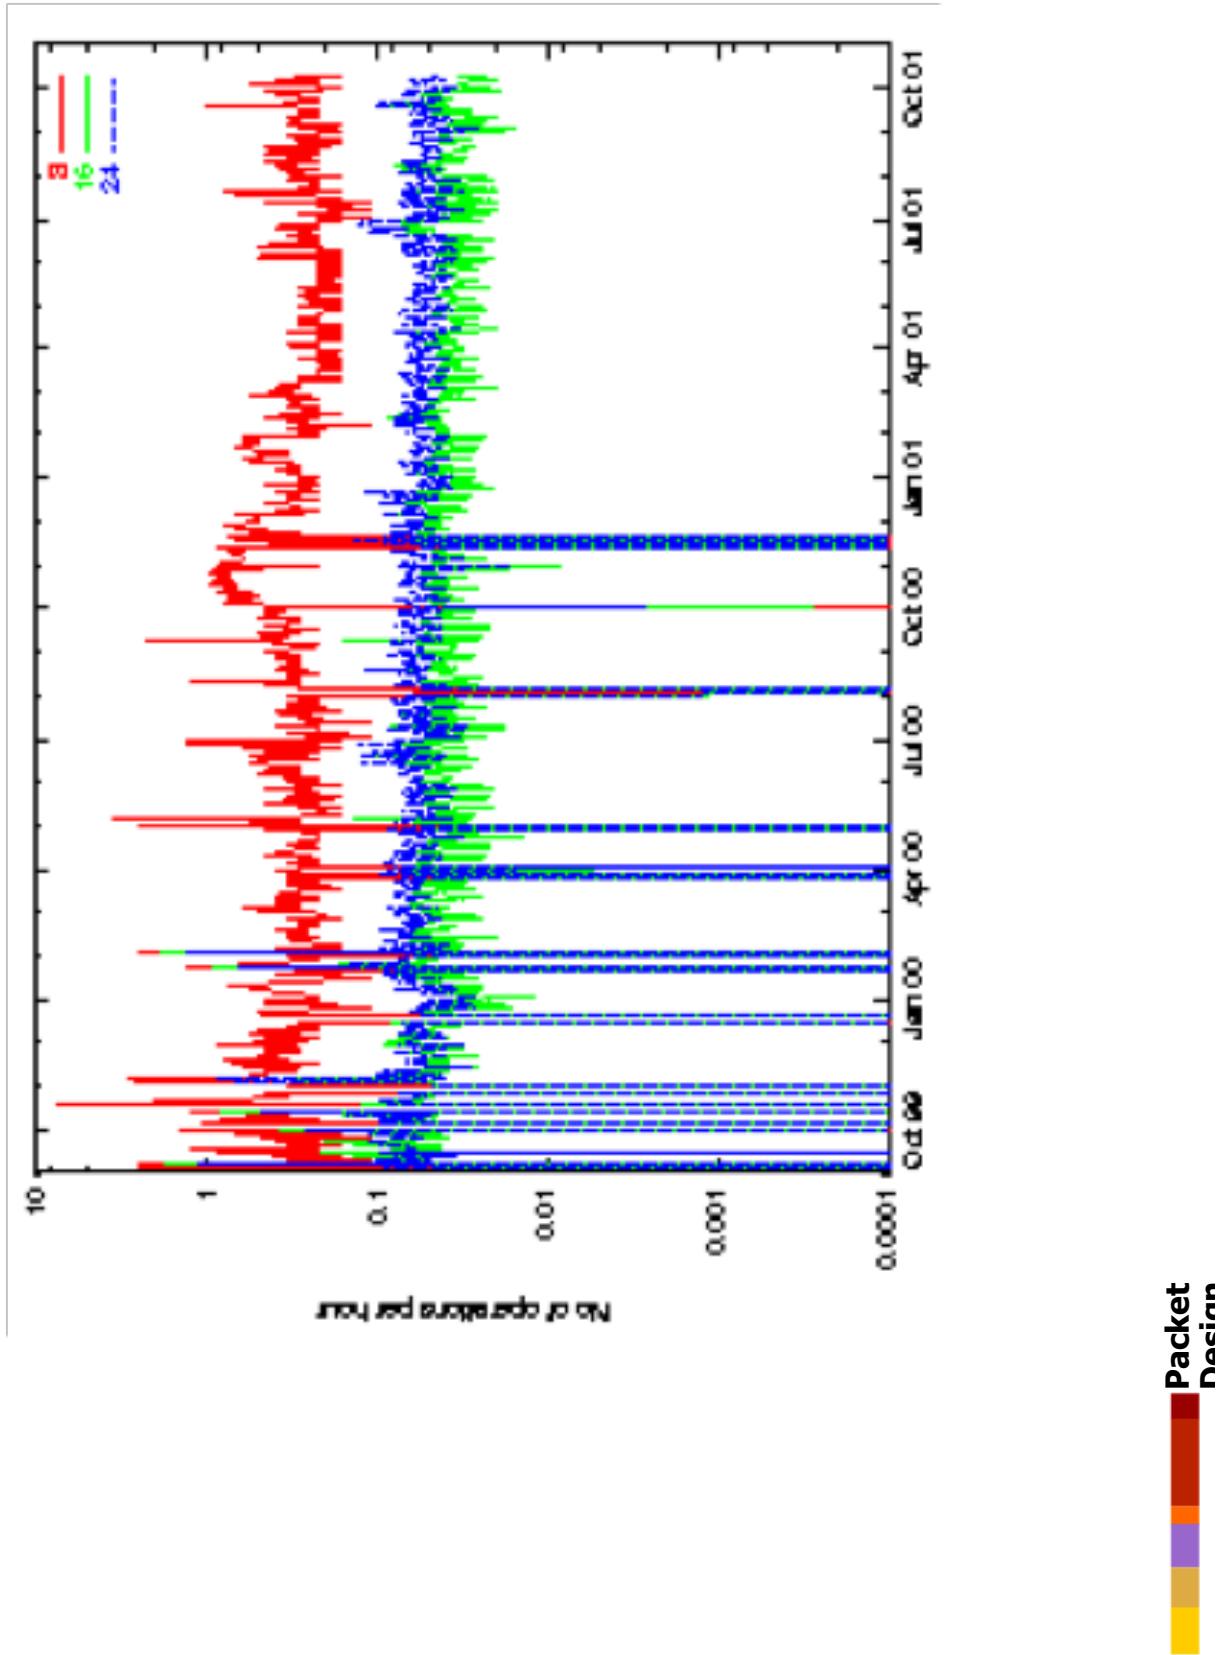
Packet  
Dacian

# Engineered Prefixes behave differently



Do people engineer prefixes dynamically?

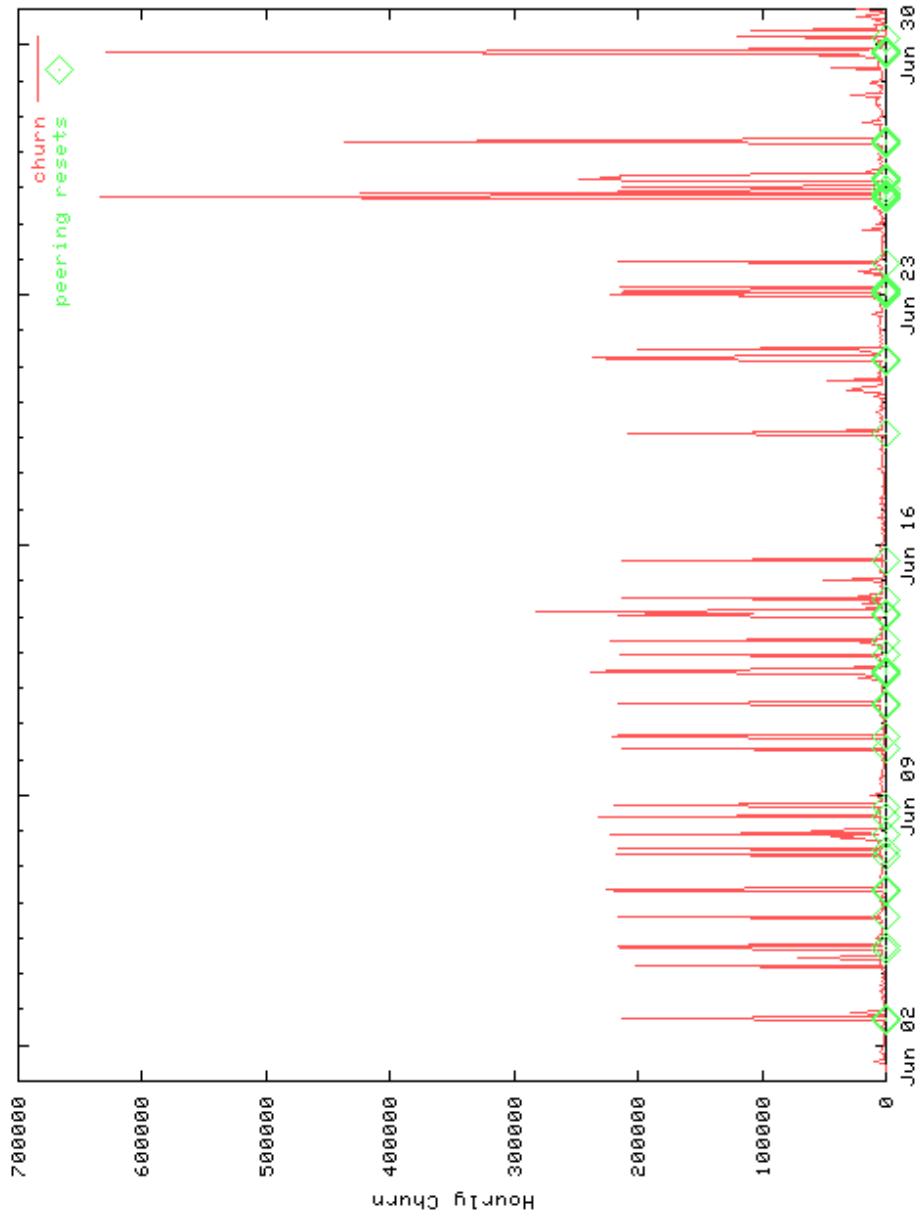
# Churn by Prefix Length



# Major cause of churn

- Each time a BGP peering is lost, up to 100K routes are withdrawn and then later re-announced
- For one router, out of 16M churn, at least 12M is because of losing and reestablishing peering

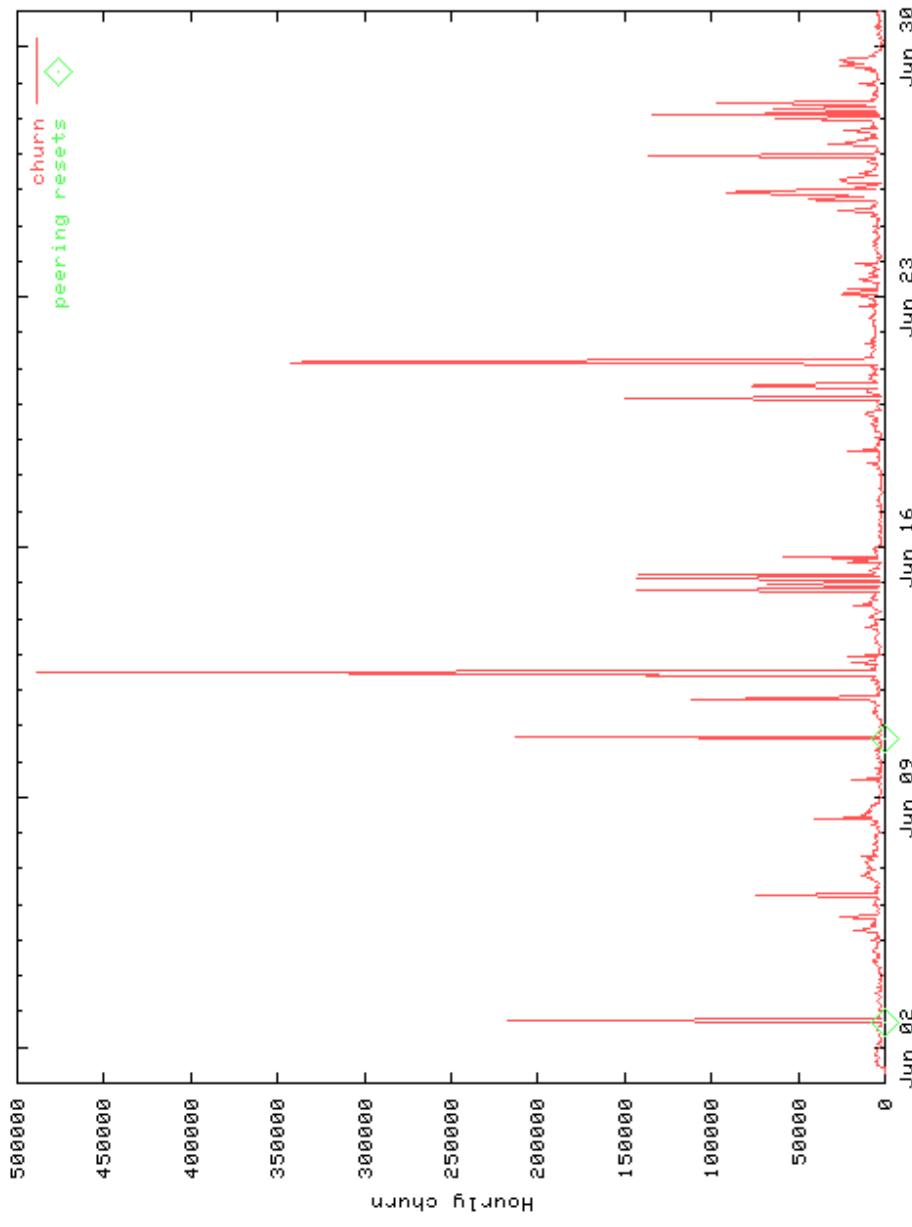
# Cause of Churn



Packet  
Dacian

Peering loss/reestablishment is very costly

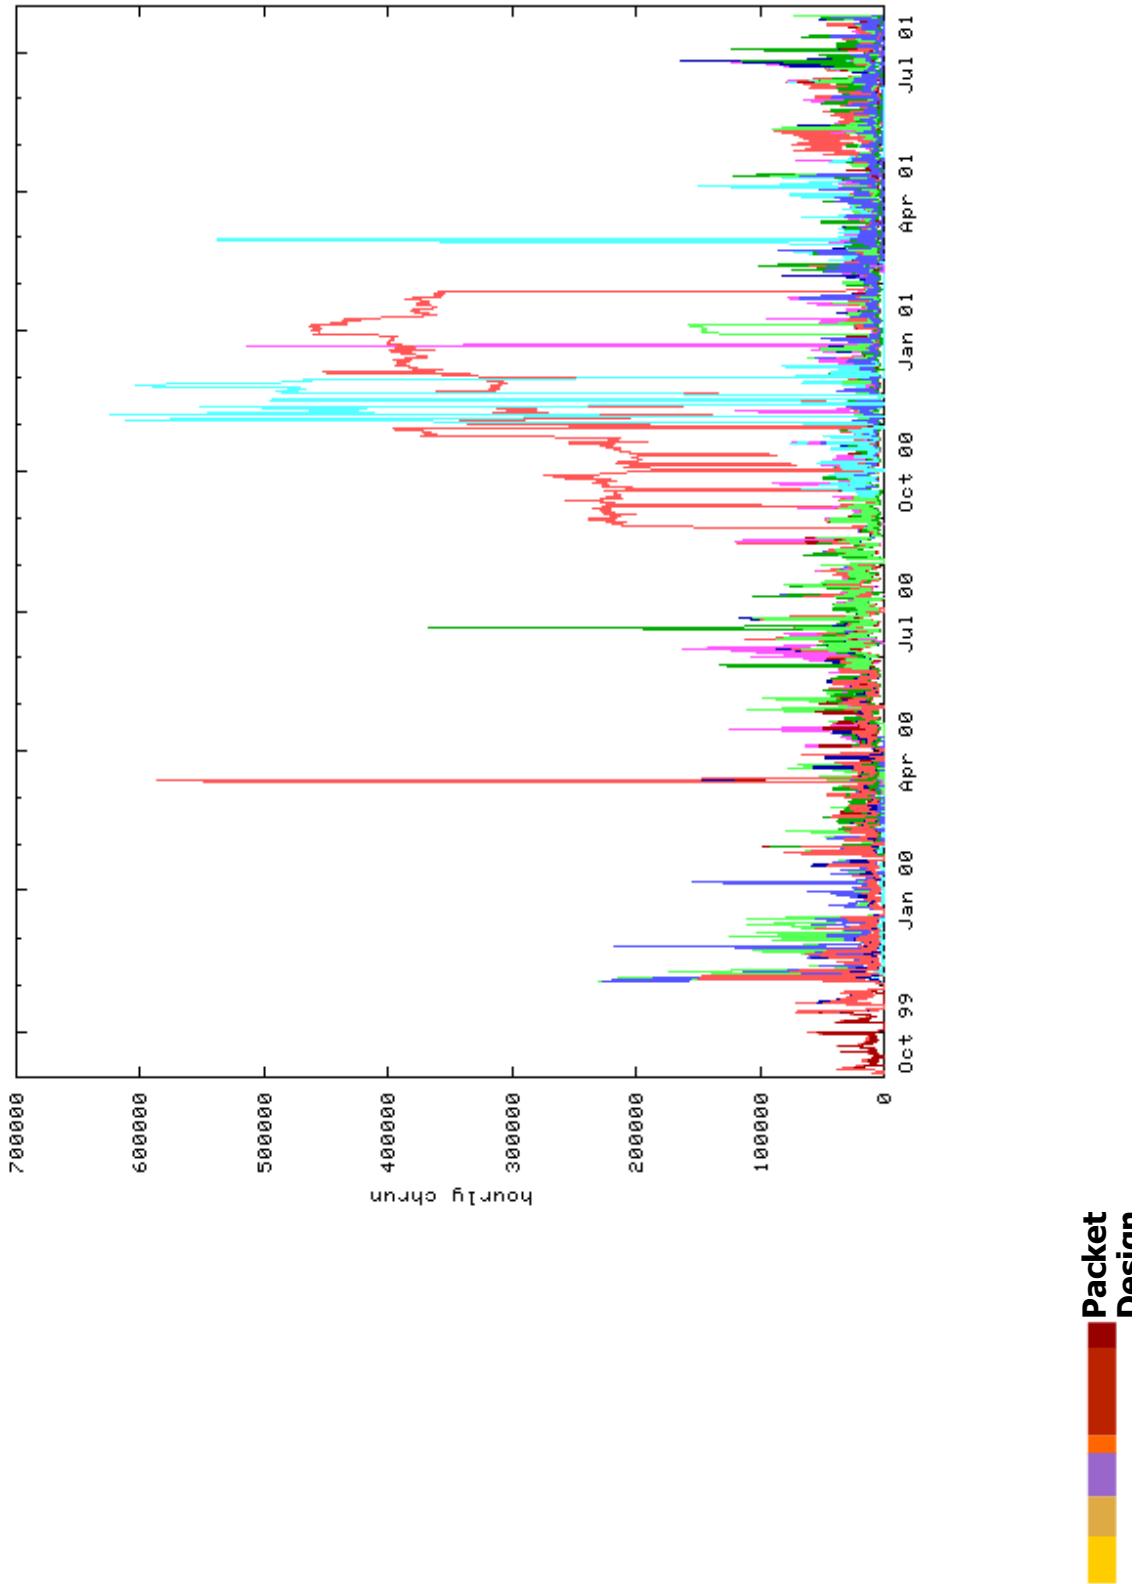
# This is carried multiple ASes away



Packet  
Dacian

# High churn due to peering reestablishment is carried multiple hops away

# High variance across routers



# Summary

## Growth

more evidence that it is quadratic than exponential

multi-homing growth: currently no evidence of exponential

Engineered prefixes: currently growing fastest

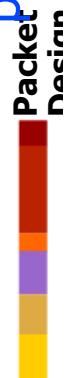
## Churn

overall there is a decrease

peering loss/re-establishment needs to improve

engineered prefixes churn more

peering loss/re-establishment dominates this by 2 orders of magnitude



# Conclusions

## Short term issues

increase router stability to remove spikes in churn  
configuration: fewer knobs, more robust filters

## Longer term

Trends can change (depends on economics/technology)

Growth is good but we need to handle it right