

RIPE/RIS Project BGP Analysis

CIDR at Work

Cengiz Alaettinoglu
cengiz@packetdesign.com
Packet Design

Revised for NANOG October 2001
IETF August 2001

BGP is not perfect!

- Too many catastrophies
 - often due to misconfiguration and bugs
- Known convergence problems
 - up to 30 minutes
- Questionable scaling properties
 - linear?, quadratic?, exponential?, *hyper-exponential?*
- *Understanding BGP's operation is crucial to getting the solution right*

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.

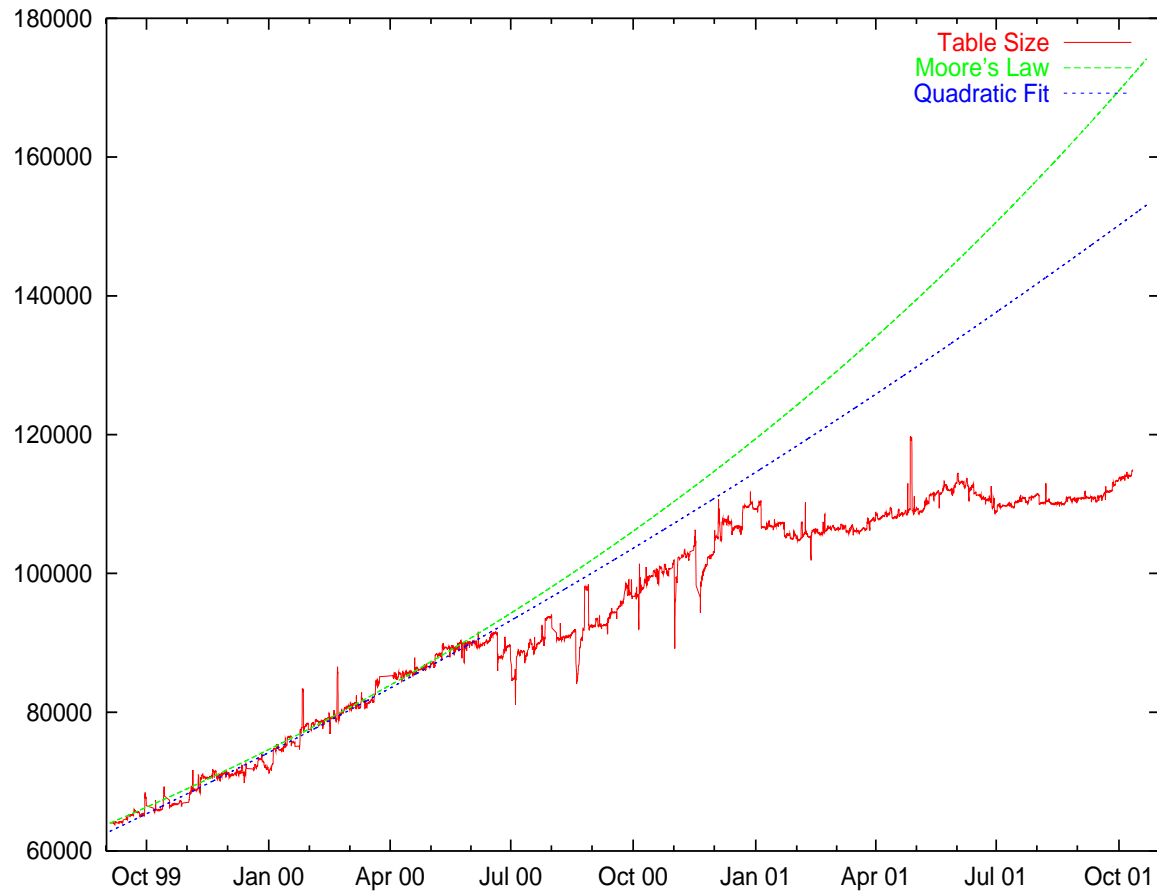
BGP Messages vs Snapshots

- 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 dont
- Growth
 - Very slow growth this year
 - Can the growth rates of the past resurface?

Routing table growth




- Smoothed w/ moving average of 5 data points

First Ten Months



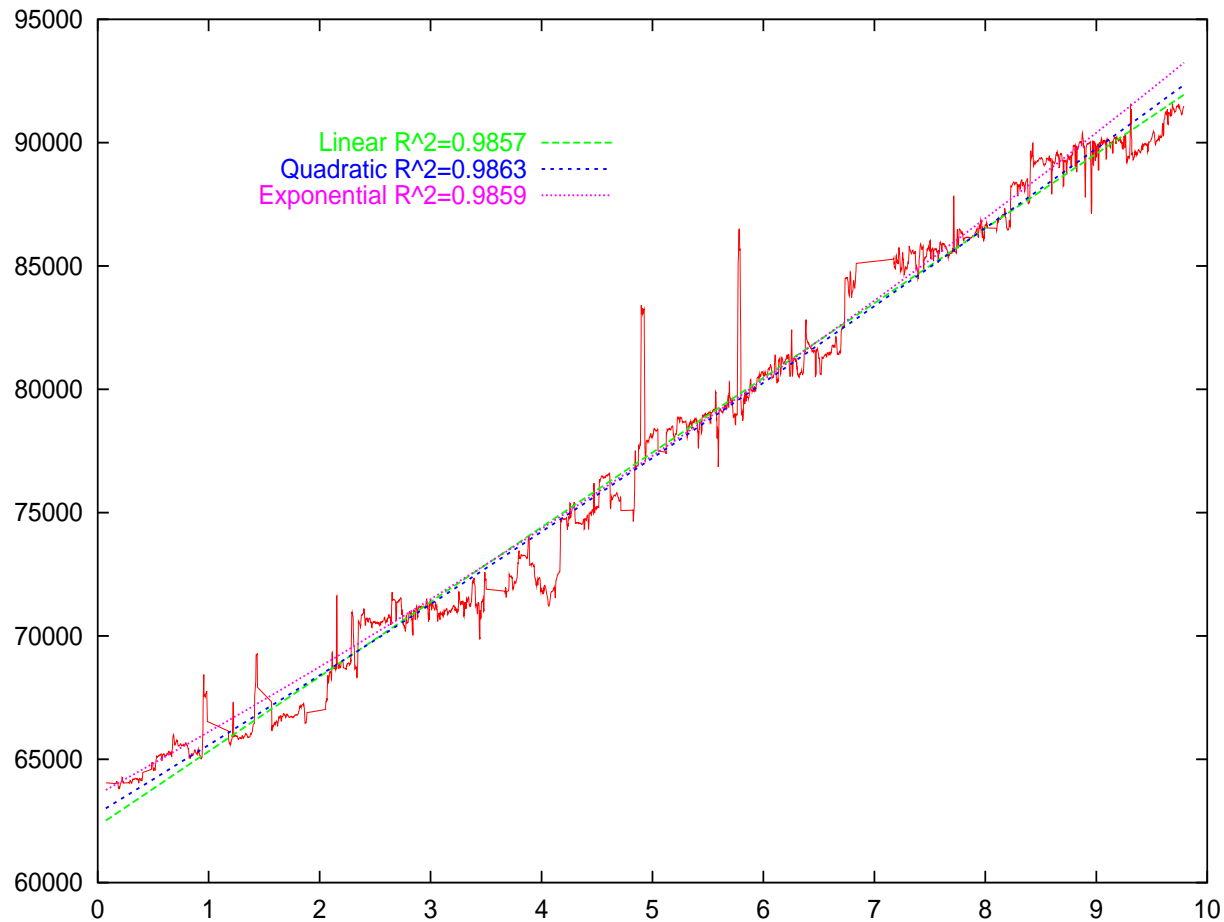
Linear y axis scale



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

Under differencing

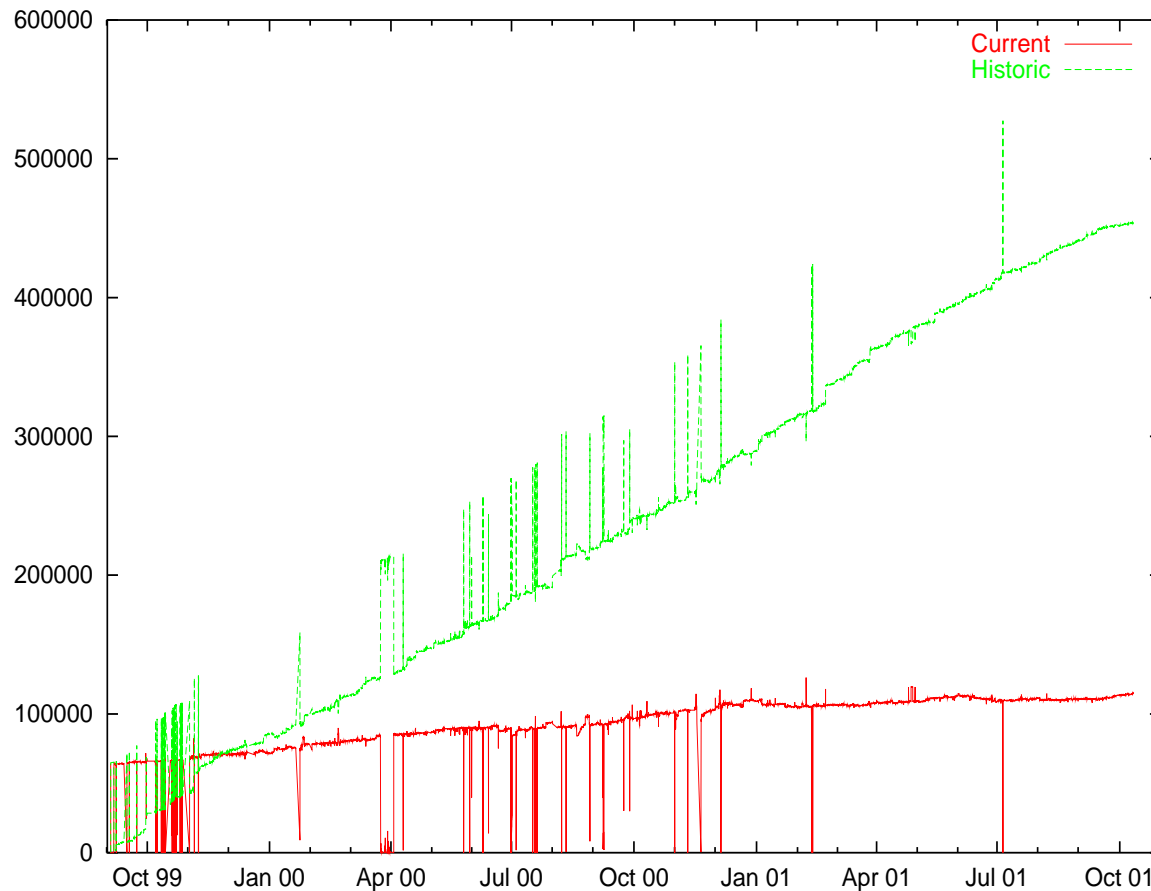


- Local trend analysis shows no evidence of an exponential

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

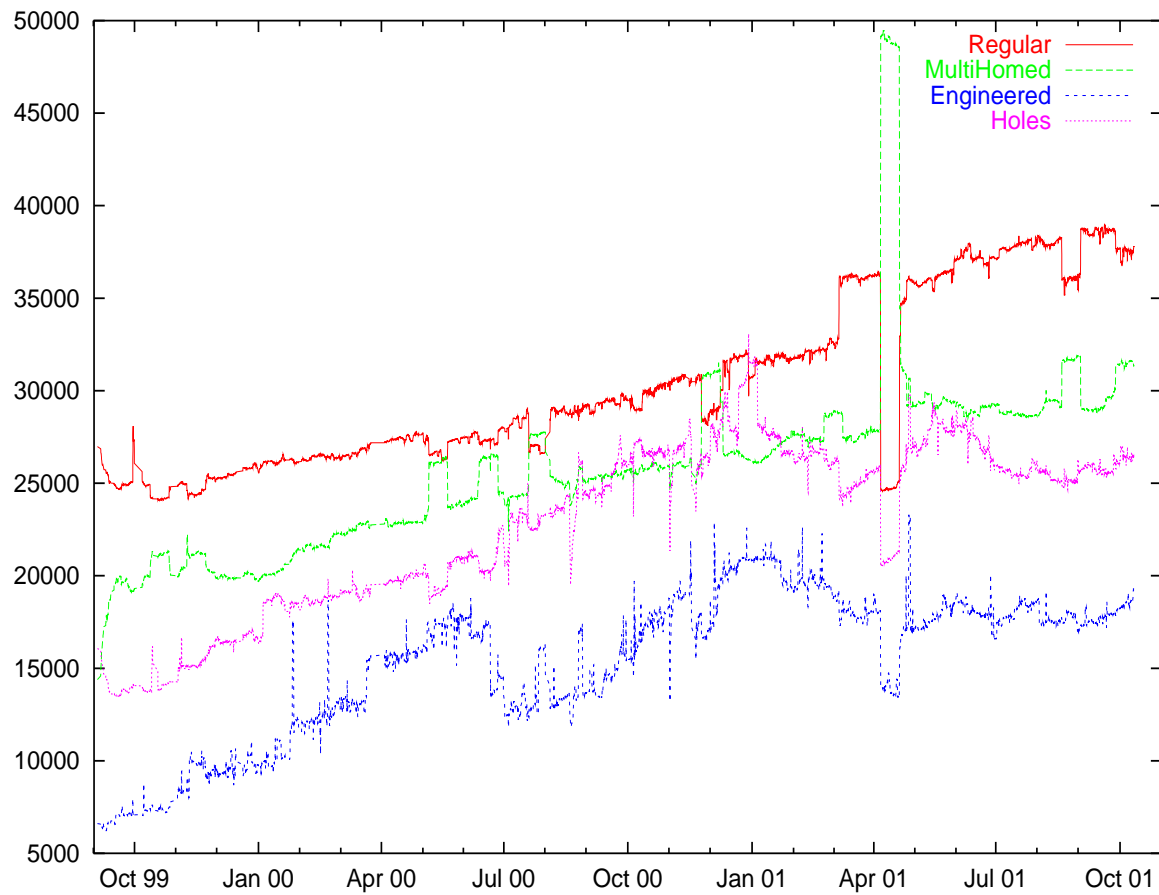


- Exponential does not fit

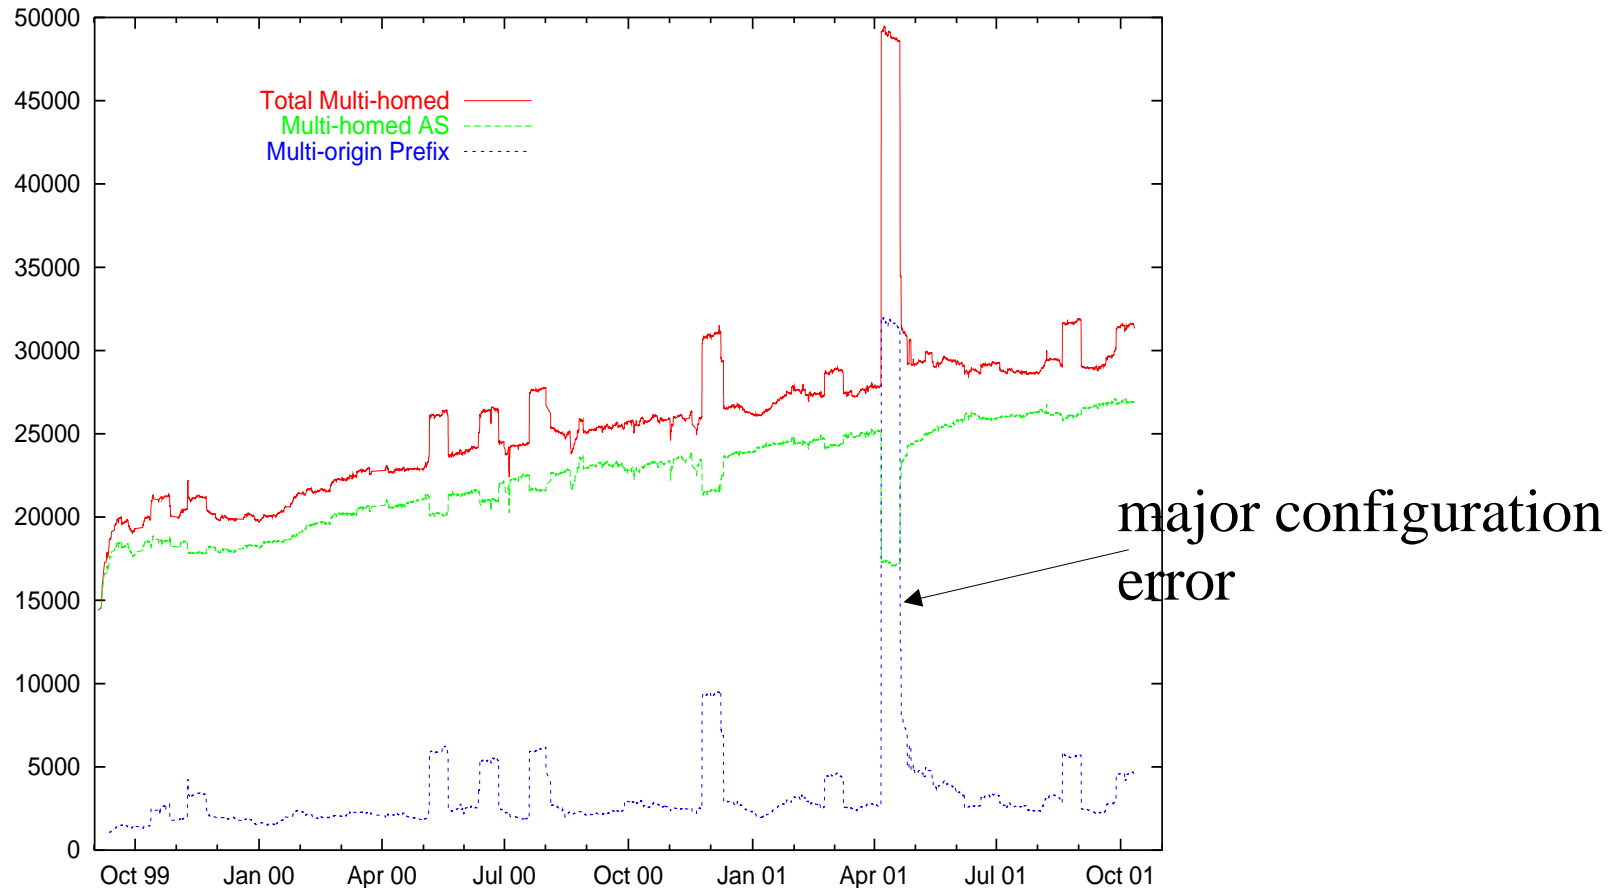
Where is the growth coming from?

- 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

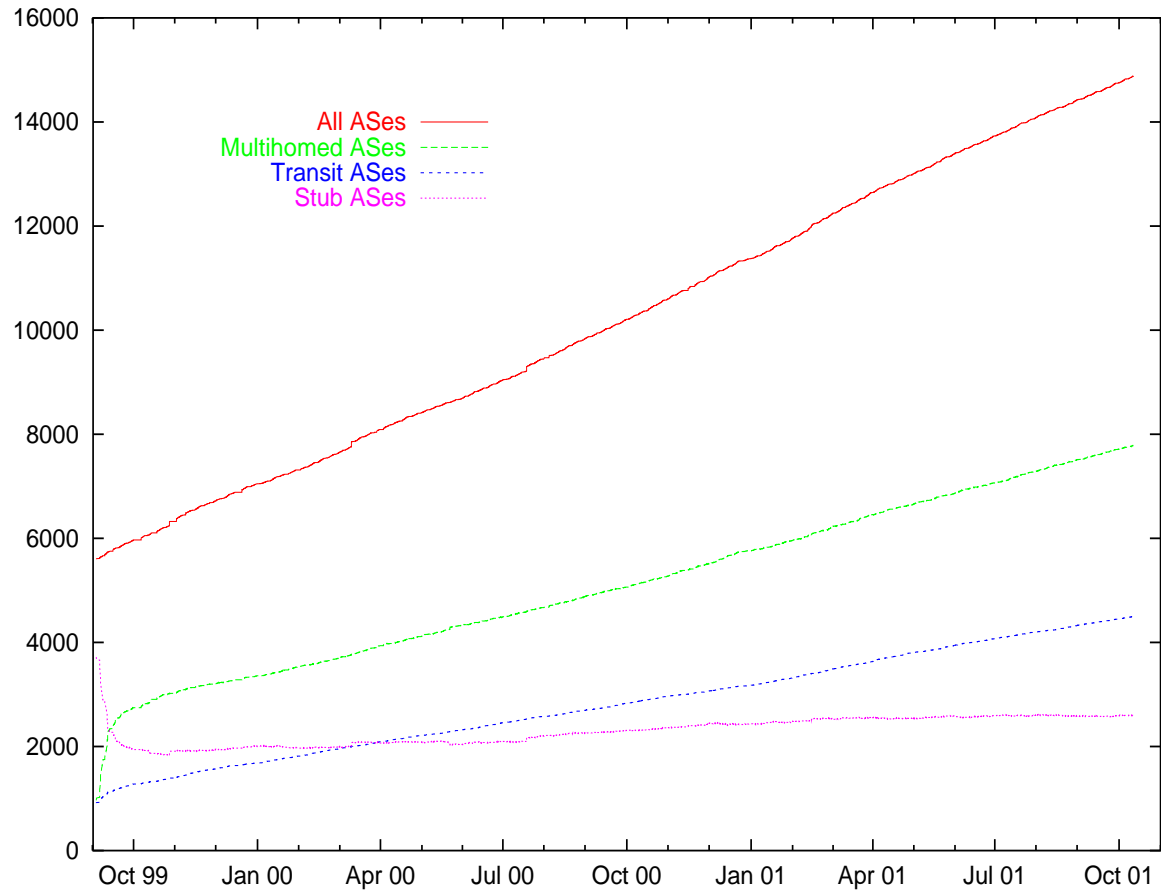
Filtering and Routing Distribution



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

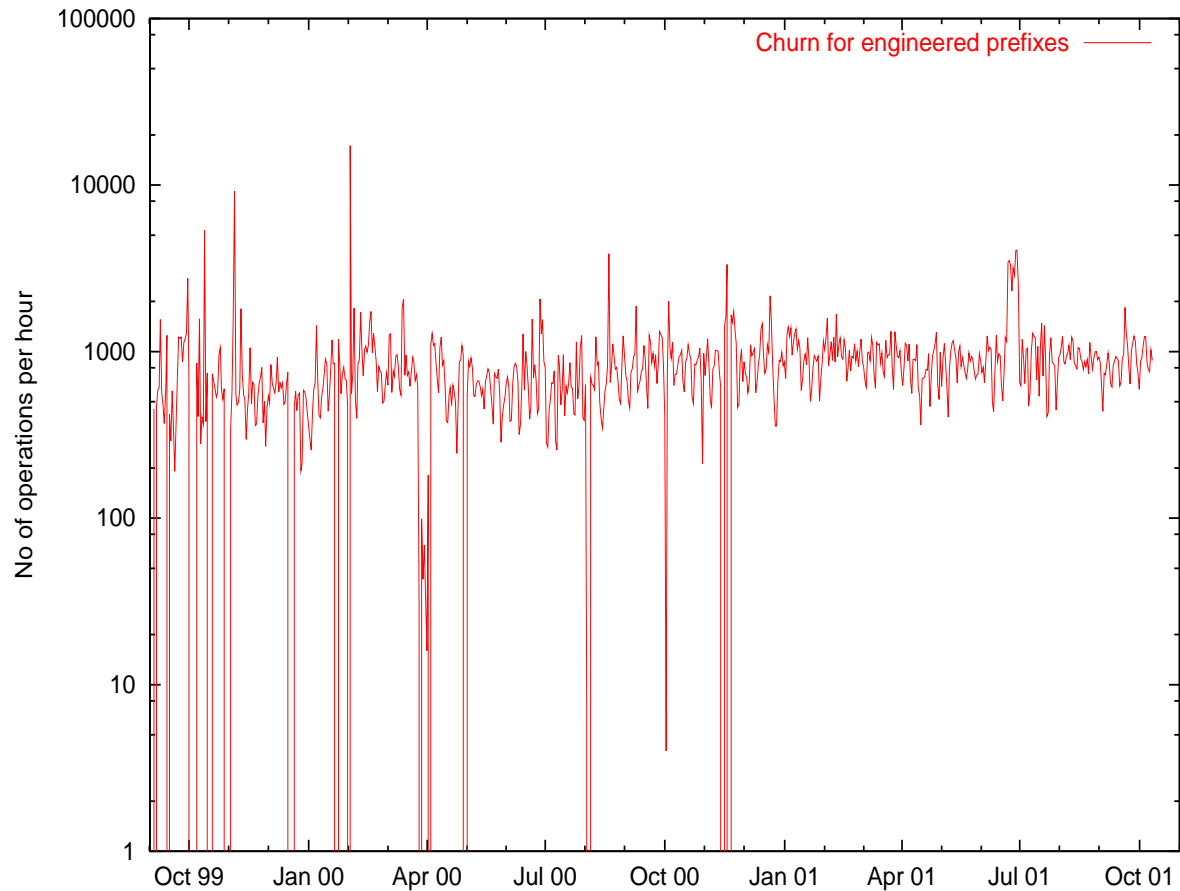
Overall churn is decreasing



Per prefix churn: stability



Engineered Prefixes behave differently



- Do people engineer prefixes dynamically?


Churn by Prefix Length



Major cause of churn


- Each time a BGP peering is lost, upto 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




- Peering loss/reestablishment is very costly

This is carried multiple ASes away



- 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
 - dont blame multi-homing: no evidence of exponential
- 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