



IPv6 Peering Thoughts

– IXPs, DOs and DONTs

Gert Döring

SpaceNet AG, München

4th European Peering Forum, Copenhagen, Sep 15th, 2009

IPv6 at the DECIX

- introduced 09 / 2001
- addressing scheme: `2001:7F8:::<isp1>:<isp2>`
 - e.g. for AS 5539: `2001:7f8::15a3:0:1/64`
- each ISP has 2^{32} bits to number their routers
- initial idea: reverse DNS delegation at /96 boundary – not followed through, not enough DNS churn to warrant the effort
- as of today, about 40% of the connected ISPs have IPv6
 - not all of them fully “production ready”, though
- DECIX route servers fully support IPv6

IPv6 at the DECIX / 2

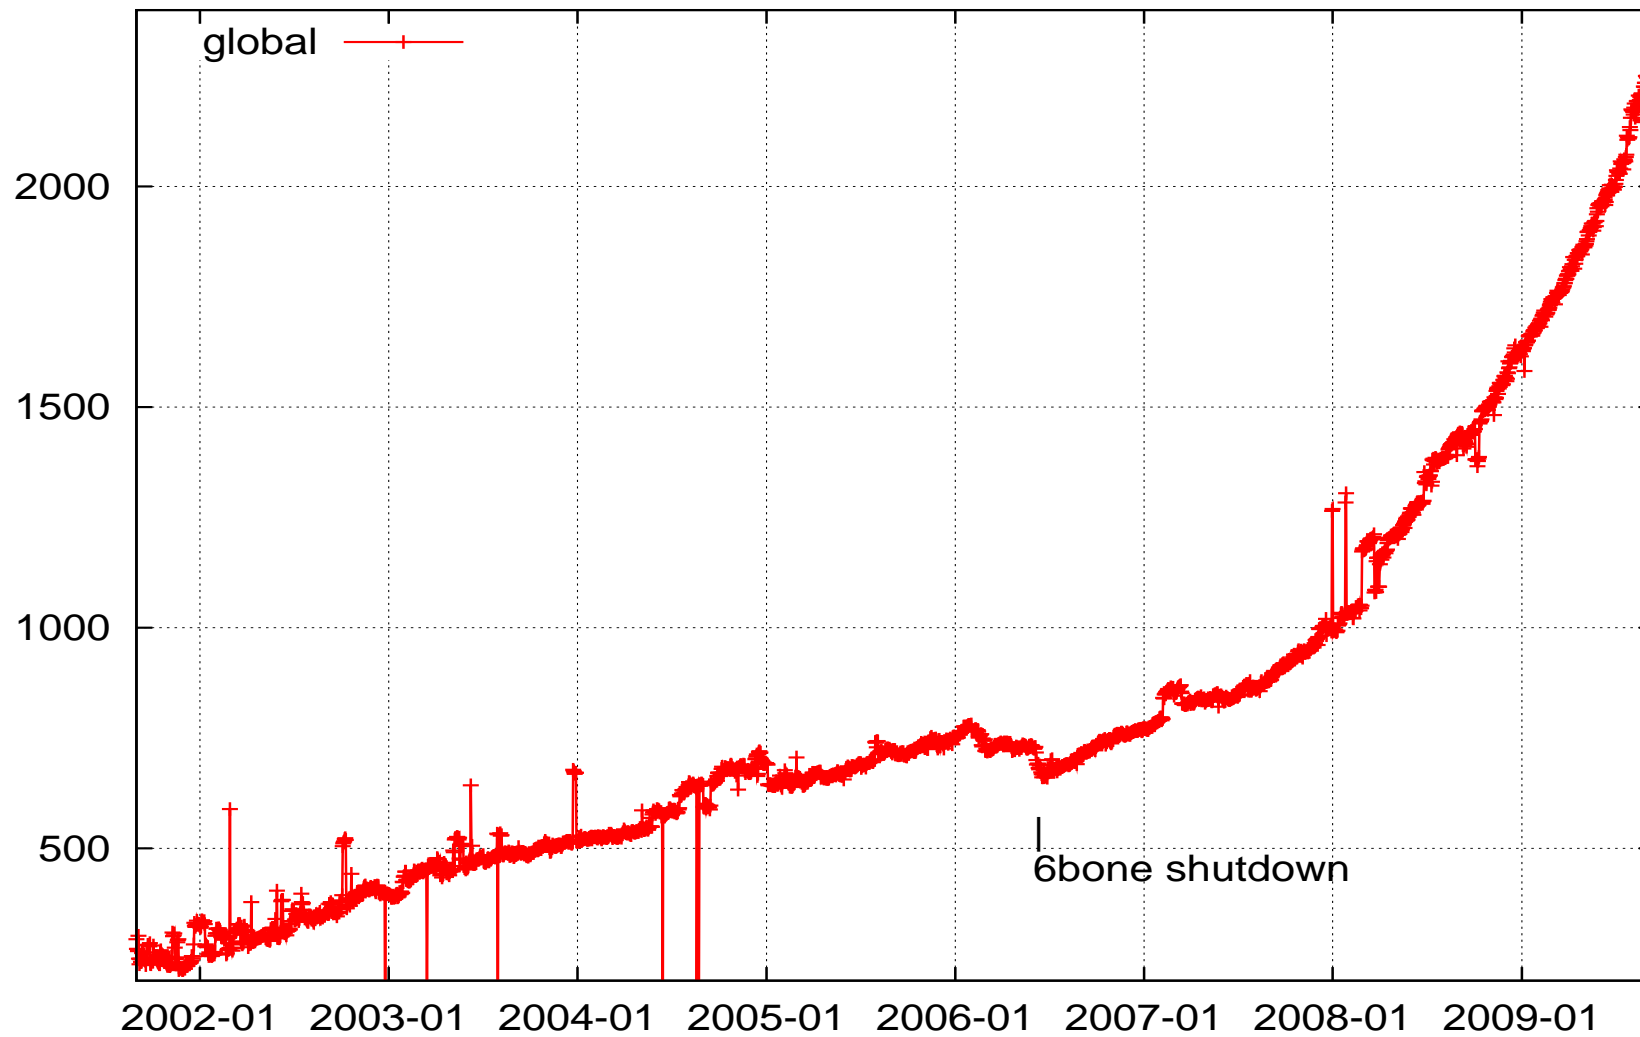
- in all this time, we had only one major problem with IPv6
 - the old Catalyst 6500 platforms had problems with IPv6 multicast (MLD) snooping and was blackholing certain neighbor discovery packets
 - no problems with current Force10 and Foundry platforms
- some issues with DECIX web site (Firewalls, RedDot CMS...), but that's unrelated to the IXP fabric itself
- traffic levels IPv6 vs. IPv4 are currently unknown

BGP thoughts...

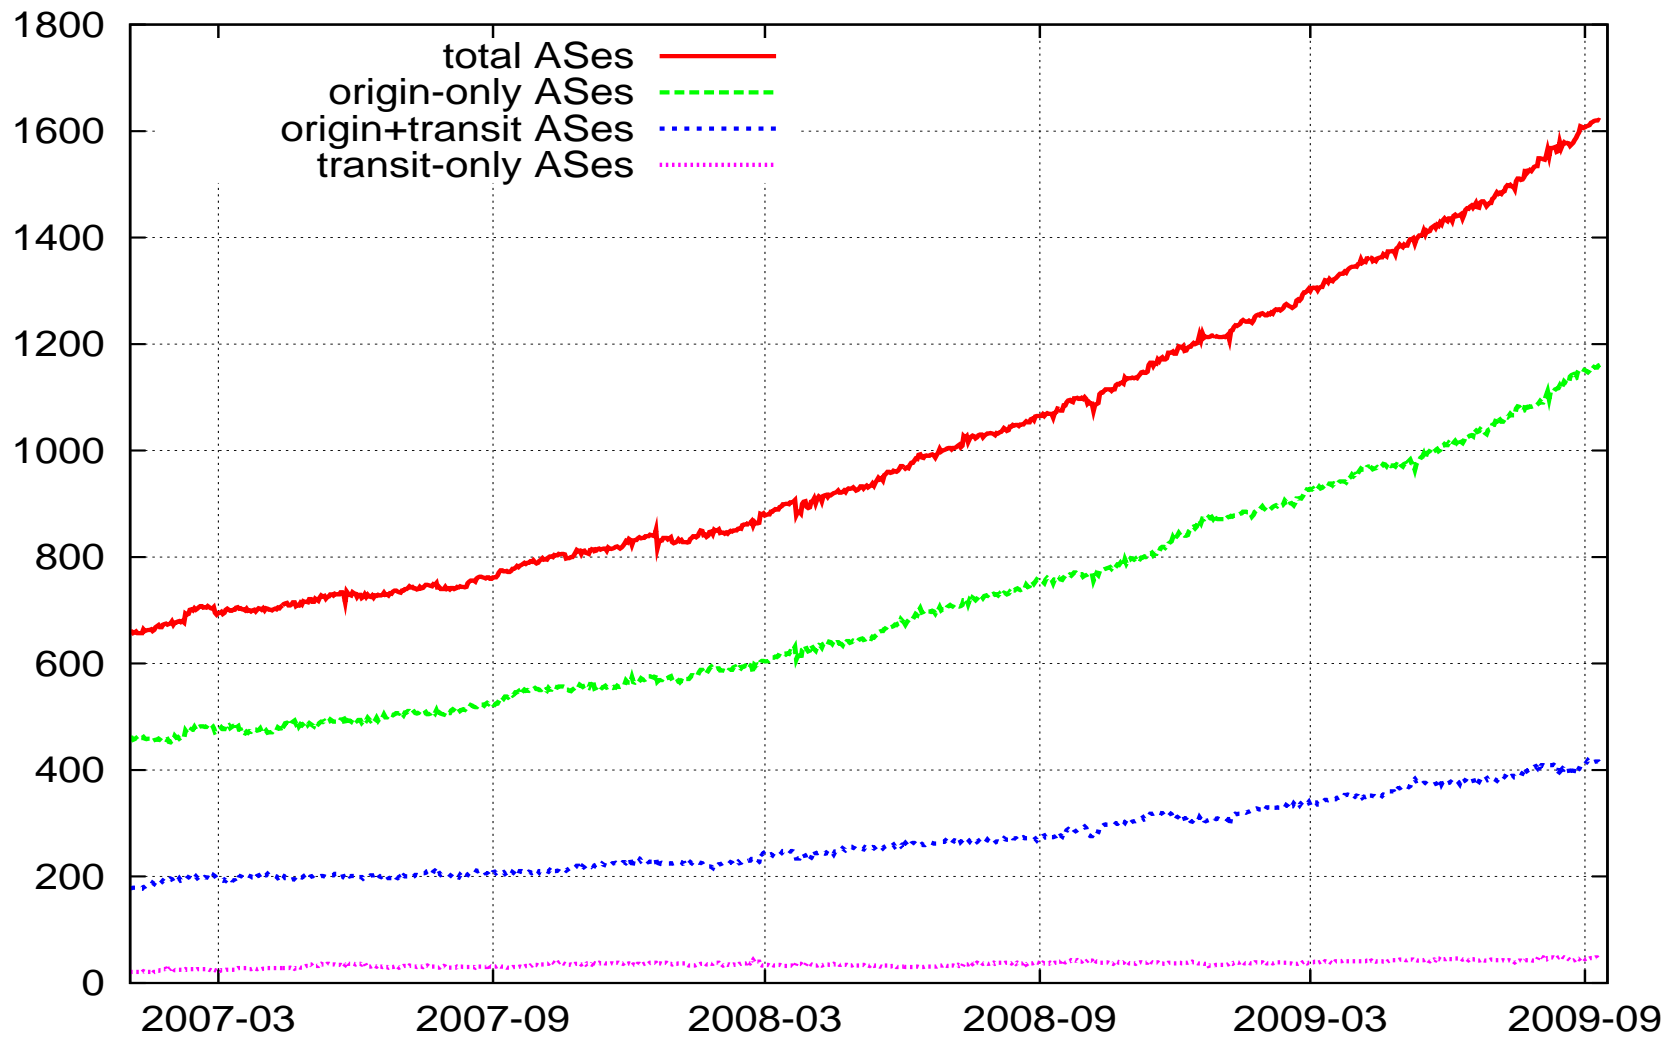
IPv6 BGP evolution - 1993

- initial goal was “get packets from A to B”,
 - no matter how many hops or how bad the latency
- BGP over intercontinental tunnels quite often the only way
- BGP “full table swaps” quite common
- at that time, a good way to start
 - build basic infrastructure to test OSes, Apps, ...
- only IPv6 early adopters would ever send or receive IPv6 packets
 - *no unsuspecting end users*

IPv6 BGP evolution - table growth (prefixes)



IPv6 BGP evolution - table growth (AS numbers)



IPv6 BGP evolution - 2009

- today: IPv6 needs to be run as professionally as IPv4
- operating systems come with IPv6 default-on
- web sites have IPv6 enabled
- \Rightarrow *unsuspecting end users*
 - if IPv6 connectivity is poor, users get hurt
 - if this happens too often, the lore will be “turn off IPv6”
- things that need to be avoided
 - (much) higher latency on IPv6 packet paths
 - weaker hardware / lower throughput on IPv6 path
 - sloppy monitoring of IPv6 network

BGP DOs

- accepting prefixes...
 - from customers: accept customer prefixes, *strictly filtered*
 - * *tag* prefixes to clearly identify as “customer”
 - from peers: accept peer network + peer customer prefixes, filtered if possible
 - from upstream/transit: accept ANY
- announcing prefixes...
 - to customers: announce ANY
 - to peers: announce customer (and self) prefixes
 - * use *community tags* to identify customer prefixes
 - to upstream/transit: announce customer (and self) prefixes

BGP DONTs

- do not send full tables to peers
- do not send full tables to upstream/transits
- *do not accept* full tables from customers or peers
- do not run IPv6 transit over high-delay tunnels (if more local alternatives exist)
- (of course there are special cases when some of this might make sense, but as a general rule: just don't do it)

The End

- now let the discussion begin...
- Questions? gert@space.net