Selected Musings on Challenges Facing the Internet

Thomas Narten
narten@us.ibm.com
September 13, 2008

Outline

- **Exponential Growth**
- **IPv4/IPv6**
- **Tragedy of the Commons and Route Scaling**
- **ICANN and Internet Governance**
- **Concerns about the future**

Challenges Facing The Internet

- **Continued exponential growth**
  - Ability of technology to scale
  - Ability of public policy/regulation to keep up (and stay out of the way!)
- **Victim of its own success**
  - Critical part of national infrastructures
  - Critical part of national economies
  - Governments tend to step in when they see threats at this level
- **Will lead to gut-wrenching changes to society**
  - Many will be good (social networks, online commerce, etc.)
  - Some will be disruptive (entire industries becoming irrelevant)

Exponential (?) Growth

- **ISC domain survey**
  - Estimated 625M hosts (i.e., individual DNS names)
  - Doesn’t include the many machines not listed in DNS
- **www.internetworldstats.com/stats.htm**
  - 1.67B Internet “users” (widely cited)
  - Defines 24.7% of world population as a “user”
  - China/India at early stages of adoption
- **Earth population**
  - 6.78B (today)
  - 9.3B (2050 projection)
- **Expect hundreds of devices per person (or more!)**

The Limitations of IPv4

- **Theory: 32 bits means 4 billion addresses**
- **Practice: addressing and routing tightly bound**
  - Can't really use 4B addresses (i.e., flat route)
  - Need to aggregate (like zipcodes, phone numbers)
  - Earth population 6.7B today, 9.3B in 2050
- **IPv4 address space exhaustion:**
  - August 2011 (IANA), July 2012 for (RIRs)
  - Assumes no “run” on addresses at end (hah!)
  - See Geoff Huston's work (http://www.potaroo.net/tools/ipv4/index.html)
- **IPv6 greatly expands the address space**

When IPv4 Address Pool Exhausts

- **IPv4 continues working, but no more allocations**
  - Major impact for (say) Cable or Wireless providers
- **Increased use of Network Address Translation (NAT)**
  - Triple NAT (or worse)
  - NAT within organizations
  - Management nightmare
- **IPv6 in use (that was the theory, reality is otherwise...)**
- **IPv4 address markets (buy/sell/trade) develop**
  - But many legal uncertainties about “ownership”
  - RIRs do not yet have policies in place to support “liberalized transfer”
  - Will mitigate but not resolve shortage of IPv4 addresses

Network Address Translation

- **IPv4 with NAT delays, but does not solve the problem**
  - Works well (mostly) in simple client/server world
  - Restricts model to client browsers contacting web servers
  - Increasingly problematic when devices also act as servers
  - Compromised performance, robustness, security, and manageability of the Internet
- **But, NAT is entrenched and seen as providing value**
- **End-to-end model at stake**
  - Critical for maintaining substrate for continued innovation
  - But entrenched interests largely happy with current model
  - Entire classes of potential applications at risk of never being deployed (peer-to-peer, ...)
- **General perception that NAT works “Just Fine” and there is no urgency to do anything.**
  - Engineers understand problem, but rest of world does not

IPv6 Benefits

- **Provides almost unlimited addresses**
  - More addresses cannot be retrofitted into IPv4
- **Plus**
  - Improved autoconfiguration
  - Improved support for site renumbering
  - Mobility with route optimization (important for wireless)
  - Miscellaneous minor improvements
  - Add security (IPSec)
    - But IPv4 has it now too…
    - And security is much harder than we thought...
- **NAT no longer REQUIRED**
  - Though likely will still be used

IPv4/IPv6 Co-existence Model

- **Dual-stack co-existence: hosts/network support both IPv4 and IPv6 simultaneously**
- **Applications use whatever is available (IPv4/IPv6)**
- **Applications use names, query DNS:**
  - A Resource Records (RRs) contain IPv4 addresses
  - AAAA RRs contain IPv6 addresses
  - DNS returns whatever has been configured
- **Long term: IPv4 quietly goes away…**
- **More likely: coexistence "forever"**
- **Key Point: invisible to the end user**
  - All happens below the covers
  - If end users have to know about IPv6, it will never happen

IPv6 Deployment Status

- **Good News**
  - Extensive research and test network deployment
  - 6Net, 6Bone, WIDE, UNH, TAHI, etc.
  - Commercial IPv6 deployment has begun
  - Policy makers, governments believe IPv6 is needed
  - Shipped with Vista, MacOS, Linux, etc.
- **Bad News**
  - Very little actual IPv6 traffic and usage
  - Stacks enabled, but most applications are not
  - Continuing chicken-and-egg problem; why go first?
  - ROI arguments not compelling (yet)
  - Consolidation of industry & economic pressures have led to general reduction in network investment
  - Little that can be done to force deployment & usage

Introducing New Technologies

- **Need credible transition plan for deployment**
- **Critical components that MUST be upgraded MUST have economic incentive to do so**
  - Altruism isn't enough
- **E.g., deploying browser is easy**
  - Just one application, few dependencies
  - Only client and server need to be enabled
  - Only client and server need to implement TLS
  - Dependency on underlying OS would be fatal

Difficulty of Upgrading Infrastructure

- **Not just an Internet problem**
- **Businesses want ROI**
  - Equipment that can be used 5+ years
- **Operationally, stability is good, change is bad**
  - Introducing any change (no matter how “good”) has associated risk
  - All changes require extensive testing ($$)
- **Internet openness is a great benefit**
  - But also means low margins, which means ROI is a continuing challenge
- **Need transition strategy that**
  - Supports coexistance between old and new
  - Yet provide compelling incentives to move to new
  - It has been said that one needs an order of magnitude benefit

Tragedy of the Commons

- **Town square has a nice patch of grass**
- **Local farmers bring their sheep (each brings one or two)**
- **Too many sheep eat all the grass ruining the commons**
- **Multiple individuals, each acting in own self-interest, ultimately destroy shared limited resource**

Challenge of Route Scaling

- **Default-Free Zone (DFZ) in Internet core:**
  - Routers must deal with explicit entries for all prefixes (i.e., destinations)
    - Every individual network requires prefix in routing table
    - Each prefix requires resources to maintain.
    - Cost of one prefix may be small, but aggregate cost of many is high
  - No “default” to pass the buck to
  - BGP used to exchange routing reachability info from which individual routing tables are built
- **Theory: use Provider Aggregate (PA) addressing**
  - Number each site out of provider's block
  - All of Duke (or Time-Warner, etc.) covered by one prefix; only one route entry per ISP in DFZ
  - Result: small & manageable routing tables

Challenge of Route Scaling

- **295K prefixes in DFZ (growing superlinearly)**
  - Update rate also growing (increased churn)
- **Aggregation conflicts with business requirements**
  - Unwillingness to renumber when changing ISPs (want to be able to change provider on moments notice with no renumbering cost)
  - Multihoming, traffic engineering, etc.
- **Larger tables require sophisticated routing hardware**
  - Some argue that Moore's Law does not apply to high-end routers
  - Challenges current business model (ROI) of ISPs
- **Biggest ISPs worried that we need to “wakeup” before it is too late**
  - If it breaks, Balkanization of Internet some sites will become unreachable
- **The Internet is the commons!**

There is No Central Authority

- **IANA (operated by ICANN) gives out addresses to RIRs**
- **Regional Internet Registries (RIRs)**
  - Hand out addresses to ISPs, end sites
  - 5 RIRs, each their own organization
  - Loosely coupled policy making
  - Not regulated by governments
- **ICANN has only loose oversight over RIRs**
  - RIRs predate ICANN
- **RIRs have no say over what ISPs actually do**
  - RIRs can't/don't say “this must be routed”
  - Leads to interesting policy discussions: end sites want liberal PI space allocation policies, ISPs say “that will kill us”

Internet Corporation for Assigned Names and Numbers (ICANN)

- **Coordinates assignment of Internet's unique identifier system**
  - DNS names (e.g., what new TLDs get created)
  - Responsibility for IP addresses (houses IANA)
    - But RIRs existed prior to ICANN's existance
  - IETF parameter registries
    - But IETF predates ICANN and controls all policy aspects of management
  - Critical goal: ensure uniqueness, stability and security of identifier system
- **Created in late 1990s as Internet exploded**
  - Need to get USG out of managing the Internet identifiers
  - Experiment in private sector management
  - Alternative to ITU/UN intergovernmental-based approach
  - May well be the worst form of Internet Governance, except for all the others that have been tried
  - Critical to have strong input from those who understand the systems and actually make it work

DNS: Why We Need a Single Root

- **End users want a single, global name space**
  - The DNS name I use should work EXACTLY the same for me as for grandmother in China, uncle Brazil, etc.
- **Fracturing of the root means that DNS name that works for me, has different meaning elsewhere**
  - Breaks key goal of interoperability
- **Folks that complain about ICANN often talk about alternative roots as being "better"**
  - In practice, they have not become popular (yet) because people understand the value of a single consistent root
- **All the "benefits" that alternative root salesmen claim ignore the hard policy questions that ICANN is designed to deal with (e.g., what happens if multiple entities want to use the same name?)**

ICANN Does Not Control Internet

- **Only “coordinates” global numbering system**
  - Done within context of Policy Development Process (PDP) - unilateral action could be disasterous
  - Changes only made after completion of process in which anyone can contribute
- **ICANN cannot:**
  - “turn off” the Internet (no influence over ISPs)
- **Does not control the root server operators, who actually run the servers that hold the DNS root zone**
  - Today, they follow USG/ICANN's lead, but might not do so if circumstances changed
- **Most of internet run by business and is unregulated**
  - Very limited ability of any one entity to “control”

USG & ICANN

- **ICANN does create the DNS root zone, but USG has a critical approval step:**
  - Nothing goes into root zone without USG approval
  - USG approval step is the real issue about USG control of ICANN/Internet
- **In theory, USG could remove ".ru" from root zone, but**
  - Would precipitate major international crisis (i.e., at UN level)
  - Would result in huge world backlash against current system & USG
  - Would not keep Russia off the net (they would quickly setup alternative roots that their sites would use)
  - Other countries likely to follow suit
  - Current Root Operators might well ignore the USG change
- **In practice, the concern about USG "control" is overblown**
  - But political discussions are rarely rational…
  - Who in congress wants to be blamed for "giving away the Internet?"

Challenges for ICANN

- **Obtain increased community and International legitimacy**
  - Much progress made in last 3 years; debate about whether ICANN should exist has finally shifted to let's work at making ICANN work better
  - Continuing focus on balance between government input vs. too much influence
- **Not a traditional regulator (so can't use traditional models)**
  - Many things that people want ICANN can do, it actually can't
- **Loosen USG's influence (both real and perceived)**
  - ICANN at another key point: Joint Project Agreement (JPA) expires October 1
- **Have the Policy Development Process work well**
- **Successfully deal with the many balls it is juggling simultaneously**
  - Deployment of DNSSEC
  - Rollout of new gTLD process
  - Rollout of IDN TLDs
  - Restructuring of its internal processes
  - Ending of the JPA and what follows

My Concerns About Future

- **Route scaling is a potential problem**
- **Not moving to IPv6 quickly enough**
- **Distributed denial of service (DDOS) attacks**
- **Misguided government regulation**
  - Not all regulation is bad, but most persons do not have sufficient understanding of what they are doing to get the details right
  - Governments notoriously bad at such things
- **Robustness and Security of Internet infrastructure on which the world increasingly depends**
  - Can we ever win the offense/defense game against SPAM?
  - What future damaging behavior will become common?

Distributed Denial of Service

- **Get a very large number of machines to send traffic to a target**
  - Botnets comprised of tens of thousands of compromised machines
  - Each compromised machine does a “little” bit of mischief (sending queries)
  - A million such machines overwhelms the target, blasting it off the network
  - Target: google, amazon, government networks, critical infrastructure (DNS servers, backbone routers), competitors, etc.
- **Problem isn’t the network, its that the machines using “the commons” are able to intentionally abuse it**
- **It’s all about motive (e.g., $$)**
  - Very sophisticated criminal groups involved, running things as a business
  - One can rent networks of compromised machines
  - Analysis of attacks show very high degree of sophistication
- **Fundamentally hard to prevent**
  - Economic motivation drives the bad guys
  - Trivially easy to locate compromised machines (number in the millions)

2008 Kaspersksy DNS Exploit

- **DNS based on caching model**
  - Client machines send DNS query to “caching resolver”
  - Caching resolver does the hard work of actual DNS resolution, returning a completed answer
- **Kaminsky found an exploit that allows one to insert false records into a caching server in just a few seconds**
  - E.g., redirect queries for bankofamerica.com to server that harvests passwords and credit card details
  - Patches made available that reduce, but do not eliminate threat
- **Only prevention is DNSSEC technology**
  - Add cryptographic signatures to data (to detect modification)
  - But not yet deployed; buisness case to deploy has been minimal
  - That has changed significanly in last year; DNSEC will be deployed in root within six months

Summary

- **The Internet has been wildly successful beyond anyone's dreams, with much more still to come**
- **Numerous challenges related to continuing growth, only some are technical**

Questions?

Backup

The Pressure to Deaggregate

- **End sites demand PI addresses (to avoid renumbering and provider lock-in)**
  - Use of PA goes against own business interest
- **Multihoming adds yet more prefixes into DFZ**
- **ISP/enterprise Traffic Engineering adds more prefixes**
- **RIR policies: for IPv4, conservation is more important than aggregation**
- **Acquisitions and mergers**
  - Selling off part of network means it no longer is part of the original aggregate

Multihoming

- **End sites want to multihome**
  - 7x24 reliability (not dependent on single ISP)
  - Load balancing (share traffic across ISPs)
- **Easy way:**
  - Advertise end site prefix to both ISPs
  - Let routing infrastructure select (best path, route around failed ISP, etc.)
  - Requires propagating individual prefixes into DFZ (which doesn’t scale)
- **How many sites want to multihome? 10M?**