

Death of the Internet Predicted!

Stephen Strowes
ENDS Seminar, 27/Feb/2008

Purpose of this talk...

- Probably more questions than answers here...
- This is a high-level, question-generating talk, intended to highlight problems usually swept under the carpet...
 - “Somebody else's problem” may soon be everybody's problem if issues are not addressed

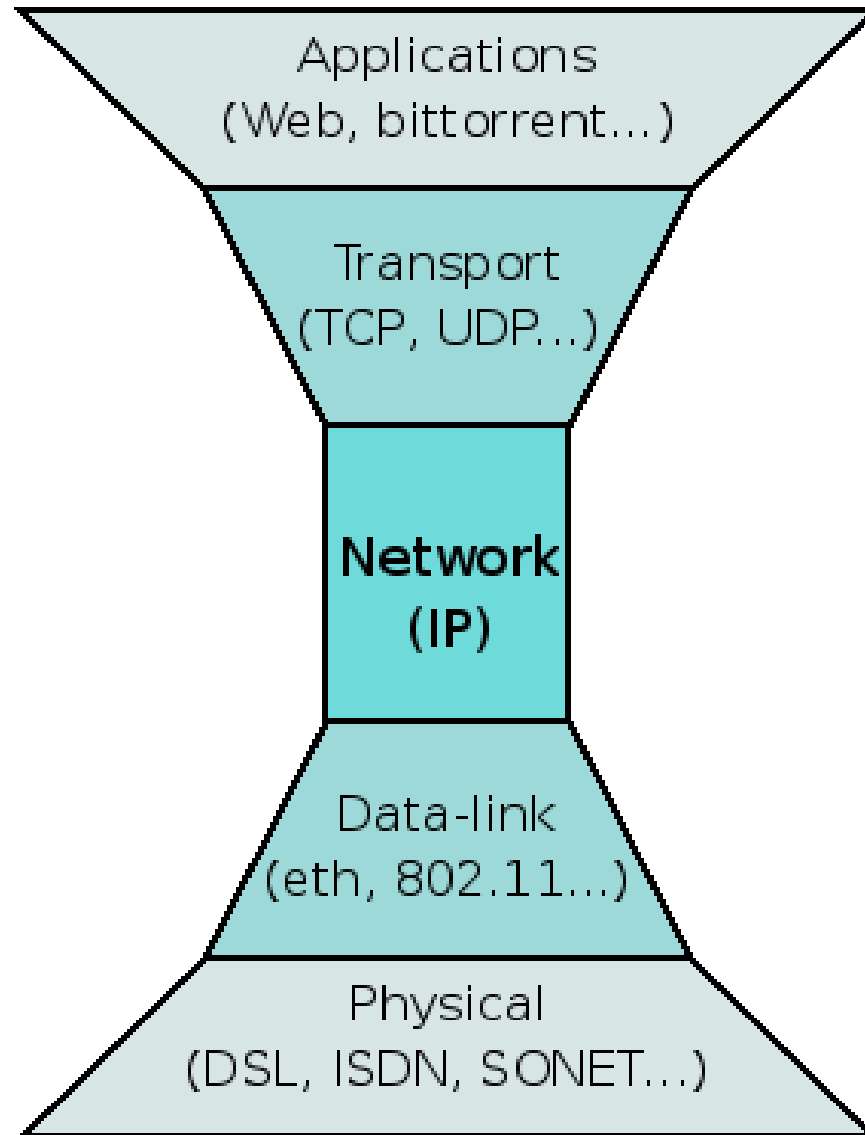
Outline...

- The existing internet architecture
- Issues facing the internet today
- Potential solutions

Outline...

- The existing internet architecture
- Issues facing the internet today
- Potential solutions

The Current Internet Architecture




The Current Internet Architecture

- IPv4 & Classless Inter-Domain Routing (CIDR)
 - IPv4 provides a 32-bit address space offering ~4billion addresses
 - Actually ~3.7 billion: reserved blocks, multicast...
 - CIDR provides variable-length network prefixes and address aggregation

The Current Internet Architecture


- Autonomous Systems
 - Coordinate with BGP
 - Path vector
- Dense mesh of nets in middle
- Aggregatable tree-like structure toward edges



Outline...

- The existing internet architecture
- Issues facing the internet today
- Potential solutions

BGP Scalability: FIB size in DFZ



BGP Scalability: Multihoming

- Why?
 - Gain in reliability
 - Two connections is better than one...

BGP Scalability: Multihoming

- Why?
 - Gain in reliability
 - Two connections is better than one...
- How?
 - Provider independent (PI) addresses

BGP Scalability: Multihoming

- PI address space is more prevalent than we'd like
 - It doesn't aggregate!
 - 59% of address allocations by RIPE over 2005/2006 were PI addresses...
 - But BGP and CIDR attempt to solve scalability issues with aggregation derived from hierarchically-defined addressing!
 - Customers want to retain IP addresses on move
 - Don't want to renumber when changing ISP
 - Don't want to be shoehorned into a technical design decision...

BGP Scalability: Trends

BGP Scalability: Trends

- Average AS-hop distance constant pretty much
 - Aggregatability works if the number of hops from source to destination increases with the network
 - Number of nodes on network continually increasing...


BGP Scalability: Trends

- Average length of advertised prefix is increasing
 - Many advertisements are for small networks
 - Average prefix length is ~/19
 - Advertised ASes are getting smaller
- Why?
 - NAT?
 - Bigger networks buying smaller blocks
 - Traffic engineering: more specific advertisements punch policy holes through existing aggregates for load balancing, policy agreements

BGP Scalability: Trends

- Key point:
 - Even if the network was not growing, BGP tables would increase in size
- BGP tables *could* be smaller... plenty of the entries could be aggregated
 - But we're always going to have bad management of data...

IPv4 Address Exhaustion



IPv4 Address Exhaustion

- IANA Unallocated Address Pool Exhaustion
 - 10-Apr-2011
 - RIR Unallocated Address Pool Exhaustion
 - 13-Aug-2012
- *Many* assumptions made in these predictions...

IPv4 Address Exhaustion

- IPv4 provides ~3.7 billion addresses
 - There are ~679M unallocated addresses
 - Of the ~3 billion addresses allocated, there are ~1.2 billion unadvertised...

Outline...

- The existing internet architecture
- Issues facing the internet today
- **Potential solutions**

Fixes?

- Reclaim IPv4 address space?
 - Many organisations have *way* more space than they need.
 - Fragmentation is inevitable...

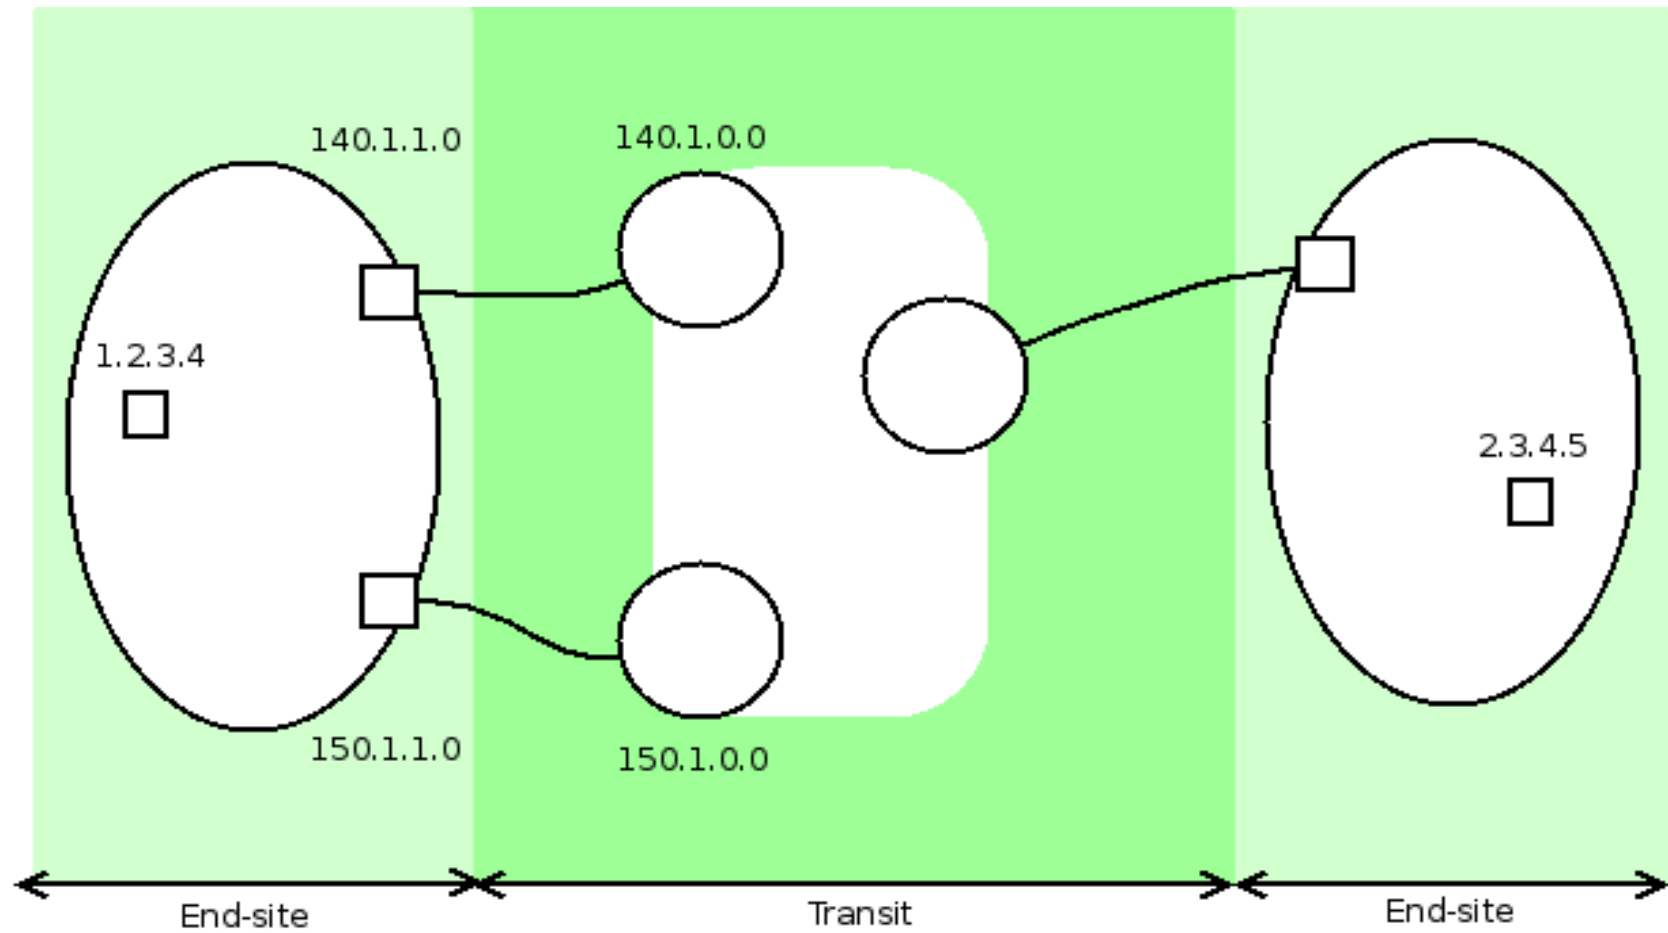
Fixes?

- IPv6
 - Provides the larger address space to solve the exhaustion problem...
 - But transition has been slow.
 - Size of IPv4 FIB: > 250,000
 - Size of IPv6 FIB: 1,104...
- IPv6 suffers from exactly the same scalability concerns as IPv4...
 - Aggregation *must* be enforced for it to work...

Fixes?

- Identifier/Locator split
 - The increased use of PI space for networks suggests that networks wish to have an identity they can take with them (much like mobile phone numbers)
- LISP: Locator/ID Separation Protocol (Cisco)

Fixes?



Fixes?

- LISP is considered a short-term fix!
 - NAT was considered by many to be a short-term kludge...

Fixes?

- New architecture?
 - Hierarchies in the IPv4 realm are shallow
 - The network is growing ever denser (peering)
 - For hierarchical routing to work, the distance between nodes must increase as the size of the network increases
 - If the average AS path length is essentially constant (around 3.5), then this is not the case...
 - Do we need a new routing mechanism in the core?
 - Hierarchical routing is the only contender *seriously* considered by the internet community

Fixes?

- Need a new routing algorithm in the core?
 - Lots of work on “compact routing”
 - Currently very theoretical, but refreshingly pessimistic with its claims
 - Would have to abandon shortest path routing
 - Is BGP shortest path?
 - No idea how policy affects compact routing
- Can a scalable routing algorithm even be designed, if policy is to be taken into consideration?

Conclusions...

- Okay, the world isn't ending, but...
 - “The Internet” is clearly not a solved problem
- There are short-term and long-term issues to be addressed:
 - Short:
 - Can we fix the address shortage problem?
 - Can we solve the scalability issues in the DFZ?
 - Long:
 - Can we build a system which is *provably* scalable?
 - Can we solve these problems in an *evolvable* manner?

Questions?

- And, some interesting places for more info:
 - <http://bgp.potaroo.net/>
 - <http://ipv4.potaroo.net/>
 - Analysing the Internet's BGP Routing Table, Huston
 - <http://www.potaroo.net/papers/ipj/4-1-bgp.pdf>
 - Locator/ID Separation Protocol (LISP)
 - <http://www.ietf.org/internet-drafts/draft-farinacci-lisp-05.txt>
 - “Why the Internet only Just Works”, Handley
 - <http://www.cs.ucl.ac.uk/staff/M.Handley/papers/>
 - “On Compact Routing for the Internet”, Krioukov, et al