20th NORDUnet Networking Conference

Problems in the Internet? Surely not!

Brian E Carpenter
Distinguished Engineer
Internet Standards & Technology
IBM Zurich Lab. &

April 2002

Topics

- Problems of success
- Problems of scaling
- Problems of thoughtlessness

  (Certainly, there are many problem areas I don't mention)
- If it isn't broken, don't fix it
- The way forward

The Next Generation of the Internet

Advancing e-business into the Future

OK, we made it, the Internet is official! Now let's see what success brings…

Computing:
The Grid

Information:
World Wide Web

Communications:
e-mail

Networking:
TCP/IP

Internet Society

ADV:
www.inet2002.org

Gold diggers – guess the year

(http://www.webcom.com/~walsh/)

The Commercial domain grew by over 10,000 in the first two weeks of Aug. Kraft Foods registered 133 product names … In the second two weeks the companies switched tactics. … Procter & Gamble started registering ailments, afflictions and body parts (e.g. diarrhea.com,pimples.com and underarms.com, etc.) 36 more.

Regulators & politicians

- National & international telecomms regulators find the Internet very tempting, but hard to get hold of. However, they are persistent.
  - When in doubt, make a regulation!
- Politicians also find it very tempting, and threatening (free speech? unwelcome material? tax free?). Also, they are unpredictable.
  - When in doubt, pass a law!
  - Never mind these geeks who say “that’s technically impossible.” Pass the law anyway.

ICANN

- Administer protocol parameters
- Coordinate allocation of address blocks to the regional registries
- Coordinate allocation of TLD names to TLD registries
- Coordinate root server operations
- *How can this **possibly** need 30 staff & $6M/year?*

Hubris

Function: noun

Etymology: Greek

Date: 1884

: **exaggerated pride or self-confidence**

(Merriam-Webster on line)

- Those who created the Internet have reason to be proud, but
  - should not lose sight of the real problems
  - should not ignore the impact of success on the original design principles of the network.

And a technical problem of success: Internationalisation

- We thought it was straightforward: rely on ISO 10646/Unicode (RFC 2277). But…
- Some uses of text are hidden entirely in protocol elements and need only be read by machines, while other uses are intended entirely for human consumption (presentation). Many uses lie between these two extremes, which leads to conflicting implementation requirements.
  - Humans can handle ambiguity, protocol engines can't
  - Humans care about cultural aspects, protocol engines are allergic to them
  - Thus, matching & folding requirements are different in the two cases

Problems of Scaling

- Reasons for continued growth and change:
  - Marketplace requirements
  - Technology and the appetite for technology feed on each other
  - Internet culture of open standards

Growth refuses to slow down

- Network costs can now beat Moore's law
- New countries are showing an interest
  - Let's bet on the 10 billion node Internet

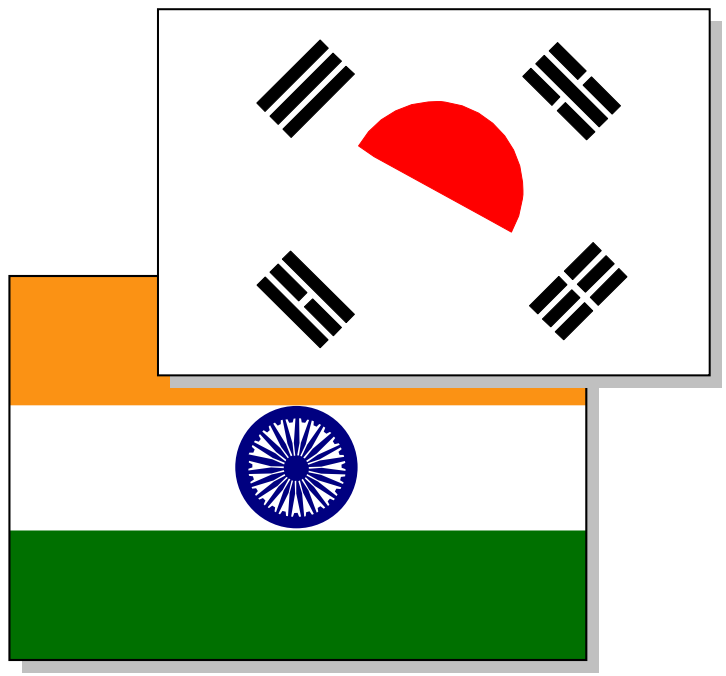

Scaling the address space

- Known problem since 1992
- Solution chosen in 1994
- IPv6 products since 1997
- Stable IPv6 standards since <2000
- So why is it so slow to start?
  - Operational costs of conversion; operational conservatism
  - Lack of strategic incentives in a fundamentally short-term industry
  - Pain from NAT is spread too thinly and not applied to the decision makers

What we miss by staying with IPv4 (1)

- The networked home & school
  - Entertainment becomes on-demand and largely interactive
  - Education… ditto
  - **Expand IT into every corner of life**
    - Needs broadband, but needs addresses too (interactive groups for learning or playing require peer-to-peer transparency)

What we miss by staying with IPv4 (2)

- Emerging markets
  - Only a tiny percentage of the world population have Internet access today
- Over the next 50 years, let's aim to get to all of them: make the market **20 to 50 times bigger.** Good for business, but good for society too.
  - Needless to say, we can't do this without enough address space

Geoff Huston's BGP graph

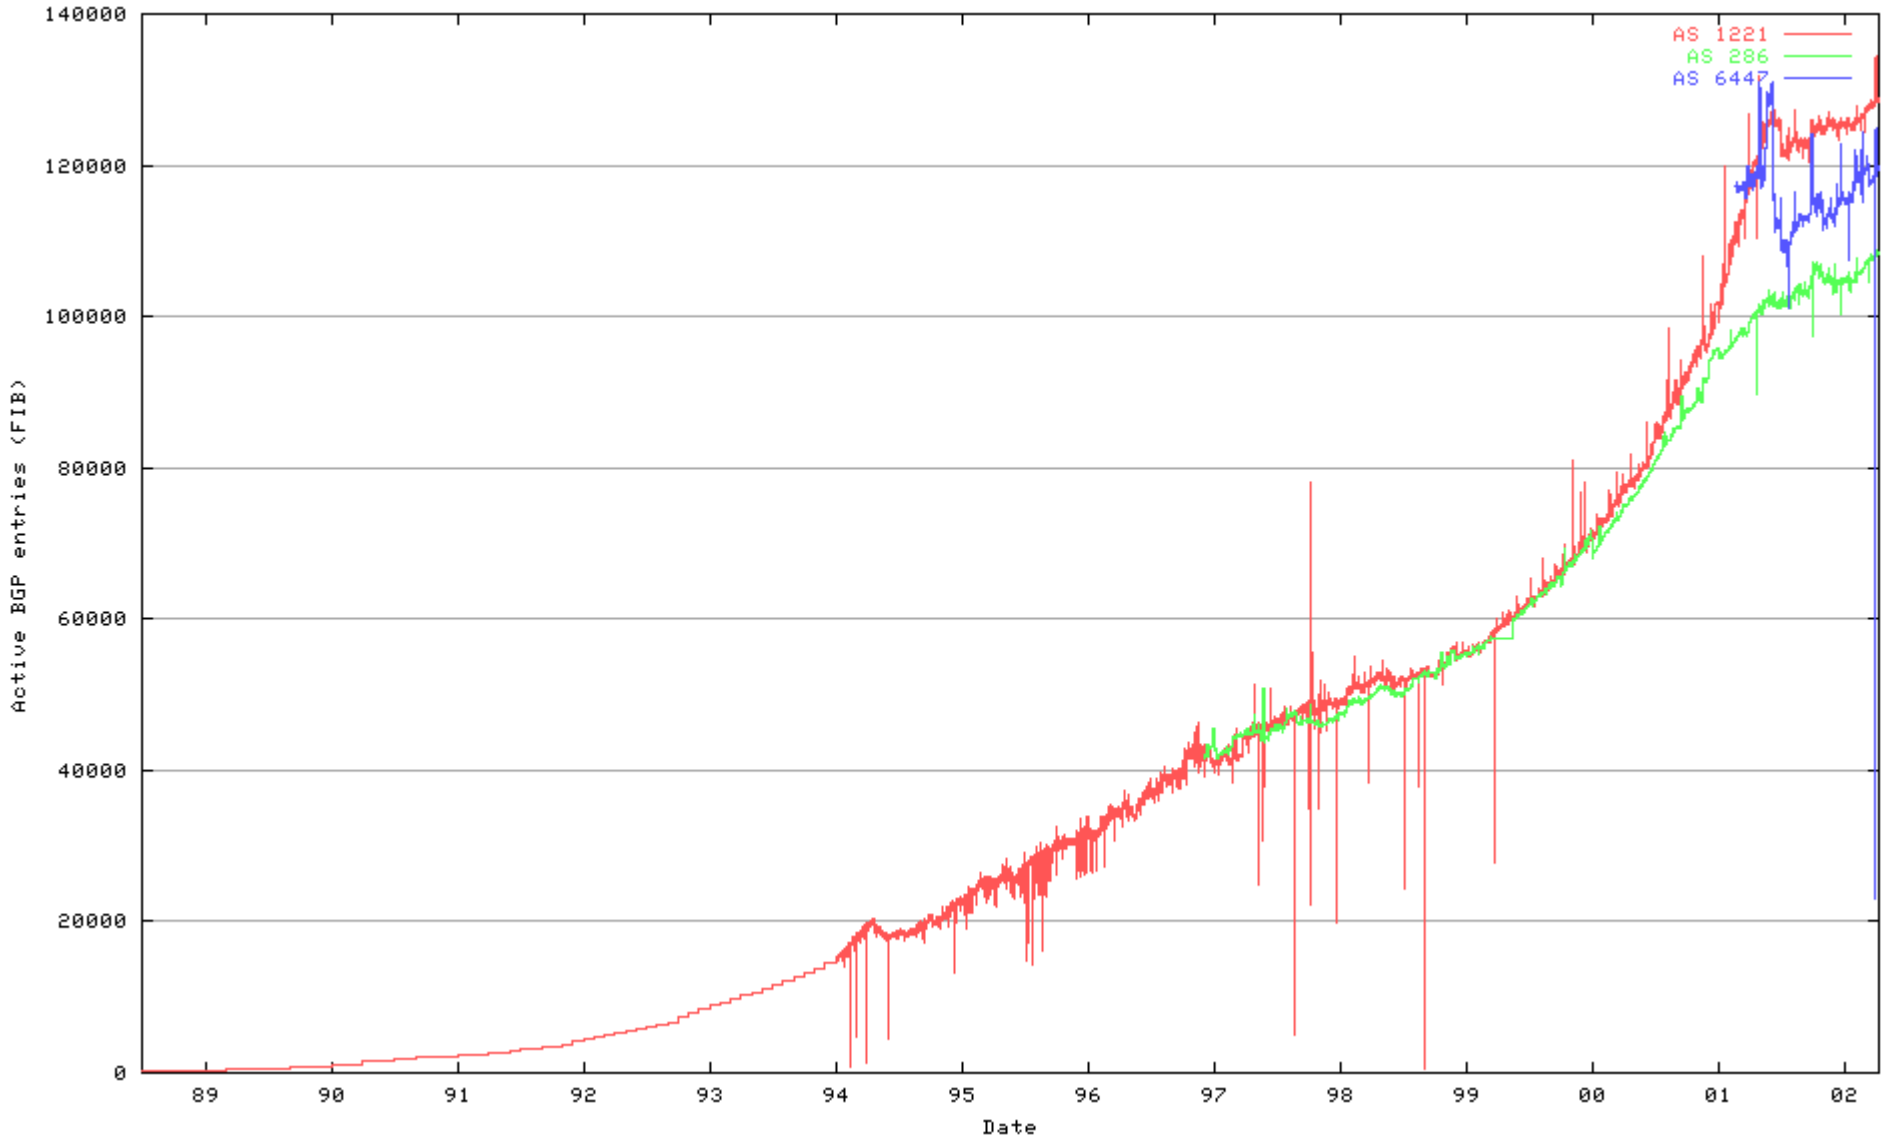

Scaling the backbone routing system

- Another problem known since 1992, but far harder in principle than scaling the address space.
  - See RFC 3221
  - See http://bgp.potaroo.net/ for the curve
  - BGP4+ is not adequate for much longer
  - Still a research topic, see http://www.irtf.org/charters/routing.html

Quality of Service viewed as a scaling problem

- We've invented session-oriented (RSVP) and stateless (diffserv) models for Internet QOS.
- Both technologies are available in widely used products. Neither has swept the world.
- Like IPv6: how can we get a new technology into the current practice of every network operator?
  - See RFC 2990

Standards organisations viewed as a scaling problem

- IETF (Internet Engineering Task Force)
- W3C (World Wide Web Consortium)
- GGF (Global Grid Forum)
- ISO JTC1 (Specific WGs of SC 2, 6, 25, 27, 29, 32, 34)
- ITU-T (various subcommittees)
- GSC (Global Standards Collaboration)
- ETSI (European Telecommunications Standards Institute)
- ECMA (formerly European Computer Manufacturers Association)
- ICTSB(European ICT Standards Board)
- CEn/ISSS(European IT standards portal)
- Telcordia
- Web Services Interoperability
- Eclipse
- OASIS
- P2P WG
- WAP Forum
- DVB (Digital Video Broadcasting project)
- IEEE
- ATM Forum
- Frame Relay Forum
- BlueTooth SIG
- Universal Plug and Play
- jini
- Salutation
- Home Audio Video Interoperability
- UMTS Forum
- 3GPP
- 3GPP2
- Network Processing Forum
- Mobile Wireless Internet Forum
- The Open Group
- New Productivity Initiative (NPi)
- OMG (Object Management Group, CORBA)
- OSGI(Open Services Gateway Initiative)
- Unicode Consortium
- JavaSoft
- IMC (Internet Mail Consortium)
- IPv6 Forum
- MPLS Forum
- Internet Software (DNS BIND)Consortium
- MINC (Multilingual Internet Names Consortium)
- IMTC (International Multimedia Telecommunications Consortium)
- Telemanagement Forum (formerly Network Management Forum)
- DMTF (Distributed Management Task Force)
- WfMC (Workflow Management Coalition)
- XIWT (Cross-Industry Working Team)

Problems of Thoughtlessness

Thoughtlessness

- Something *must* be done!

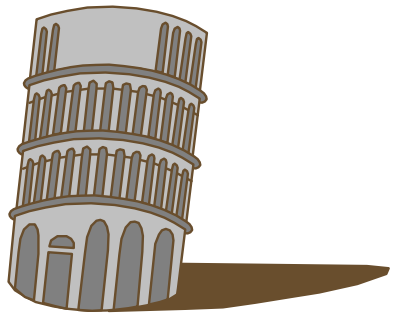

- This is something…

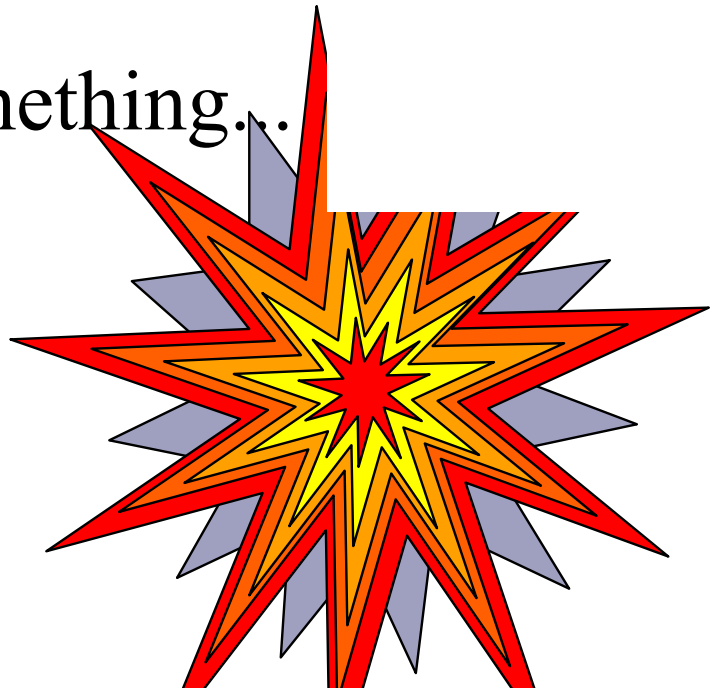

- Let’s do it!

Network Address Translation

- It was such a tempting quick fix…
- It could even be marketed as a security system (by pre-configuring it to allow nothing)
- And it breaks many non-client-server applications as well as network level security
  - See RFC 2993

Layer Violation Boxes ("Level 4 switches" etc.)

- Let's just peek into application layer headers…
- Let's just send this packet to a different server…
- Let's just proxy this request without being asked...
- Let's just rewrite this little piece here…
- They were all such tempting quick fixes
- Result: unpredictable, inexplicable glitches & failures
  - See RFC 3234
  - Middleboxes should be architected, not thrown together

Let's just put it in the DNS

- The DNS was narrowly designed, as a replacement for */etc/hosts* with distributed update and distributed lookup
- It was also designed to be extensible
- But it wasn't designed as a directory
- It is abused as a directory (pimples.com)
- It still isn't secured
  - See draft-klensin-dns-role-02.txt

Let's just run it over HTTP

- HTTP was narrowly designed, to carry HTML requests and responses
- It was also designed to be easy to use
- Firewall operators are bound to let it through
- But it wasn't designed as a transport protocol
- It is abused as a transport protocol & firewall penetration technique
  - See RFC 3205

The mythical PKI

- It was thoughtless to imagine that by creating technology capable of supporting a universal public key infrastructure, such an infrastructure would come into existence.
- As a result, we have a big challenge in actually deploying public key based solutions except within closed worlds.

If it isn't broken, don't fix it

- The basic Internet architecture remains remarkably robust under all these assaults, and under physical assaults.
- A current review of the basic architectural principles hasn't seriously faulted them.
- Most important, the end-to-end principle remains valid (i.e., don't trust the network to do anything you can do for yourself)
  - See draft-iab-arch-changes-00.txt

The way forward

- Engineers have responsibilities beyond the short-term. Just as this applies to civil engineers, electrical engineers, transport engineers, it now clearly applies to Internet engineers.
- In tackling all the challenges listed above, and new ones, we must apply an ethic of thinking about *long-term robustness and growth* as well as short-term operational and economic issues.
- We don't have the right to be thoughtless.