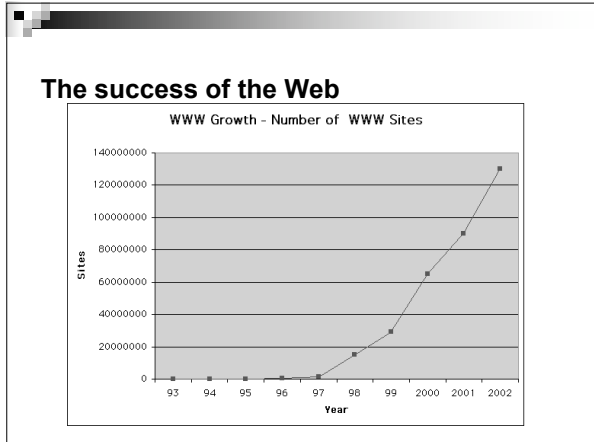
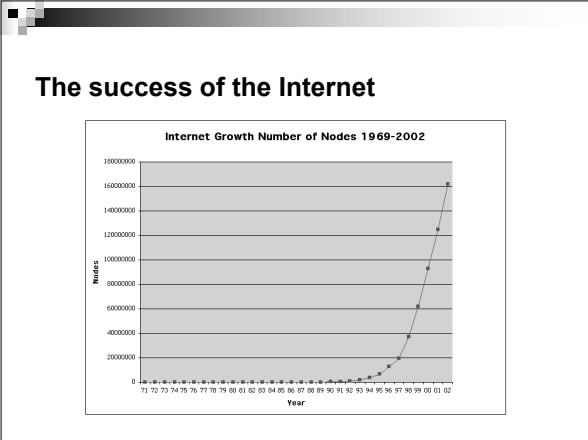


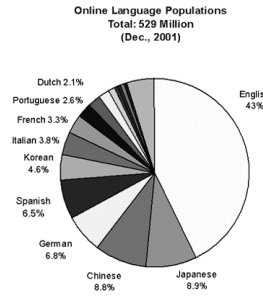
What's wrong with the Net?

Mark Handley
UCL Department of Computer Science



People Online

DATE	NUMBER	% POP	SRC
August 2001	513.41 million	8.46	Nua Ltd
August 2000	368.54 million	6.07	Nua Ltd
August 1999	195.19 million	4.64	Nua Ltd
Sept 1998	147 million	3.6	Nua Ltd
November 1997	76 million	1.81	Reuters
December 1996	36 million	.88	IDC



Online Language Populations
Total: 529 Million
(Dec., 2001)

Language	Percentage
English	43%
Japanese	8.9%
Chinese	8.8%
German	6.8%
Spanish	6.5%
Korean	4.6%
Italian	3.8%
French	3.3%
Portuguese	2.6%
Dutch	2.1%

The net is a success!

- The problem:
 - In almost every way, the Internet only just works!

The net only just works?

It's always been this way:

- 1975-1981: TCP/IP split as a reaction to the limitations of NCP.
- 1982: DNS as a reaction to the net becoming too large for hosts.txt files.
- 1980s: EGP, RIP, OSPF as reactions to scaling problems with earlier routing protocols.
- 1988: TCP congestion control in response to congestion collapse.
- 1989: BGP as a reaction to the need for policy routing in NSFnet.

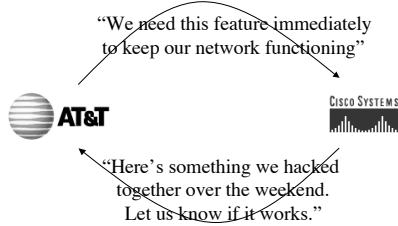
Changing the net.

- 1st Jan 1983.
 - Flag day.
 - ARPAnet switched from NCP to TCP/IP.
 - About 400 machines need to switch.
- As the net got bigger, it got harder to change.

Before web...

- Prior to the 1990s the Internet was primarily academic and scientific.
 - Common goals.
 - Low cost of failure.
- Then came the web, and commercialization of the Internet.
 - Exponential growth.
 - Financial costs of failure.
 - ISPs struggling to keep ahead of demand.
 - Huge innovation in applications.

Development Cycle




Running out of addresses...

- The current version of the Internet Protocol (IPv4) uses 32 bit addresses.
 - Not allocated very efficiently.
 - MIT has more addresses than China.
- IPv6 is supposed to replace IPv4.
 - 128 bit addresses.
 - We don’t need to be smart in address allocation.
 - How do we persuade people to switch?

Network Address Translators

- Scarcity of addresses has made addresses expensive.
- NATs map one external address to multiple private internal addresses, by rewriting TCP or UDP port numbers.



Network Address Translation

- Introduces asymmetry: can’t receive an incoming connection.
- Makes it very hard to refer to other connections:
 - Signalling, causes the phone to ring.
 - On answer, set up the voice channel.
- Application-level gateways get embedded in NATs.
 - It should be easy to deploy new applications!

The sky is falling!!!




- No.
- But we're accumulating problems faster than they're being fixed.

Imminent problems

- Address space exhaustion.
- Congestion control.
- Routing.
- Security.
- Denial-of-service.
- Spam.
- Architectural ossification.

Congestion Control

- The Internet only functions because TCP's congestion control does an effective job of matching traffic demand to available capacity.



Limitations of AIMD

- Very variable transmit rate is fine for bulk-transfer, but hard for real-time traffic.
 - TCP-Friendly Rate Control (TFRC)
 - Datagram Congestion Control Protocol (DCCP)

Limitations of AIMD

- Failure to distinguish congestion loss from corruption loss.
 - Wireless
- Limited dynamic range.

$$\square R \approx \frac{s}{RTT \sqrt{p}}$$

AIMD: Limited Dynamic Range

- One loss every half hour, 200ms RTT, 1500bytes/pkt.
- 9000 RTTs increase between losses.
 - peak window size = 18000 pkts.
 - mean window size = 12000 pkts.
 - 18MByte/RTT
 - 720Mbit/s.
- Needs a bit-error rate of better than 1 in 10¹².
- Takes a very long time to converge or recover from a burst of loss.

High-speed Congestion Control

- High-speed TCP (S. Floyd)
- Scalable TCP (T. Kelly)
- FAST (S. Low)
- Fair queuing + packet pair (S. Keshav)
- ATM ABR service.
- XCP (D. Katabi)

Routing

- BGP4 is the only inter-domain routing protocol currently in use world-wide.
- Lack of security.
 - Ease of misconfiguration.
 - Policy through local filtering.
 - Poorly understood interaction between local policies.
 - Poor convergence.
 - Lack of appropriate information hiding.
 - Non-determinism.
 - Poor overload behaviour.

Replacing BGP?

- BGP works!
- BGP is the most critical piece of Internet infrastructure.

- No-one really knows what policies are in use.
 - And of those, which subset are intended to be in use.
- No economic incentive to be first to abandon BGP.

Security

- We're reasonably good at encryption and authentication technologies.
 - Not so good at actually turning these mechanisms on.
- We're rather bad at key management.
 - Hierarchical PKIs rather unsuccessful.
 - Keys are a single point of failure.
 - Key revocation.
- We're really bad at deploying secure software in secure configurations.
 - No good way to manage epidemics.
 - Flash worm: infect all vulnerable servers on the Internet in 30 seconds.

Denial of Service

- The Internet does a great job of transmitting packets to a destination.
 - Even if the destination doesn't want those packets.
 - Overload servers or network links to prevent the victim doing useful work.
- Distributed Denial of Service becoming commonplace.
 - Automated scanning results in armies of compromised zombie hosts being available for coordinated attacks.

Denial of Service

- Traditional security mechanisms are useless for defending against DoS.
 - Attacker can force you to do expensive crypto operations.
- Many DoS point solutions have side-effects that can be exploited by an attacker:
 - Sendmail SPAM blocking.
 - BGP Flap Damping.
- The death of ping.

Security and DoS

- We need architectural solutions, not point mechanisms.

What are the right steps forward, architecturally, that doesn't turn this into the "Information SuperSkyway" (where you have to present your ID in triplicate, board an inconveniently scheduled and uncomfortable packet-herder, have an accredited professional guide you along prescribed and often-congested channels, to get approximately where you wanted to go)?

- Leslie Daigle, IAB

Architectural Ossification

- The net is already hard to change in the core.
- IP Options virtually useless for extension.
 - Slow-path processed in fast hardware routers.
- NATs make it hard to deploy many new applications.
- Firewalls make it make to deploy anything new.
 - But the alternative seems to be worse.
- ISPs looking for ways to make money on "services".
 - They'd love to lock you into their own private walled garden, where they can get you to use their services and protocols, for which they can charge.

Summary

- In almost every way, the net only just works.
- This is a *critical* time.
 - The net is moving out of it's infancy.
 - The problems are significant.
 - *We* get to influence it's future.