

# The Internet Five Year Plan

Geoff Huston

February 2003

# Why a Five Year Plan?

- Set useful goals
  - It's a long enough timeframe to think about more than the stuff in my mailbox
- Set realistic goals
  - Yet its probably a short enough time frame to understand what could be achieved
- And it has a fine set of precedents!
  - Many controlled economies have used the five year plan as the basis for national planning, so why not the Internet? 😊

# A personal set of views

- Predicting the future is easy – anyone can do it – and many do!
  - And this is just one set of views and opinions about the next 5 years on the Internet
  - There are as many views as there are people looking forward.....

# The Previous Plan

- Lets look at what's been achieved between 1997 and 2002

# 1997 Industry Issues



# 1997 Issues

- Transmission bandwidth was still relatively scarce and relatively expensive
- ATM was seen as the switching technology of choice for many
- Last mile technology was all about making 56K modems more reliable
- The fuss about the online economy was often seen as either too much enthusiasm or too much skepticism, or even a mix of both
- IPv6, VPNs, vLANs, VOIP, QoS, Mobility, GigE all hot technical topics

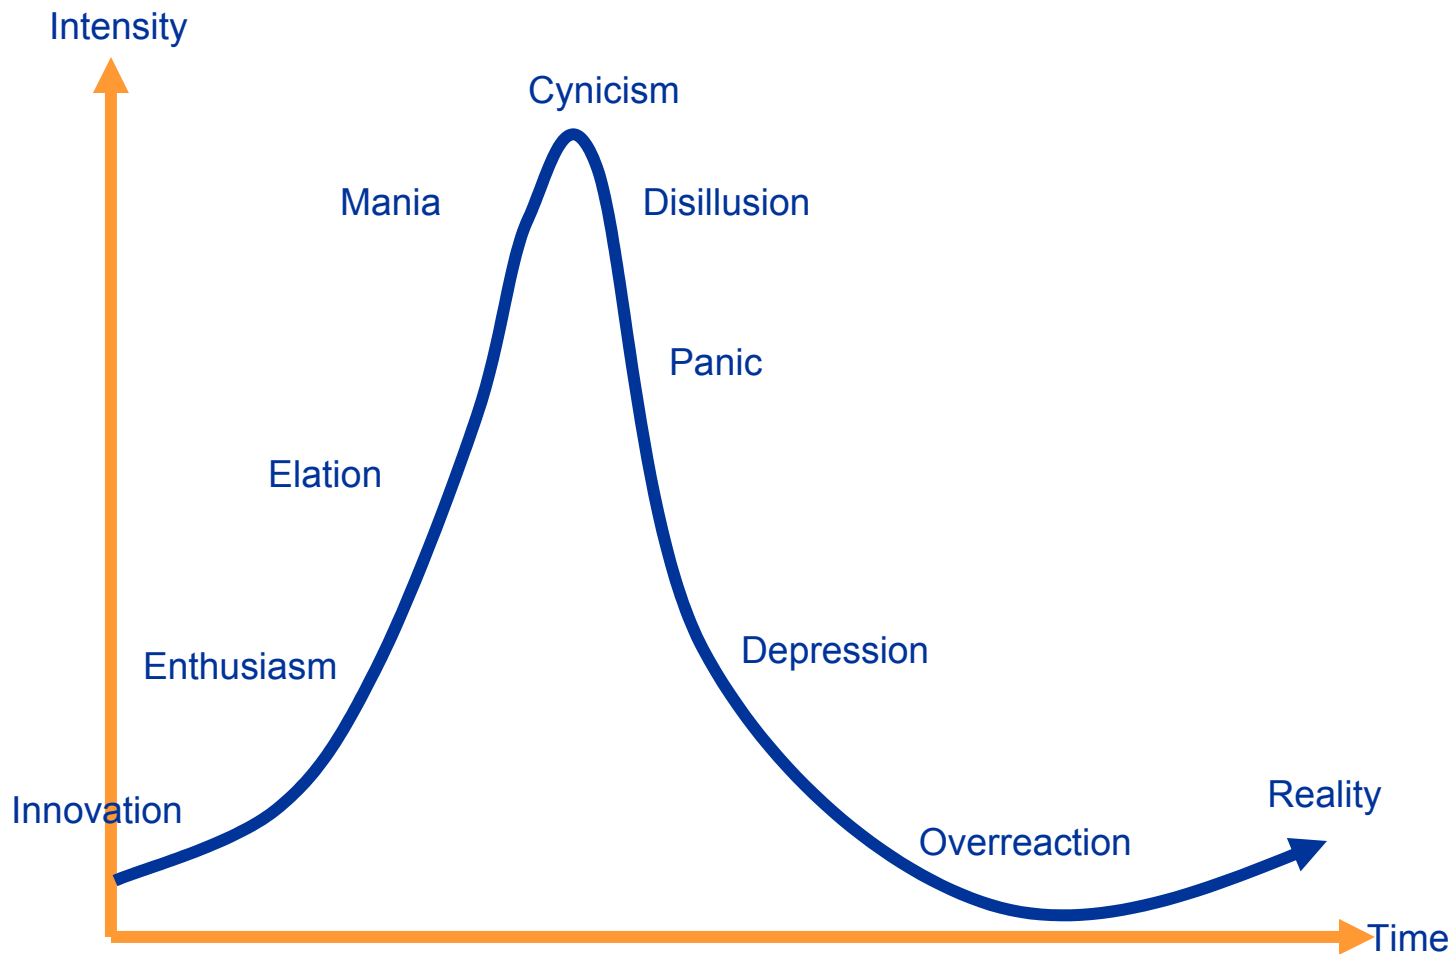
# 1997 Issues

- The value of the Internet was not proven or even widely assumed
- Non-Internet online products were still around at the time
  - Such as, for example, MSN
- ISPs and Telcos were generally on opposite sides of a pretty large dividing wall
  - ISPs viewed the telco with suspicion as a massive gorilla who was going to destroy their market
  - Telcos viewed the ISPs with suspicion as a competitive entrant with no background or experience in the sector with requirements that were often at odds with telco planning and expectation
  - A severe clash of cultures was inevitable

- The settings for an Internet boom:
  - Surprise factor
  - Progressive industry deregulation
  - High disruptive capability
  - Unknown market value
  - Unbounded utility model




# The Emotional Stages of a Boom



# The Last Five Years on NASDAQ

- Boom and Bust!

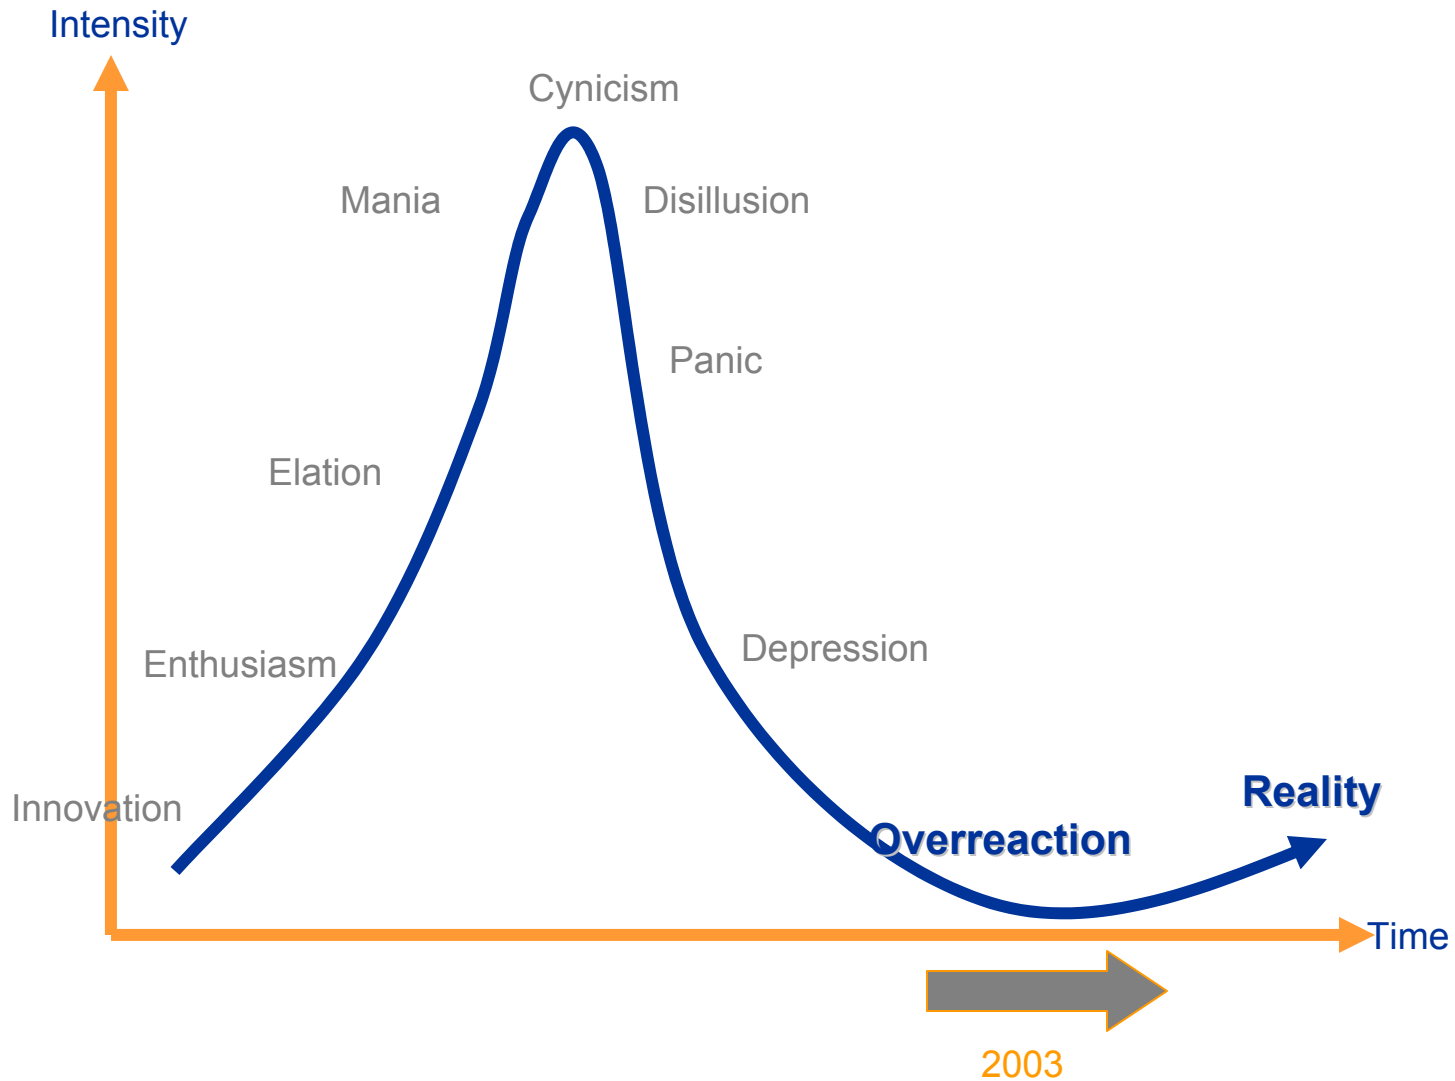
NAS/NMS COMPOSITE  
as of 16-Jan-2003



Copyright 2002 Yahoo! Inc.

<http://finance.yahoo.com/>

# Today



After: Gartner

- It takes more than wild-eyed optimism
- It's a business in a strongly competitive service market:
  - Know your customer
  - Know your product
  - Know your costs
  - Know your margins
  - Know your direction
  - And work hard!

- Some things will never go away!
  - Current IP technology-related issues:

IPv6, VPNs, vLANS, VOIP, QoS, Mobility, 10GigE

- And some things have been added
  - Added IP technology-related issues:

MPLS, DSL, WDM, optical switching, 802.11, Bluetooth, 3G

# The Next Five Years

- Predicting the Future is easy
  - The tough bit will be getting it right!
  - What follows are purely personal visions from the crystal ball ....

# From Optimism to Conservatism?

- A conservative period of controlled expansion rather than explosive growth
- Accompanied by greater emphasis on service robustness rather than service differentiation

# Security, Security, Security

- A highly visible security focus for the next few years
  - Encryption and authentication at many levels of the IP protocol suite
  - Dismantling aspects of the Internet's distributed trust model and attempting to replace it with negotiated conditional trust



# Shifting product emphasis

- A recognition of the existence of price/ quality tradeoffs in the IP market, similar to other markets
  - Differentiation will be based on quality of the product, not on definable attributes of the offered data that attempt to generate different responses from the network in terms of service quality

# Goodbye to Convergence?

- A shift in engineering planning?
  - Move away from the ideal of a mythical single converged switching platform
    - ATM is not a panacea for converged networks
    - And nor is IP
  - Move to accept an engineering base of multiple platforms with service specialty in each platform
    - Multiple wavelengths, multiple service platforms

# Multiple Networks

- Recognise IP's strengths and weaknesses
  - ★ IP allows adaptable traffic sessions to operate extremely efficiently over wired networks
  - ★ IP is not the optimal approach to support:
    - mobile wireless traffic
    - resource management requirements
  - ★ IP is not strong in supporting:
    - real time traffic under localized congestion events
    - various forms of traffic engineering applications
- Continued use of multiple networks to provide specialized service environments for various application sectors for some time yet

# Last mile considerations

- Concentrate on deployment of fibre and DSL based last mile IP networks
- Gradual phase out of modems as the dominant IP access device with a shift to emerging ubiquity of broadband last mile access
- Wireless is probably not a logical contender for last mile
- Hybrid Fibre Coax systems are capital intensive and often rely on a strong pay-TV market to provide some capital leverage
- Fibre is great – but its also capital intensive.
- DSL is a reasonable compromise for many environments

# Fibre Capacity – Five years of abundance?

- Dense Wave Division Multiplexing is lifting per-strand optical capacity
  - from 2.5Gbps to 3.2Tbps (320 wavelengths, each of 10Gbps per lambda) per optical strand
- Theoretical ceiling of around 75Tbps per strand
  - Probably achievable in five years using 40G or 100G per wavelength
  - But no obvious projected market demand levers to motivate high volume deployment of such technologies

# Unit cost movements

- Transmission costs per Gbps per Kilometer are coming down due to DWDM deployments – likely to continue
- Switching costs per Mpps are coming down due to Moore's law at work – likely to continue
- Installation costs per route mile are remaining constant – and likely to remain so, or possibly rise
- Per user volumes are rising – future trends of this metric are influenced by the adoption of various tariff structures and content patterns
- Growth in number of users will probably stabilize

# Network Management

- SNMP-based architecture
  - In-band management model
  - Query-response polling architecture using a structured set of query variables
  - Problems:
    - Insecure
    - Vulnerable implementations
    - Too simple?
  - Efforts underway to create a successor architecture to SNMP to incorporate better security, lock and confirm actions (mutex plus confirm), shared management state and greater levels of device control within the managed device

# IPv6 and the next five years

- V4 remains the overwhelmingly dominant protocol choice
  - 32 bit (4G) address space
    - 65% allocated
    - 32% deployed
    - 5%- 10% utilization density achieved
    - Consumption at a rate of 32M addresses p.a.
  - Anticipated lifespan of a further 10 years (at most) in native mode
  - Indefinite lifespan in NAT mode



# IPv6 and the next five years

- “IP with larger addresses”
- Address space requirements are no longer being easily met by IPv4
- This is an issue for high volume deployments including:
  - GPRS mobile
  - Pocket IP devices
  - Consumer devices



- Playback Zoom
- i.LINK (IEEE1394) IN / OUT
- Video IN / OUT
- S-Video IN / OUT
- Audio IN / OUT (Stereo)
- USB Terminal
- Intelligent Accessory Shoe
- Headphone Jack (Stereo)

- NPQM91: 370 min

#### Network Function

- Bluetooth Standard: Ver 1.1
- Email: SMTP, POP3
- Web Browser
- HTML: HTML3.2, Frame, JavaScript, SSL (V2/3)
- Image: GIF, JPEG, XBM, PNG

Sony DCRTRV950

# IPv6 Weaknesses

- Not sufficiently “different” from IPv4
  - No ‘value add’ to fuel investment in transition
  - Reuses large amounts of V4 infrastructure to there’s an expectation of identical outcomes
    - <http://www.kame.net>
- Not sufficiently “similar” to IPv4
  - The coupling of address and identity functions in the IP architecture makes transparent address translation a challenge
  - Referential integrity issues – is the DNS protocol independent or loosely/tightly coupled between V6 and V4
- Still working on the technology
  - Address architecture
  - Site-Local addressing
  - Multi-homing
  - Mobility
  - Transition mechanisms

# IPv6 Futures

- IPV6 appears to offer reasonable technology solutions that preserve IP integrity, reduce middleware dependencies and allow full end-to-end IP functionality
- Issues are concerned with co-existence with the IPv4 base and allowing full inter-working between the two protocol domains
- The next couple of years are make or break for V6
  - Either we start to deploy in volume
  - Or we better figure out how to live with the NAT mess!

# Routing

- IP uses a de-coupled routing architecture
  - Routing architectures can (and do) change without disrupting the service platform
- Two level hierarchy
  - Interior routing to undertake topology maintenance and best path identification
  - Exterior routing to undertake connectivity maintenance and conformance to external policies

# Routing – Interior Routing

- Predominant use of SPF algorithms for topology maintenance
  - OSPF
  - IS-IS
- Overlay external routes with iBGP
- Little evidence of takeup of MPLS-based approaches

# Routing – Exterior Routing

- BGP is the protocol of choice for exterior routing
  - Operator base highly familiar with BGP characteristics and capabilities
  - Easily disrupted
    - Poor security model with massive levels of distributed trust and no coupled authentication mechanisms
  - Poor scaling performance
  - Highly unstable (oscillation and damping)
  - Unresponsive to dynamic changes
  - No TE / QoS Support
    - And none likely!
  - No alternative to field!
    - And none likely!

- Where ATM collides with IP
- MPLS is an encapsulation technology that adds a network-specific egress label of a packet, and then uses this for each hop-by-hop switching decision
- Originally thought of as a faster switching technology than IP-level switching. This is not the case
- Now thought of as a more robust mechanism of network-specific encap than “IP in IP”, or “IP in L2TP in IP”
- Has much of the characteristics of a solution looking for a problem:
  - IP-VPNs? IP-TE? IP-QoS? Multi-protocol variants of these?
  - TE is the most likely outcome

- Sharing of a common base packet switching platform by a collection of IP networks
- Issues of integrity of the platform and integrity of the offered IP service to the VPN client
- Critical areas of technology development include
  - MPLS – Multi-Protocol Label Switching
  - MPR – Multi-Protocol Routing
  - VLANs – Virtual LAN Packet Frame formats
  - IPSEC – end-to-end IP authentication and encryption services
  - QoS – various forms of Quality of Service network mechanisms
  - PPP / MPLS / VLAN / VC inter-working – the enterprise-wide VPN service model
  - Dynamic VPN technologies – secure edge-based discovery tools



- Is community traffic isolation a network function or an edge-to-edge function?
- Network-based VPNs
  - Move from A-B circuits to meshed routing
  - Likely to continue this trend despite weaknesses in the current routing and provisioning model
- Edge-based VPNs
  - Continue to improve throughput rates for packet encapsulation/decapsulation and payload encryption
  - Possible breakthrough with automated point-to-point tunnel broker systems

- In theory voice is just another IP application
- In practice it's a lot harder than that - if you don't have heaps and heaps of excess bandwidth
  - Issues of Quality and Signalling
  - And VOIP is often most heavily used on the thinnest trunk routes
- VOIP is just one solution
  - It works well in small volumes on today's networks
  - Will it continue to operate at relatively small volume?

# VOIP as a QoS problem

- Quality
  - Voice is a low jitter, low loss, low latency, constant load application
  - TCP is a high jitter, medium loss, variable load transport
  - The problem is to get VOIP into the network without it being unduly impaired by TCP flows
    - Either overprovision the network and minimize the impacts or
    - differentiate the traffic to the network and allow the network elements to treat VOIP packets differently from TCP packets

# VOIP as a signalling problem

- How can you map the E.164 telephone number space into the Internet environment?
  - Allow VOIP gateways to operate autonomously as an agent of the caller rather than the receiver
  - ENUM technology to use the DNS to map an E.164 number to a URL service location
  - Use the DNS to map the URL service location to an IP address of the service point
  - What happens with NATs?

# The VOIP Plan

- No – the telephone network will still be there in 2007
- And it will probably still carry the bulk of the world's call minutes
- But
  - VOIP last mile may become common in the office environment
  - Hybrid control environments (Megaco, ENUM) may emerge
  - The PSTN may carry the call, but the IP network may manage the call setup characteristics.
  - Call features could become yet another desktop application

# Network Abuse

- Wouldn't it be so good if all spam suddenly disappeared?
  - It would be good if we could work out how to stop this form of abuse anti-social behaviour?
  - The true value of the Internet for most customers lies in person-to-person messaging
  - This value must be protected from abuse
- Either we may have to
  - give up some level of personal privacy,
  - increase the cost of the service,
  - pass more control to operators and regulators,
  - allow greater levels of regulatory oversight on the Internet,
- or
  - allow the network to be massively devalued by such abuse

# Security

- Just how ‘secure’ is IP?
- Not anywhere near secure enough!
  - DNS
  - Routing
  - Addressing
  - Content
  - Vulnerabilities are just about everywhere

- *In theory*
  - IP makes minimal assumptions about the nature of the transmission medium. IP over wireless works well.
- *In practice*
  - high speed TCP over wireless solutions only works in environments of low radius of coverage and high power
  - TCP performance is highly sensitive to packet loss and extended packet transmission latency



# Wireless

- 3G IP-based wireless deployments will not efficiently interoperate with the wired IP Internet
- Likely 3G deployment scenario of wireless gateway systems acting as transport-level bridges, allowing the wireless domain to use a modified TCP stack that should operate efficiently in a wireless environment
- 802.11 is different
- Bluetooth is yet to happen (or not)

# Wireless and Mobility

- IP is not comfortable over radio systems
  - IP requires low bit error rates, constant RTTs and high onset thresholds in order to operate efficiently
  - Adaptations to wireless have not been wildly successful to date
    - WAP
  - Likely adoption of nomadism as an adequate IP mobile solution
    - Few assumptions about the nature of mobility itself
    - No assumptions about suspended sessions
    - Incremental deployment models

# IP Transport

- Requirement to carry increasing volumes of payload at increasing carriage efficiency
- Current focus at the IP transport technologies of
  - POS (Packet over Sonet)
  - 10Gigabit Ethernet
- Issues of operations and management of these technologies, including robustness, resilience and progressive failure modes of operation
  - IP Routing protocol robustness, convergence and stability
  - Traffic Engineering technologies
  - Optical Wavelength Switching and Optical Cross-Connect technologies (GMPLS-based control architectures)

# IP Extensions & Refinements

- IP Multicast technologies
  - Extension of IP into support of common broadcast / conferencing models
  - Large-scale multicast
  - Small-scale multicast – conferencing
  - No widescale deployment as yet
- IP Mobility
  - IP support of mobility functions for mobile hosts and mobile subnets
  - Difference between nomadic operation and roaming operation
- IP QoS
  - IP support of distinguished service responses from the network
  - Per-flow responses or per-traffic class response models exist
  - No real uptake of either approach so far

# Services and Middleware

- WWW caching technologies
- Interception technologies
- Open pluggable edge service technologies
- Service provision and IP Anycast
- Directory technologies
- Public Key Certificate structures?

# So what can we expect?

- My personal list of expectations for the next five years:
  - No repeat of boom and bust
  - Conservative business objectives with conservative returns
  - Continued levels of regulatory interest to ensure that public objectives are being achieved
  - Continued expansion of the underlying infrastructure
  - Sector members with longer term objectives phrased more modestly than may have been the case in the past five years

# Thank You

- Questions?