A Decade of Technology Pitfalls and Successes

<*dino@procket.com*>

Define Decade

- I will take you back 0x10 years
- Too much good foundational stuff in 1985-1990 to not talk about
- IMHO
  - What failed
  - What succeeded
  - Where did the ideas *really* come from

Agenda

- What's the communication model?
  - Layer 2 or Layer 3
  - Datagrams or Connections
  - Point-to-point, Multi-point, Any-point, Broadcast
- What network layer do we use?
  - IP, XNS, DECnet, SNA, OSI, Appletalk, IPX, IPv6
- What routing protocols do we use?
  - RIP, GGP, EGP, Hello, [E]IGRP, OSPF, IS-IS, NLSP, BGP
- What management/security protocols do we use?
  - SGMP, CMIP, CMOT, SNMP, CLI, XML
  - IPsec, SSL, SSH

What's the communication model?

- Layer-2 or Layer-3
  - Mid-80s Transparent Bridging (TB) and Source Route Bridging (SRB) were popular
  - IP, XNS, and Appletalk were layer-3 competitors (all the way up to the app level)
  - DECnet Phase IV had been around and conversions were happening to Phase V
  - Novell IPX deployments increasing in "PC networks"
  - LANBridge 100 was the first "multi-port" bridge
  - IBM came to the scene with a 16mbps Token Ring Bridge (SRB)

What's the communication model?

- Layer-2 or Layer-3
  - DEC started doing 100mbps on the Gigaswitch (i.e. FDDI bridges)
  - Vitalink introduced non-LAN bridging by extending bridging to the WAN
  - Cascade built Frame Relay "Switches"
  - Fore built ATM "Switches"
  - These bridges could support all the "host-based" protocol packet forwarding
  - These bridges "tried" to support unlike media
  - Protocols like LAT and Netbios don't have a network layer

What's the communication model?

- Started seeing merchandise like:

What's the communication model?

- Routers did exist for a long time
  - BBN built "packet-switches" for the Military and National Labs (ARPAnet and MILnet)
  - DEC build DECnet routers for the enterprise
  - Apple built routers for the plug-and-play workgroups
  - Novell/3Com build devices for "PC-LANs" for brokerage/accounting firms
  - Proteon and cisco starting to build "multi-protocol" routers for research and education
  - Proteon and cisco built both SRB and TB bridge functionality too to enter lucrative enterprise market

What's the communication model?

- Routers were better than Bridges:
  - Routers dealt with unlike media easier
  - Routers dealt with arbitrary topologies better
  - Routers dealt with addressing that had opportunity for hierarchy and therefore scalability
  - Routers dealt with inter-organizational connectivity better

What's the communication model?

- Bridges were better than Routers:
  - They could be built cheaper, price-per-port was order of magnitude cheaper
  - They *had* little configuration to make work
  - Mind-set was you had to be an expert to manage routers
  - Lot more bridge vendors than router vendors, price is driven down even further

What's the communication model?

- Battle for the LAN data-link
  - Ethernet versus Token Ring
- Battle for the high-speed WAN data-link
  - ATM versus POS
- What device type to use
  - Routers or Bridges

What's the communication model?

- Battle of Ethernet versus Token Ring existed for nearly 10 years
  - IBM was really the only proponent for Token Ring
  - Workstation vendors and PC vendors endorsed Ethernet
  - Who won? Ethernet due to simplicity and huge commercial support

What's the communication model?

- Battle of ATM versus POS existed for ~5 years
  - Promise of ATM to the desktop was too ambitious
  - Building a routing architecture at layer-2 could be done, but we already had one at layer-3
  - ATM as a core fabric to connect routers was useful (DS-3 running out of steam), OC-3/OC-12 only on ATM switches
  - POS (IP over SONET) allowed high-speed connections between routers, no topology abstraction in the core fabric (routers were the core)
  - Who won? POS due to simplicity and router support (now more chip vendors support POS than ATM), today we do IP and IPv6 on 10gbps WANs

What's the communication model?

- What device type to use, Routers or Bridges?
  - "*Bridge where you can, route where you must*"
  - Bridges are good for workgroups, don't go with too many bridge-hops (you'll get into trouble)
  - Bridges are good as a core fabric to connect LAN-based Routers, don't go geographical (you'll get into trouble)
  - Don't forget your broadcast and multicast applications, that is what protocols you intend to run (you will unscale your bridges)
  - I'd say "*Bridge to make localized subnets, route everywhere else*"
  - Who won? Routers will never lose, Bridges are a commercial success due to Ethernet ubiquity (used in local-area only)

What's the communication model?

- Datagrams or Connections
  - X.25, SNA, ATM, and Frame Relay used connections below the network-layer
  - IP, XNS, IPX Appletalk, DECnet used connections at the transport-layer but datagrams at the network layer
  - OSI wanted to support both models
    - CLNS and CONS
    - CONS even over simple media like Ethernet
  - IEEE 802.2

What's the communication model?

- Connections
  - SNA was proprietary so one company could control the technology (this made it successful)
  - X.25 worked because of like minds (pre-EU or CCITT)
    - Telenet, Tymnet, Transpac, Datapac
  - ATM/Frame Relay like minds where circuit-switch oriented (all signals/data over one technology over the entire earth)
  - Connections could run over a datagram network
  - Harder to have datagrams over a connection oriented network

What's the communication model?

- Connections
  - Hard to have connections on each media type known to man
    - Microwave links, now WiFi
  - Hard to have QOS on each media type known to man
  - Reliable connections are not needed on reliable links
  - Most data-link connection protocols run in datagram mode (HDLC, 802.2, PPP)
  - Most of the dynamic connection-mode protocols (Frame Relay and ATM) run in PVC mode

What's the communication model?

- Datagrams
  - Most of the host based vendor protocols are datagrams at the network layer
    - IPX, Appletalk, XNS, DECnet, CLNS, IP, IPv6
  - Ubiquity of Ethernet and other datagram oriented data-links allowed datagram network layers to run effortlessly over links
  - Datagrams in packet-switches require less state
  - Datagrams allow fast rerouting with less overhead

What's the communication model?

- Datagrams versus Connections
  - Who won?

Datagrams!

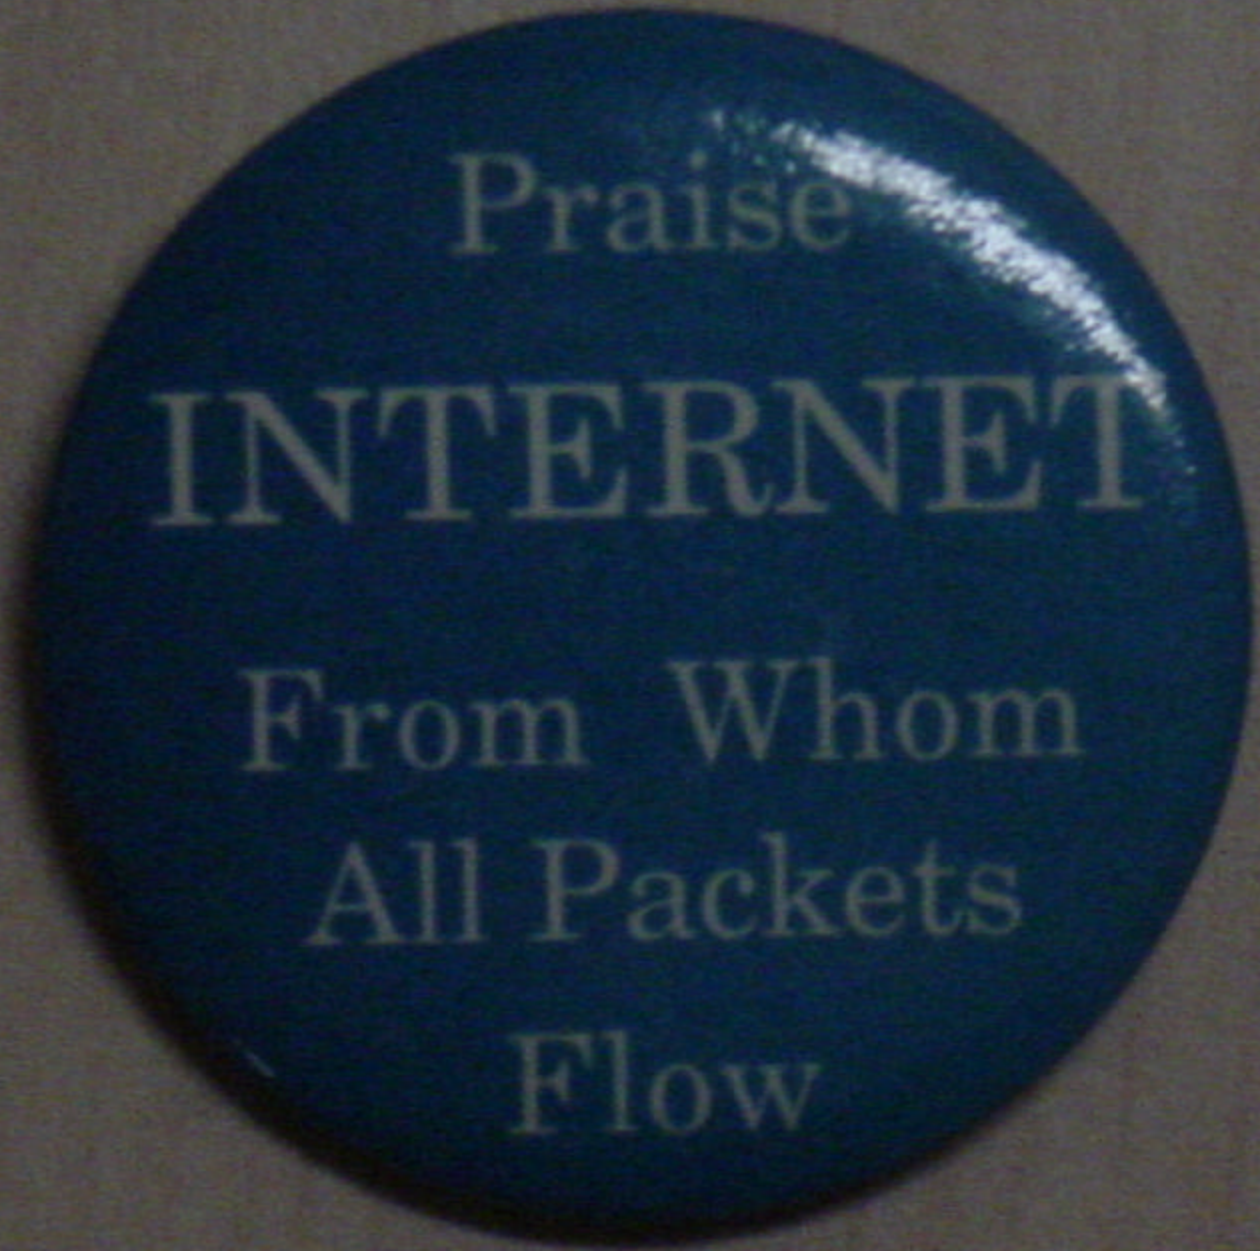

What's the communication model?

- Communication modes
  - 2-party communication (point-to-point)
  - Multi-party communication (multi-point)
  - Closest-party communication (any-point)
  - All-party communication (all-point/broadcast)
    - Within a reasonable scope

What's the communication model?

- Communication Modes
  - Layer-2 datagram protocols could do unicast, multicast, anycast, broadcast
    - LANs
  - Layer-2 connection protocols could only really do unicast and multicast
    - ATM PVCs, Frame Relay (didn't do multicast well)
    - SMDS and X.25 could only do unicast
  - Multi-mode Layer-3 protocol would have to simulate equivalent Layer-2 services
    - IP replication over Frame Relay (wastes bandwidth)
  - Only network layer protocols could support multi-hop, scalable, metric-based modes

What's the communication model?

- Communication modes
  - XNS
    - Point-to-point and link-local broadcast
  - DECnet
    - Point-to-point, link-local multicast, link-local broadcast
  - Appletalk
    - Point-to-point, broadcast, and later multicast
  - IP
    - Point-to-point, multicast, anycast, directed and limited broadcast
  - IPv6
    - Point-to-point, multicast, anycast

What's the communication model?

- For mode support, who won?
  - Appletalk really had the applications to back the multicast/broadcast modes
  - IP has the scalability and the more robust network layer support, mostly because it was the interoperable network layer of choice
  - All other variants failed
  - IPv6 desires to give the promise IPv4 gave

What network layer do we use?

- Why so many network layer protocols
  - No rapidly defined standards
  - IBM around for a long time and focused on reliable financial apps, hence SNA
  - Xerox came out with XNS but no real protocol other than SPP for transport
  - Novell needed a PC-workgroup solution so did a XNS copy-and-modify to yield IPX
  - DEC hit scientific community and really had the vision for datagram/dynamic routing protocols
  - Apple wanted devices to simply plug-and-play

What network layer do we use?

- TCP/IP
  - Military and Research Lab community used TCP/IP
  - UNIXes left XNS and went with TCP/IP
  - Sun, SGI, and other newer WS vendors commercially supported TCP/IP
  - Independents had TCP/IP implementations on PCs and VAXes
  - Router vendors like BBN, ACC, Bridge, Proteon, and cisco were early pioneers of network/routing layer protocols
  - DARPA and NSF dubbed IP defacto standard for building Internet infrastructure
  - When Microsoft jumped on the bandwagon there was no turning back – intro of Windows 95

What network layer do we use?

- What was wrong with TCP/IP?
  - But TCP/IP was not an international standard
  - ISO was building full-blown standards from layer-1 to layer-7 (i.e. OSI reference model)
  - US Government proposed GOSIP mandate for new contracts
  - DEC, a huge supporter of OSI, released DECNET Phase V, a complete OSI stack implementation
  - What about Internet applications, could they run over OSI?
  - Bottom line: took too long for vendors to field OSI products, the IETF was winning the "working standards" race
    - IETF motto: "*rough consensus and working code*"

What network layer do we use?

- Who won?
  - Should be obvious to everyone in this room


"The Power
of TCP/IP"
/usr/net '87
Sponsored by Excelan

What network layer do we use?

- Mid-90s problem
  - IPv4 has 32-bit addresses, running out
  - Too hard to reassign addresses to routing domains
  - Routing domain multi-homing problem is exploding the routing table size
  - “The sky is falling”

What network layer do we use?

- How to solve the IPv4 problem
  - LNAT, IPAE
  - OSI
  - SIP, PIP, CATNIP, TUBA, NIMROD, BigTen
- Who won?
- SIP had 8-byte addresses
  - Compromise by committee to make 16-byte addresses
  - IP version 6 was born

What network layer do we use?

- More merchandise became available:

What network layer do we use?

- What's good about IPv6
  - No adhoc standards, good documentation for all protocols day-1
  - We will never ever run out of addresses in anyone's lifetime (keep doubling 4 billion 96 times) ;-)
  - Designed by both host and router vendors together (and you gotta love the header format)
  - Like the IETF way, build standards through implementation

What network layer do we use?

- What's bad about IPv6
  - Did we solve enough problems in this shot
  - Routing architecture is a copy from IPv4
  - Too many transition mechanisms documented (no one learned from the DECnet IV to V conversion)
  - Design by committee made protocols more complicated than necessary (i.e. ND is not simple ARP)

What network layer do we use?

- Summary
  - IPv4 is here to stay for a while
  - IPv6 is gaining momentum and there is a swell of infrastructure support
  - You can certainly buy IPv6
  - We finally have seen the end of proprietary network layer protocols

What routing protocol do we use?

- What about the plethora of routing protocols?
- Did we have too many?
- Did we learn from our past sins?
- The industry is converging
  - We have less protocols now then ever

What routing protocol do we use?

- Chronology, start with RIP
  - Everyone had RIP
    - XNS, IPX, Appletalk, IP, and yes even for IPv6
  - Basic understanding of DV concepts
    - You have direct routes, you pass them on, your neighbor continues to pass them on
  - Metrics were weak
  - DV lends itself well to route filtering
  - DV lends itself to rerouting instability

What routing protocol do we use?

- Chronology, GGP and EGP
  - Milnet and Arpanet connected mostly single LAN domains together
  - Inter-AS protocol required (was simply a up/down detection protocol)
  - The advent of the “IGP/EGP” split
  - No metrics in these protocols
  - Needed AS-loop detection mechanism

What routing protocol do we use?

- For a long time
  - RIP/EGP was your multi-vendor interoperable routing protocol suite
  - cisco introduced IGRP due to weak RIP metric
    - Solved counting to infinity, hold down, and split horizon sooner than public domain RIP implementations
    - Vectored metric
  - gated happened
    - RIP, Hello (delay metric), and EGP

  - This was the state of routing protocols from '85-'91

What routing protocol do we use?

- Under development
  - Proteon working on OSPFIGP
    - Pronounced 'OS P FIG P'
    - Brought 'OSPF' to IETF
  - IETF working on BGP
  - DEC using IS-IS in DECnet IV and V networks
  - ANSI X3S3.3 standardizing IS-IS for ISO standard
  - At this point, too many protocols, interworking issues came up
    - Networks were melting down from route redistribution errors

What routing protocol do we use?

- War between IS-IS and OSPF
  - Integrated versus Ships-in-the-Night routing
    - OSPF worked only for IPv4
    - IS-IS carried IPv4, NSAP, and Phase IV addresses
  - ISPs felt too much machinery in OSPF, went to IS-IS
  - Enterprise liked the bells-and-whistles of OSPF, stayed away from IS-IS
  - CPU consumption for link-state, forced people to stay with IGRP


ACC
Advanced Computer Communications
3Com
OSPF
PROVEN
INTEROPERABILITY
proteon
WELLFLEET

What routing protocol do we use?

- BGP was progressing in IETF
  - cisco lead in BGP, no real interoperability other than with the public domain BGP
  - Proteon lead in OSPF, lots of interoperability support by new router vendors
  - Interworking issues with OSPF/BGP
    - Should you redistribute BGP into OSPF
    - How could you use OSPF tag to carry BGP info
    - Quickly learned to carry BGP info in iBGP
- IDRP was progressing in ISO
  - BGP3 didn't carry subnet masks
  - Use IDRP for IP (and CLNS) and don't rev BGP

What routing protocol do we use?

- BGP4 deployment
  - Take a couple of implementations and a dozen motivated individuals in the right organizations
    - You get a turn-key BGP3 to BGP4 conversion
    - Amazing how this ran so smoothly
  - The Internet now had CIDR based routing
    - Sorry IP doesn't get credit for this
    - See Ross Callon's NSAP longest match routing spec

What routing protocol do we use?

- Mid-90s, router vendors seek alternatives for their customers
  - cisco comes out with EIGRP for IGRP-deployed customers
  - Novell does a copy-and-modify of IS-IS to yield NLSP
  - OSPF increases deployment in multi-vendor networks
  - IS-IS increases deployment in ISP backbone networks

What routing protocol do we use?

- Late 90s to present
  - More IGRP networks go to OSPF or IS-IS
  - IPv6 gains popularity
    - IS-IS becomes integrated again and carries IPv4 and IPv6 routes
    - BGP4+ carries IPv4 unicast/multicast prefixes as well as IPv6 unicast/multicast prefixes
    - OSPFv3 used for IPv6 while OSPFv2 used for IPv4 in SIN mode
  - People continue to put more baggage into BGP4

What routing protocol do we use?

- No mention of LDP and RSVP-TE?
  - They are not routing protocols
  - They are label binding protocols
  - Nothing more needs to be said

What routing protocol do we use?

- Summary
  - BGP is here to stay --- it is the hammer for every nail
  - Both OSPF and IS-IS will continue to support the base IGP infrastructure
  - Proprietary routing protocols nearly gone
  - No surprise to anyone in this room

What management protocols do we use?

- Just a brief word about network management protocols
  - There were too many
  - That has been fixed for a long time
  - MIBs are good, you gotta have them
  - Can MIBs alone completely manage a router
    - No, but we can name a lot of objects
    - Is good for proprietary extensions

What management protocols do we use?

- How did we start
  - HEMS and SGMP
  - Proteon had first SGMP implementation
  - SGMP was the basis for ubiquitous SNMP
  - OSI had CMIP, more functionality than SNMP
  - War between CMOT (CMIP over TCP) or SNMP
    - OSI camp lost again
    - Too many vendors went the SNMP path

What management protocols do we use?

- General rule of thumb
  - SNMP used for monitoring
    - Standard, you can count on
    - For your GUI interfaces
  - XML for standard polling and customizing output formats
    - Scripting in the hands of the customer and not the vendor
  - CLI to debug hard problems (or anything else)
    - Development engineers have their hidden commands
    - CLI output format different based on protocol implementation choices

What management protocols do we use?

- What the industry has said
  - People are used to the industry standard CLI
    - cisco IOS CLI
    - For reduced training cost to operators
  - People like the consistency of a structured CLI
    - Juniper JUNOS CLI
    - gated syntax
    - For network engineers to program/customize

What security do we use?

- Just a final word about security protocols
  - IPsec is great!
  - Why is it not implemented everywhere?
  - And since it is not, we have application level solutions (i.e. SSL and SSH)
  - Why don't vendors default it to on?
    - If there was a key distribution problem why don't I see it with SSL and SSH
  - If we want a secure network, let's stop the lip service and deploy end-to-end IPsec ASAP
    - No interim or transition mechanisms please

Final Summary
- The Internet Way -

- Open source
- Public domain
- Learning through experiment and evolution
- Can't live without it
- Thank you ~~Al Gore~~ Vint Cerf ;-)

Questions?