



X-Bone + 10

A Decade of International Collaboration

Joe Touch

Postel Center Director,
USC/ISI

Research Assoc. Prof.,
USC CS & EE/Systems



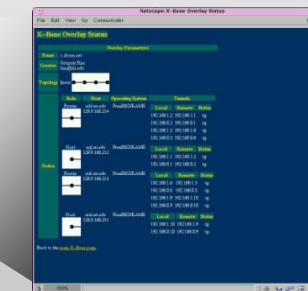
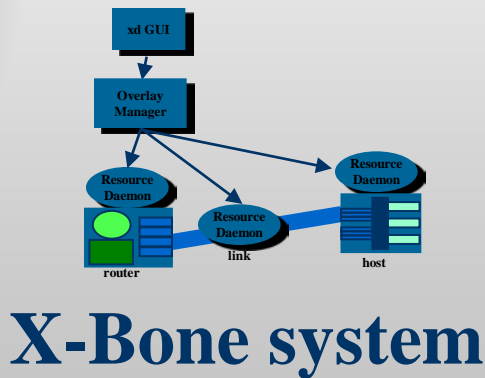




X-Bone Overlay System

Web GUI



Multiple views




Automated monitoring




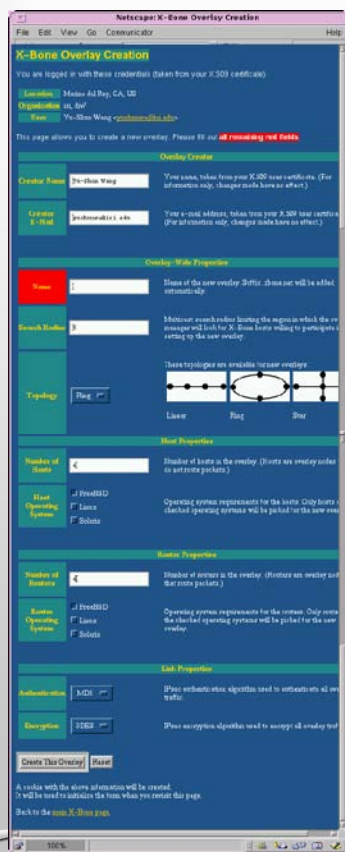
Timeline

- 1997 – first whitepaper
- 1998-2001 – X-Bone (DARPA)
 - IP overlays with revisitation, recursion (LISP)
 - 2000 – running code (FreeBSD, Linux)
 - 2000 – application deployment
 - 2001 – TetherNet “NAT-buster” to support demos
- 2001-2004 – DynaBone (DARPA)
 - 800-way spread-spectrum parallel overlays
 - 15-level deep overlays
- 2001-2003 – NetFS (NSF)
 - File system configuration of network properties
- 2002-2005 – X-Tend (NSF)
 - X-Bone for testbed uses
- 2003-2005 – DataRouter (int.)
 - Support for overlay P2P forwarding
- 2005-2006 – Agile Tunnels (NSA)
 - Partial overlays for DDOS safety
- 2006-2009 – RNA (NSF)
 - Extending X-Bone Choices model to general protocol stack architecture




Creating a Ring

Request



VI – definition

- Virtual Internet is network composed of:
 - Virt. hosts, virt. routers, virt. links (**tunnels**), i.e., an end-to-end system
 - provides at least the same services as IA
 - in a virtual context
- First-principles extension
 - More than a patch
 - More than interim





VIA Principles

- *TENET 1. Internet-like*
 - $VIs = VRs + VHs + \text{tunnels}$
 - Tunnels are links; separate net addresses
 - Emulating the Internet
- *TENET 2. All-Virtual*
 - decoupled from their base network
- *TENET 3. Recursion-as-router*
 - some of VRs are VI networks

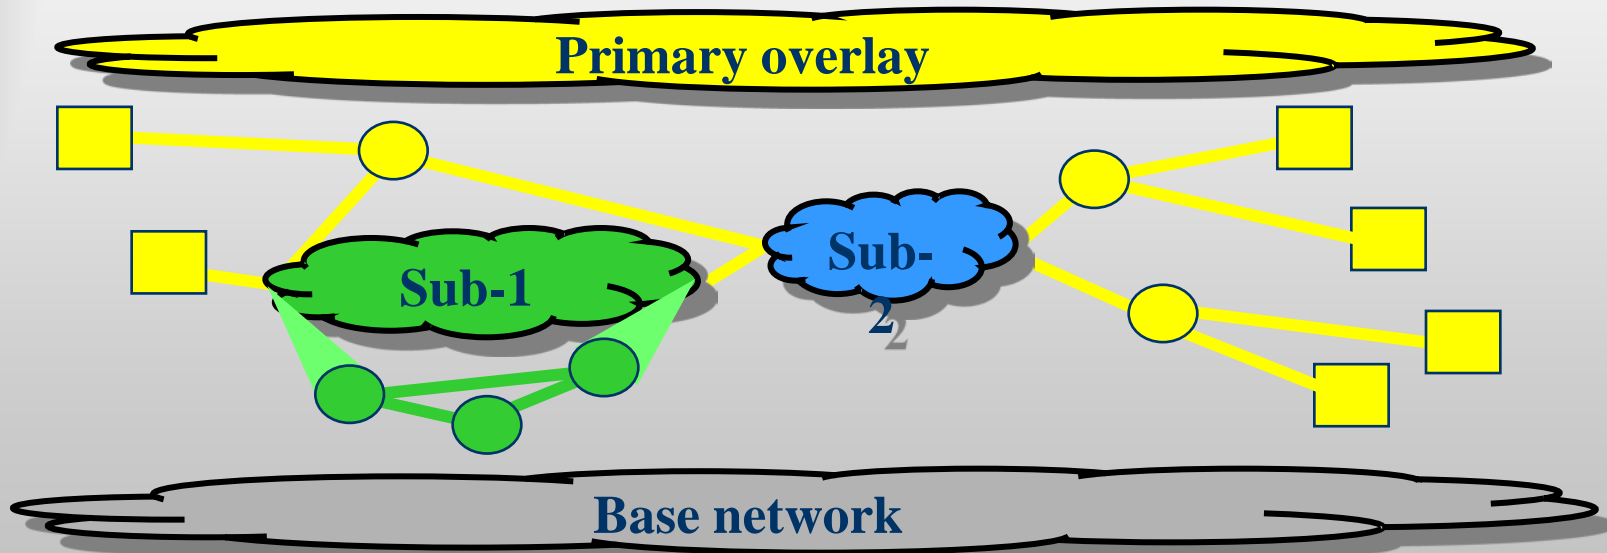


Extra constraints

- Internet-based
 - Routing (link up) vs. provisioning (link add)
 - *...one header to bind them all...*
(use IP, provide IP recursion)
- Complete E2E system
 - All VNs are E2E
- VN “Turing Test”
 - A net can’t tell it’s virtual
- Use existing protocols, OSs, apps.

Recursion-as-Router

- **Sub-overlays look like routers**
 - L3 version of *rbridges* (*IETF TRILL WG*)
 - *Similar to LISP/NERD/etc.*










Software Features

- Running code since 2000
 - FreeBSD port, Linux RPM, Cisco via buddy host
- Secure
 - TLS control plane, IPsec data plane
 - “Red teamed” software
- IPv4, IPv6 (both with IPsec)
 - Integrated with DNS updates, dynamic routing
- Per socket association to overlays
 - Allows process to bind to multiple overlays
- Application deployment
 - slice configuration, control distributed services
- Supports recursion, revisitation

Potential Uses

- **Test new protocols**
 - Test denial-of-service solutions
- **Deploy new services incrementally**
 - Dynamic routing, proxylets, security
- **Increase lab & testbed utility**
 - Overlapping nets, add delay & loss
- **Scale to very many nodes**
 - Simplify view of topology
- **Support fault tolerance**
 - Added level of recovery



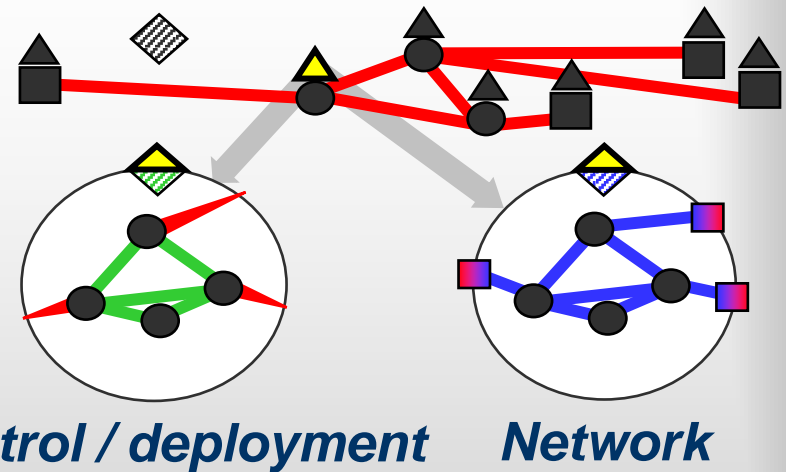


What We Don't Do...

- **Optimize the overlay topology**
 - we use a plug-in module (AI folk can provide)
 - it requires network status (not quite mature)
 - fault tolerance only via ground truth (future work)
 - X-Bone is capability more than performance (now)
- **Non-IP overlays**
 - Single, common interoperation layer
 - IP recurses / stacks nicely
 - "The lowest level at which experimentation is permitted is also the highest level at which experimentation occurs." – J.Touch, 1996


VI Features

- Recursion
 - Control (like BGP AS's)
 - Network (like LISP/NERD)
 - BARP (label distrib)
- Revisitation
- Integration of resolution, choices
 - Shims and glue layers as fundamental
- Service for deploying & managing VIs
 - Language for describing VIs

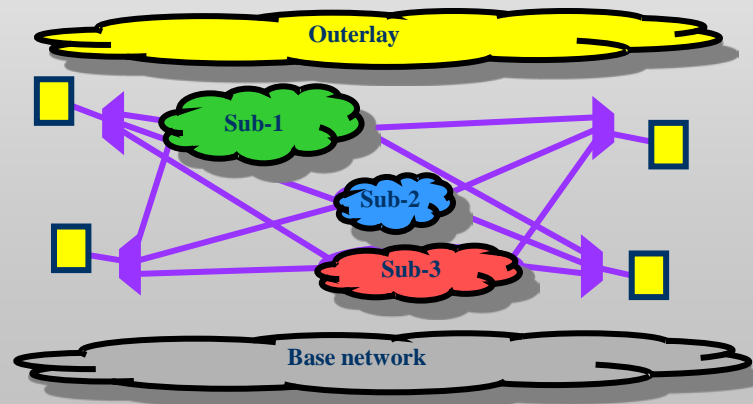


More VI Features...

- Compose:
 - DTN, Plutarch




- Alternate:
 - Control Plane, FEC, Boosters,
 - Dynabone





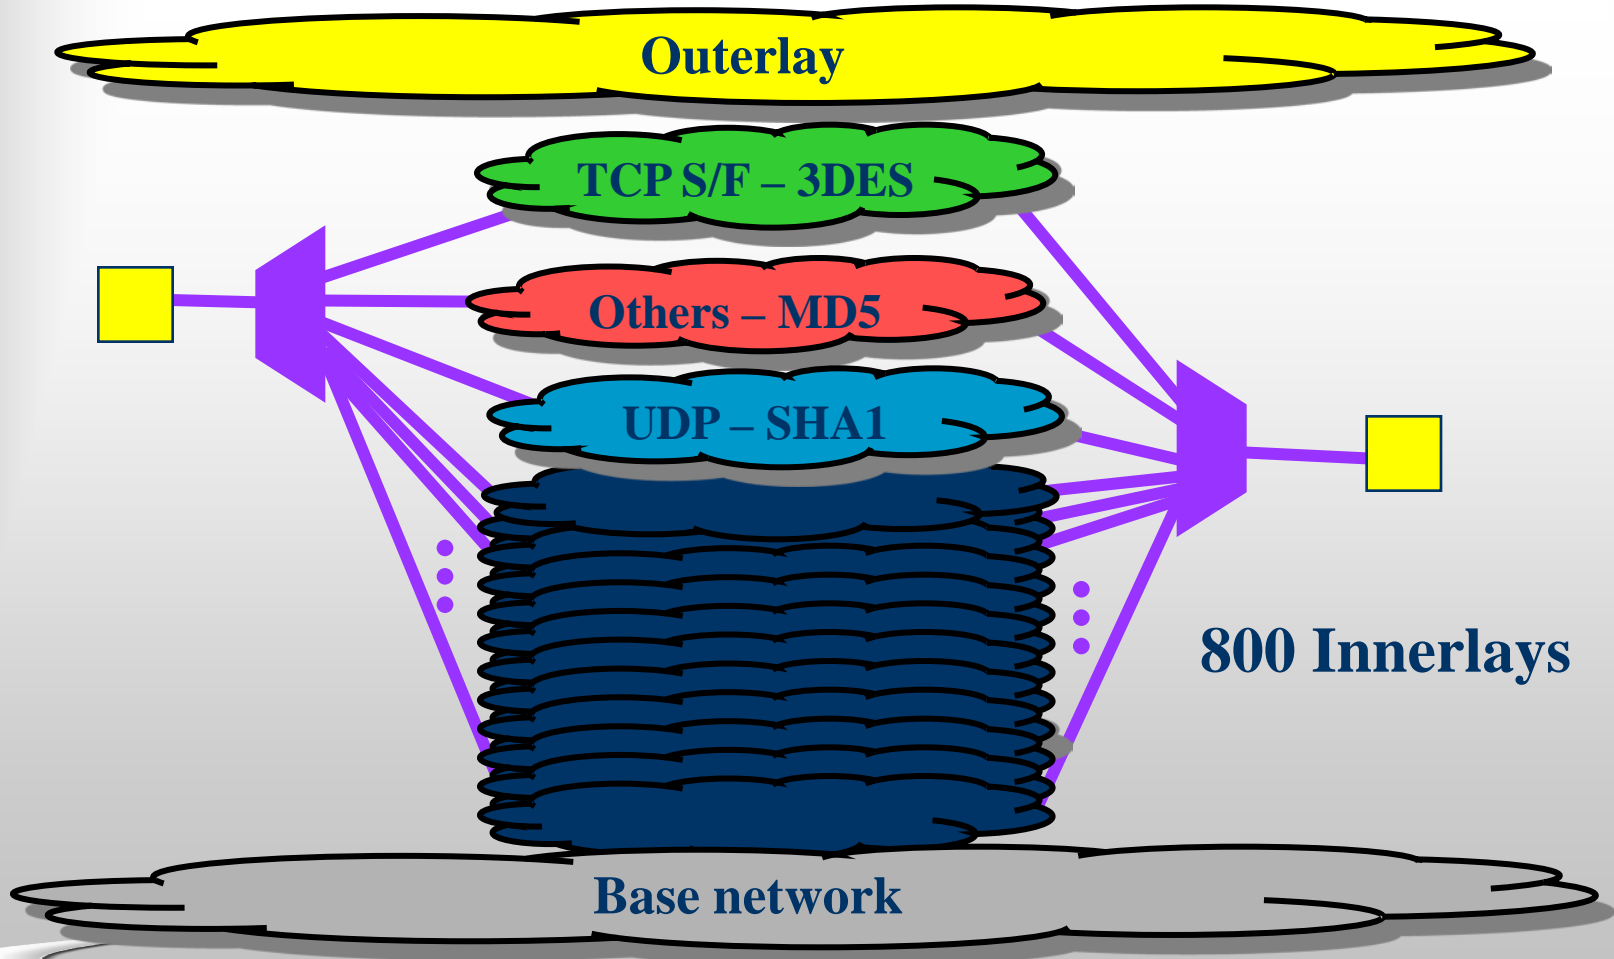
TetherNet

- Rents a block of addresses
 - Auto-configures secure tunnel
- Undoes effect of NAT/NAPT
 - Also effect of net non-neutrality






DynaBone Spread Spectrum



Agile Tunnel Protocol (ATP)

- Client
 - > **tunnel head @client**
 - > **roaming tunnel tail**
 - > server (hidden)

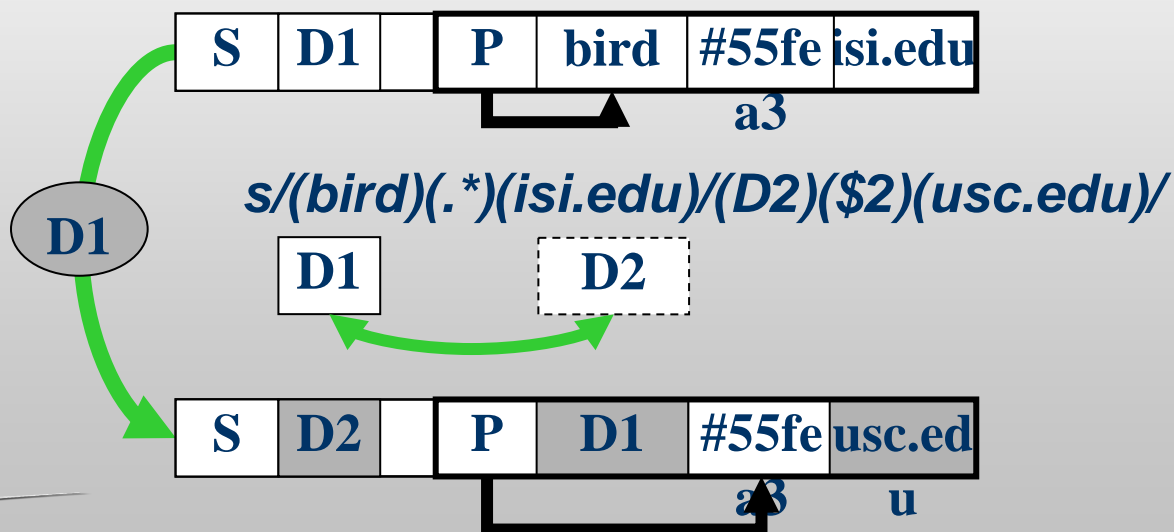
- Works like an inverse tunnel:





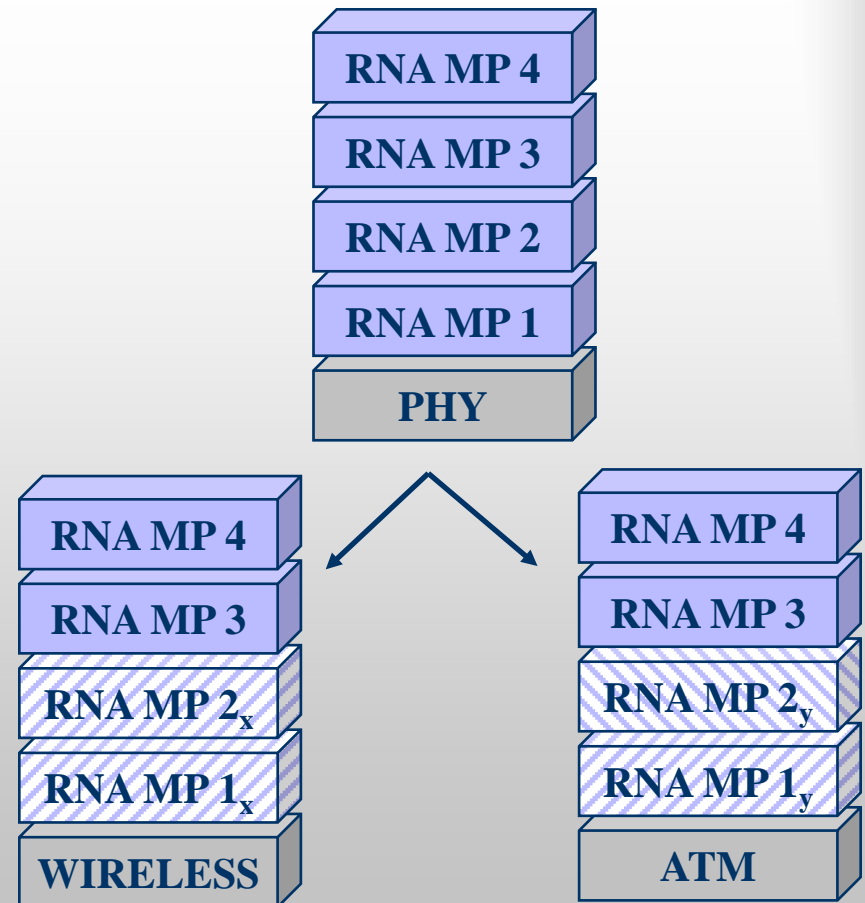
DataRouter for P2P

- P2P re-implements network arch.
- Need app.-layer forwarding at net layer
- Add string-based forwarding



Recursive Net Arch

- Layering as more than software engr.
 - Layers defined by scope, context
- Create a one layer 'stem cell' protocol
 - Integrate resolution, "choices" from X-Bone
 - Template of basic functions, ala J. Day





VI Observations

- Virtualization *changes* the architecture
 - Hosts are really processes, everything else is really a router or system
 - Devices aren't localized
 - Subnet as a router
 - NAT as a host front-end
 - Link and net layers are tightly coupled
- Core concepts from previous glue/shims
 - A single model yields layering, forwarding, routing, and dynamic composition



Current Efforts

- IRTF NetVirt BOF / NVRG mailing list
 - Met in Phila. March 2008
 - A- 'next gen' testbeds / GENI, 4WARD, AKARI, ...
 - B- 'next generation' research / FIND, FIRE, ...
 - C- virt extensions / PPVPN, L*VPN, X-Bone, LISP/NERD, PWE
- Future Internet meetings
 - ICCCN 2008 "FIAPP" (future Internet arch and protocols)
 - CoNext 2008 "ReArch" (re-architecting the Internet)
 - ICCCN 2009 "NAP" (net arch & protocols)



US-EU Collaboration

- UPC / O. Ardaiz Villanueva (Spain)
 - Visiting researcher, overlay service deployment
- UCL MICE / Peter Kirstein (UK)
 - X-Bone collab. within DARPA Active Nets
 - Deploy a MICE teleconferencing overlay
- EU 4-WARD / Martina Zitterbart (Germany)
 - RNA collaboration with Univ. Karlsruhe, NC State NetSilo / Rouskas
 - Unifying approaches to protocol layering



US-Japan Collaboration

- WIDE / Jun-ichiro (Itojun) Hagano
 - X-Bone, FreeBSD virtualization support
- WIDE / Hideaki Imaizumi
 - Dynabone, Visiting researcher, IPv6 transition mech.
- WIDE / Shigeya Suzuki
 - X-Tend, Visiting researcher, RF tags
- NEC / Norihito Fujita
 - X-Tend, Visting researcher, P2P network overlays



Project Issues

- Coordinating projects
 - Separate preconditions
 - Separate goals (esp. funded independently)
 - Separate funding trajectories
 - *very* hard to coordinate a-priori
- Support for travel
 - Meetings are more international anyway
 - Direct support for guest presentations



People Issues

- In-person visits
 - Visiting researchers
 - Guest presentations
- Regular meetings
 - "T-Troup" participation, group talks & paper reviews
 - External review collaboration
- Social support
 - Lunches, coffees, Friday bagels
 - Off-time - "WoW"s, holidays, sharing local favorites



Things that help...

- Support / piggyback on other meetings
 - Sigcomm 2007, Kyoto (Jun Murai/Keio Univ., chair)
 - CoNext Research 2008, Madrid
- Skype + Powerpoint
 - Cheaper than telephone
 - Easier than teleconferencing
- Be conscious of language issues
 - Alternate descriptions
 - Simpler description/words...
without simplifying the ideas