The X-Bone & its Virtual Internet Architecture 10 Years Later

Joe Touch, Greg Finn, Lars Eggert, Amy Hughes and Yu-Shun Wang

Workshop on Overlay and Network Virtualization
Kassel, Germany
March 6, 2009

Talk Outline

history

Virtual Internets

- why
- what
- architecture highlights

related projects at ISI (time permitting…)

- X-Bone, DynaBone, TetherNet

History

X-Bone was a series of research projects at USC/ISI

- X-Bone, DynaBone, TetherNet, X-Tend, NetFS, GeoNet, …
- 1997-2005+
- initial funding from DARPA, follow-on funding from the NSF
- http://www.isi.edu/xbone/

key results

- an architecture (the “Virtual Internet” architecture)
- a deployment/management system (the “X-Bone”)
- follow-on work using virtual nets:
  - DynaBone — spread-spectrum virtual networks
  - TetherNet — rent real Internet behind firewall + NAT
  - GeoNet — geographically-routed virtual networks

Prior & Related Work

new services & protocols

- Cronus, M/6/Q/A-Bone

multi/other layers

- Cronus, Supranet, MorphNet, VANs

partial solutions

- VPN, VNS, RON, Detour, PPVPN, SOS

virtualization, revisitation, recursion

- X-Bone, Spawning, Netlab/Emulab

OS virtualization

- VMware, jails, vserver, XEN, PlanetLab

Virtual Internet – Why

“network equivalent of virtual memory”

protection

- separate topology, optionally secured
- test + deploy new protocol/service

sharing

- increase utility of infrastructure

abstraction

- adapt topology to application

Virtual Internet – What

network = hosts + routers + links

virtual network =

	virtual host	→ packet src/sink
+	virtual router	→ packet gateway
+	virtual link	→ tunnel *X* over *Y*

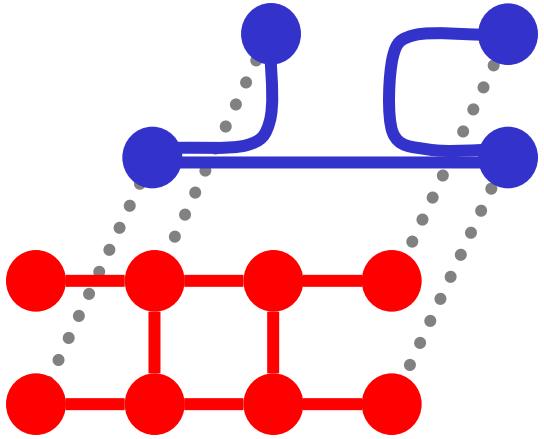

virtual Internet – ”network of networks”

- use Internet as physical media
- create virtual link & network layers
- strong L2 vs. weak L3 host model

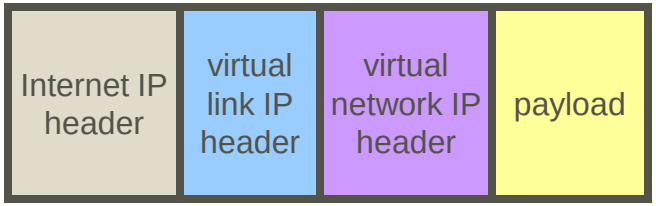

a virtual Internet should look exactly like the real thing

- “if an app can know it runs in a VI, we did it wrong

VI Architecture Feature – Recursion

virtual Internets on top of virtual Internets

our litmus test:

- system should be able to do recursive VI-in-VI without hacks

recursion has real uses cases

e.g., allows transparent reconfiguration

- change outer VI w/o affecting inner
- fault tolerance, basis for DynaBone

also allows VI “embedding”

- “router is a network inside”

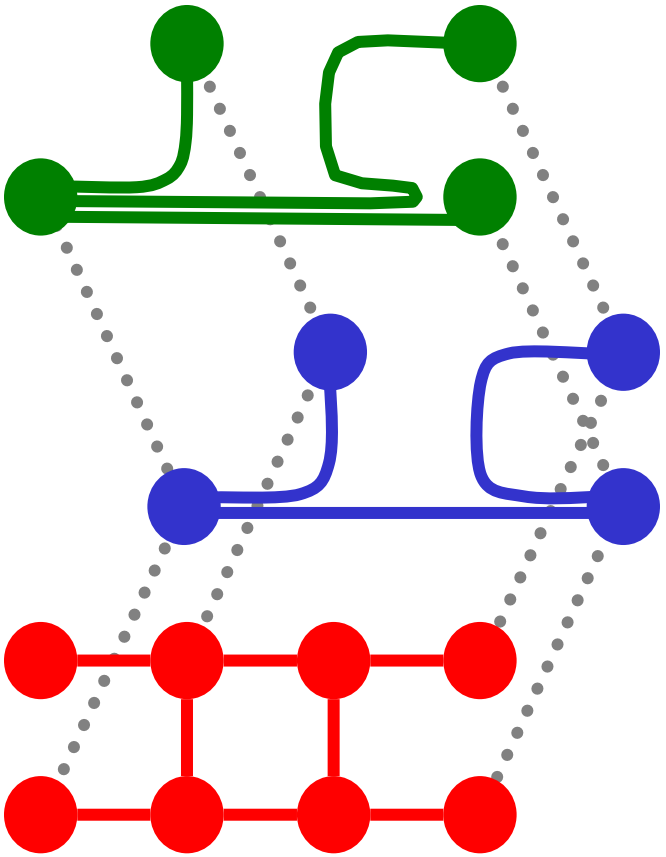

VI Architecture Feature – Concurrency

one node participates in multiple virtual Internets at the same time

basis for isolation & abstraction

bind different apps/VMs to different VIs on the same physical node

VI Architecture Feature – Revisitation

one node participates in the same virtual Internet but multiple times

allows creation of VIs larger than physical resources

fully decouples virtual from physical topologies

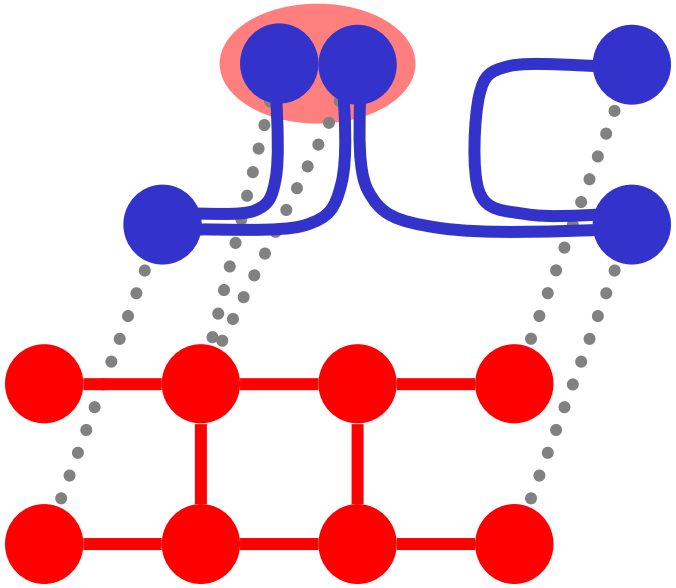

VI Architecture Feature – Hop-by-Hop Security

security in the Virtual Internet architecture is a virtual link property

- decoupled from topology
- transparently coexists with end-to-end security inside the VI
- transparently coexists with security underneath a VI

IPsec tunnel mode

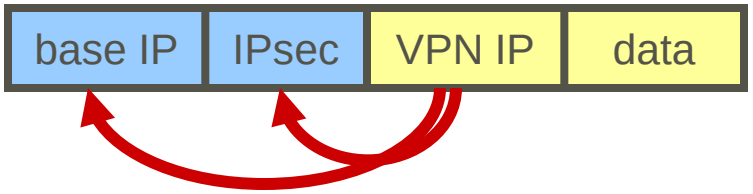

IPIP tunnel + IPsec transport mode

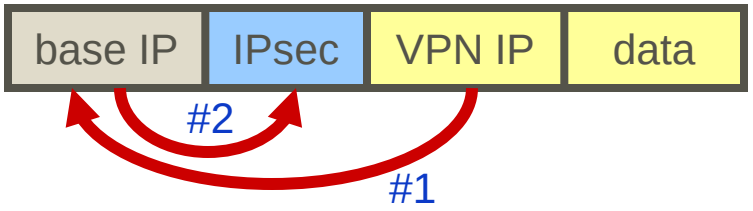

IPIP tunnels + IPsec transport mode

- modular tunnel mode equivalent
- huge IETF debate around 2000 (draft-touch-ipsec-vpn-05.txt)

The X-Bone System

deployment + management system for virtual Internets

- programs → standardized API
- humans → web interface

high-level virtual network description language

- express virtual topology + services
- XML

collaborating, distributed management daemons

- multicast expanding-ring discovery
- distributed resource reservation
- instantiate + manage virtual network

non-goals: topology optimization, non-IP VIs, …

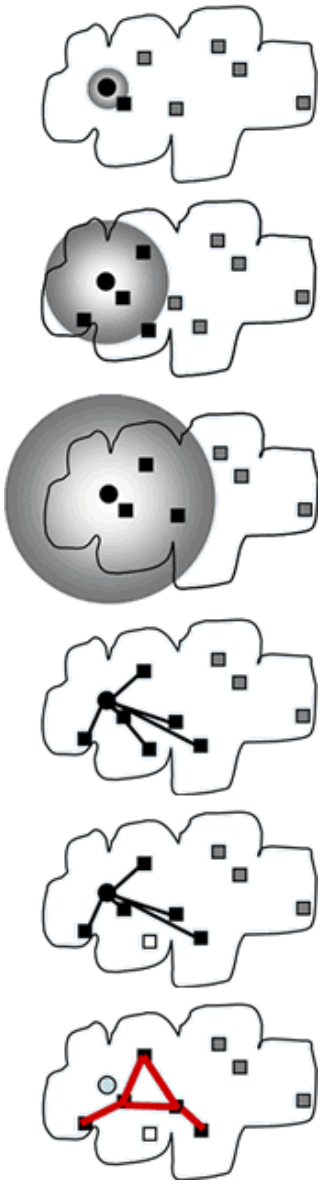

X-Bone Screenshots

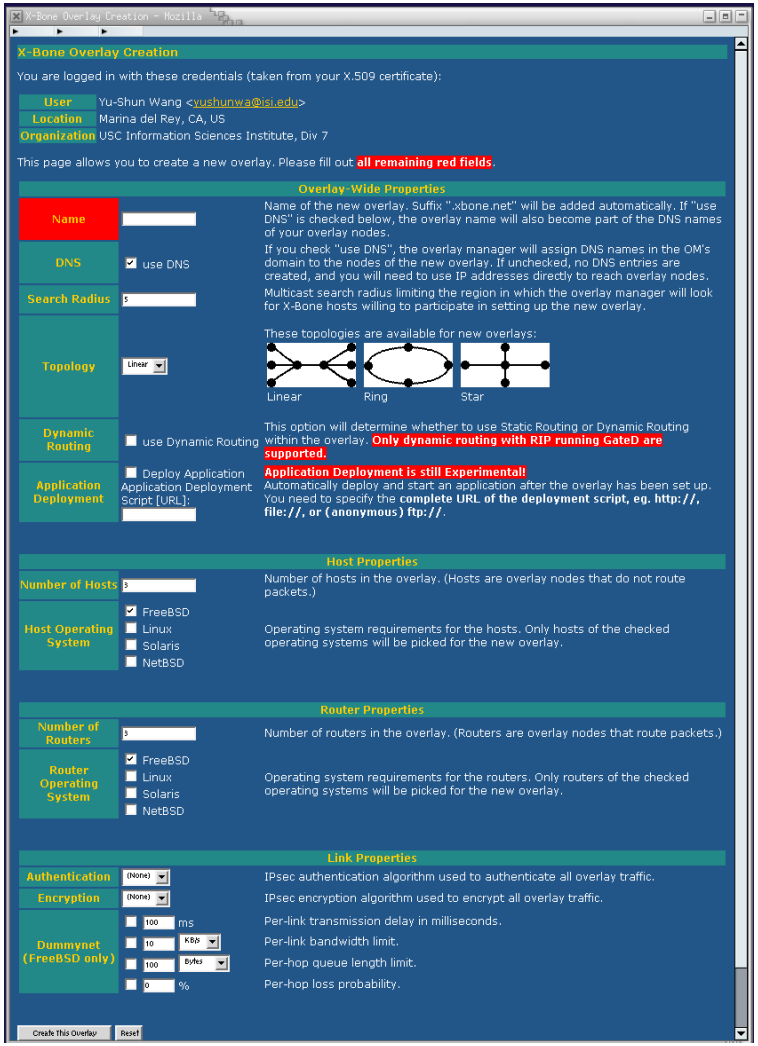

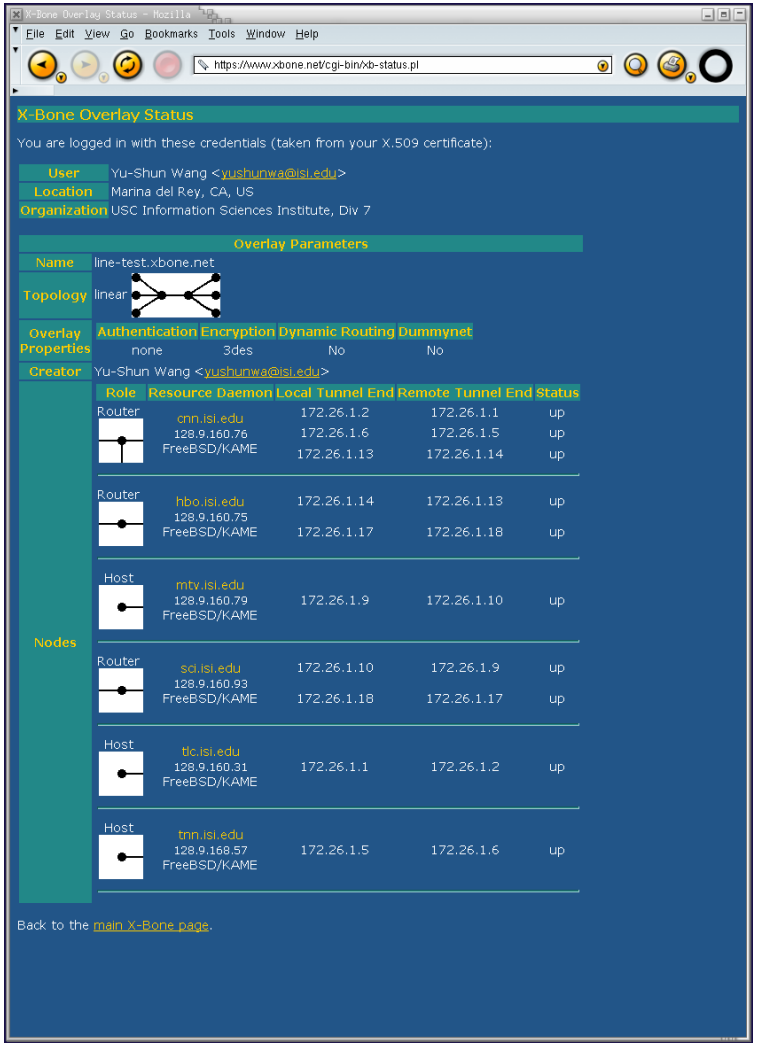

X-Bone Status

current release: 3.2

- mature: 10 years of open source availability

platforms: FreeBSD, Linux

- unofficial: NetBSD, Cisco

widely used (by 2003):

- UCL, UPenn, Aerospace, DOD Canada, Sinica Taiwan + more

Related Work at USC/ISI

DynaBone

parallel inner virtual networks = algorithmic & protocol diversity

spread-spectrum multiplexer, wrapped inside outer virtual network

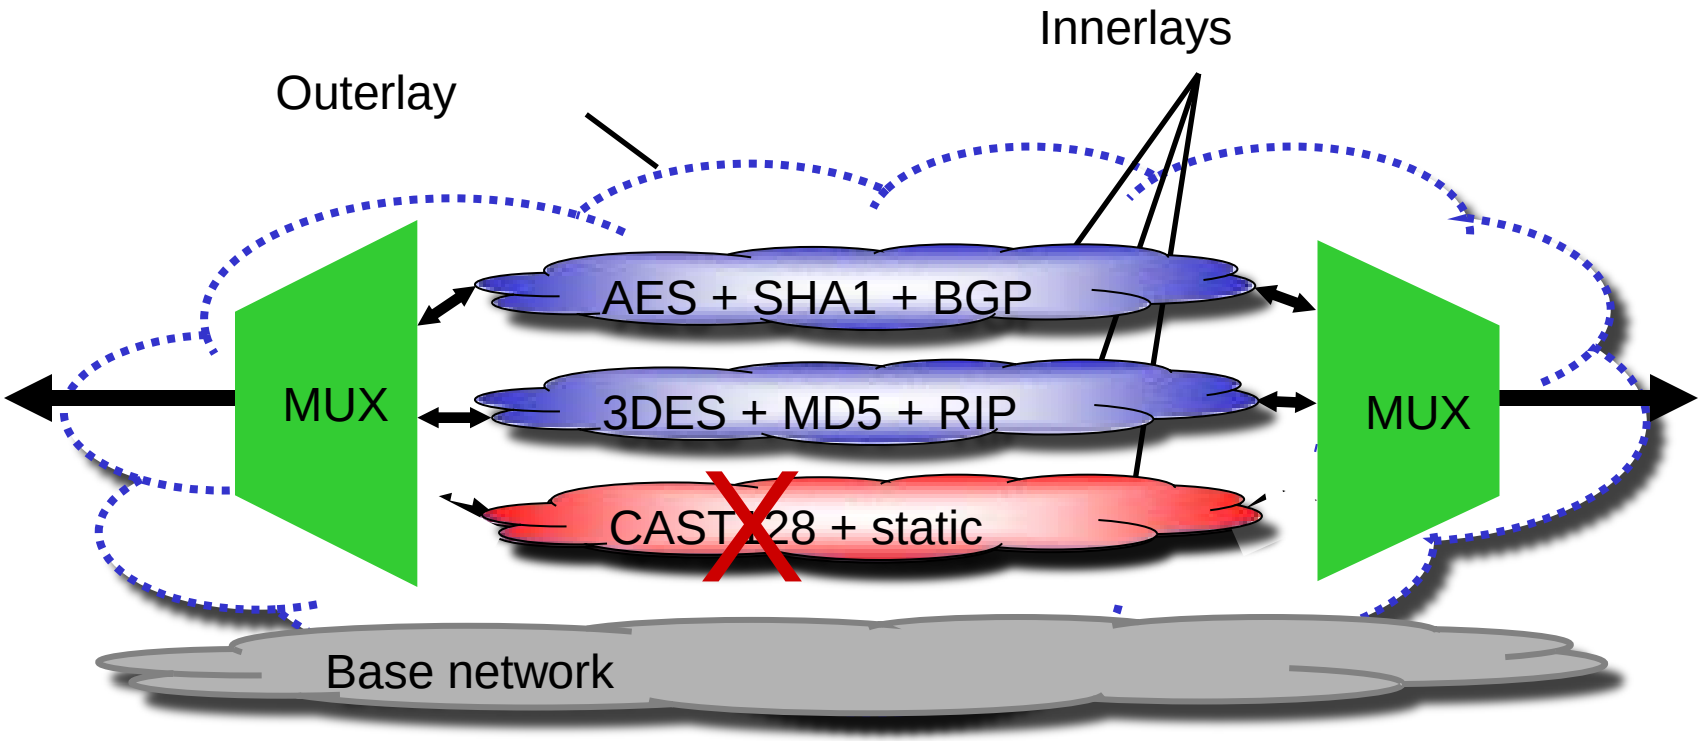

TetherNet

issue: firewalls, NATs, clueless ISPs

broken end-to-end connectivity

solution: relocate real Internet subnet

real = routable IP + DNS + no fw + …

tunnel subnet from anchor router to tether router at remote site

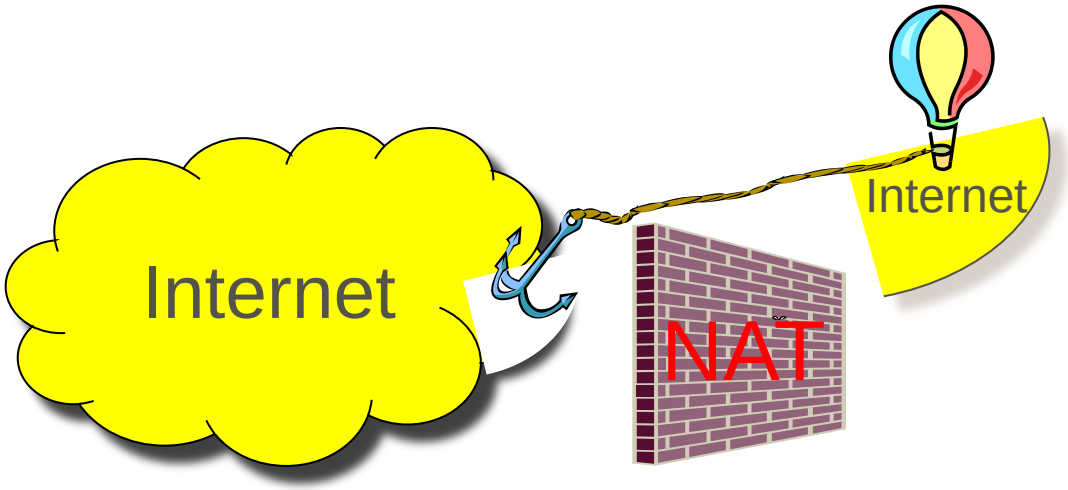

TetherNet Features

true Internet behind NATs and firewalls

- IPv4 + IPv6
- multicast
- fwd/rev DNS
- traffic shaping
- 802.11b AP
- secure: IPsec for traffic, X.509 for user auth
- web interface configuration

U.S. patent filed, talks with licensees

TetherNet Screenshots

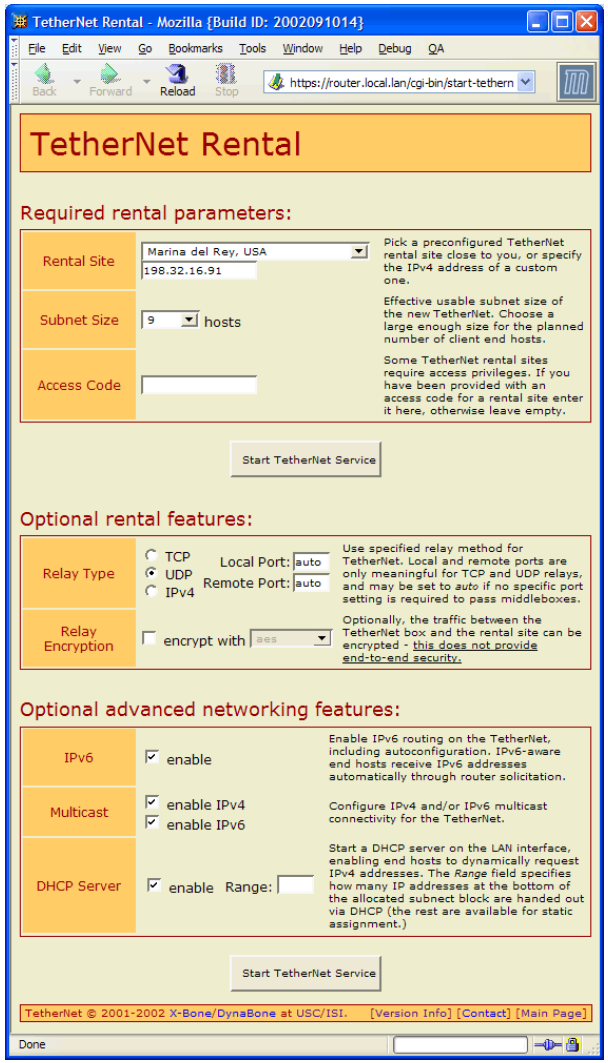

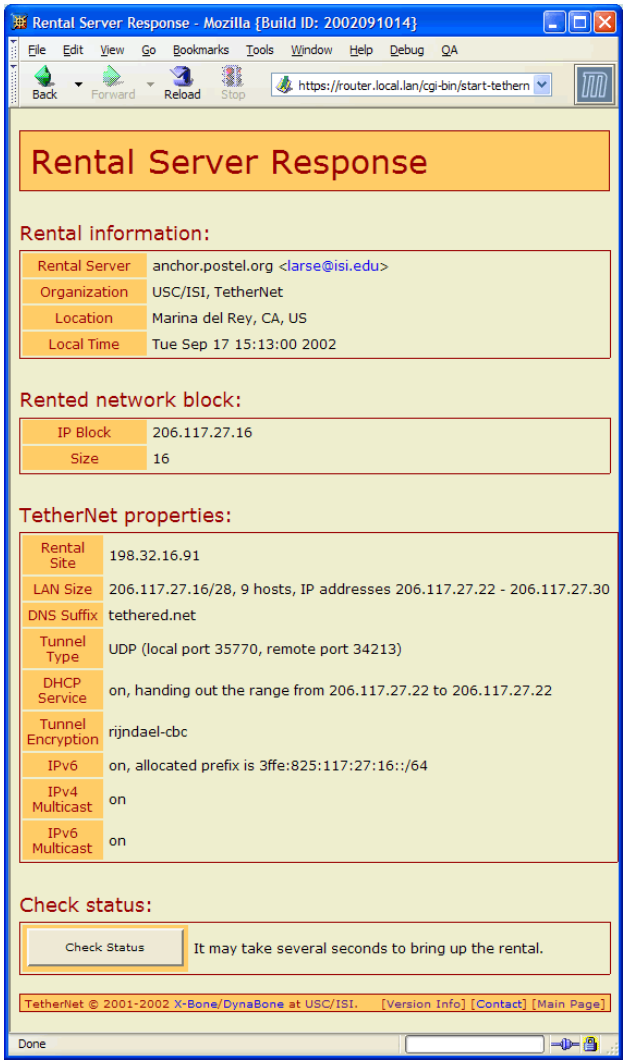

Other Projects

X-Tend

maintain + extend X-Bone as tool for research + education

GeoNet

geographically-addressed overlays

NetFS

access control for the network stack via a pseudo file system

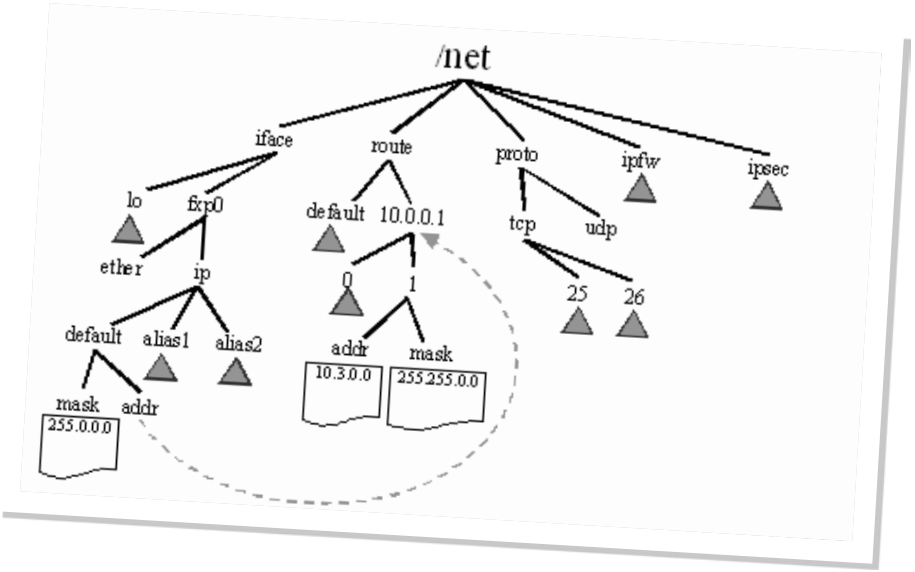

THANK YOU!