

# The X-Bone

Joe Touch / touch@isi.edu

USC / Information Sciences Institute

NGI Workshop White Paper - March 1997

## Introduction

Experimental networks are increasingly being implemented using overlay networks. Overlay networks are protocol or application specific subnets connected by IP encapsulation tunnels over production networks. Overlays are configured and managed manually and independently of each other and the underlying infrastructure. Recent overlays support new protocols such as the M-Bone for multicast IP [5], and the 6-Bone for IPv6 [1], application networks such as the distributed Harvest cache system [3], and experiment networks such as for the Digital Libraries Initiative [4].

The X-Bone is a system for the rapid deployment and management of multiple overlay networks. The X-Bone's coordinated management supports both long-lived and short-lived large scale overlay networks with little manual intervention. The X-Bone provides a configurable virtual networking infrastructure, critical to the development of both network and application services, and is equally useful for deploying isolated infrastructure for restricted purposes.

## The X-Bone

The X-Bone consists of existing overlay technology combined with teleconferencing-style management tools and advanced coordination services. Current overlay networks are composed of per-overlay routing daemon software installed at selected sites, connected by IP-encapsulation tunnels [8]. Application overlay networks use existing networking infrastructure together with application-layer 'routing', *e.g.*, to interconnect proxies of a cache distribution system as in Harvest [3]. Interference between overlays is currently managed by system administrators.

The X-Bone uses an interface and resource discovery protocol adapted from the sd and sdr M-Bone tools [5]. Sd and sdr are session directory tools that advertise multiple M-Bone multicast sessions on a single, global multicast channel. A similar X-Bone directory tool, 'xd,' provides an equivalent interface to multiple X-Bone overlay networks.

Whereas sd/sdr support users joining audio/video M-Bone sessions, xd supports local daemons joining X-Bone overlay sessions by creating IP-encapsulation tunnels between the corresponding daemons. When a site joins an overlay, its xd tool requests an authenticated copy of the appropriate daemon software and configuration information. These overlay networks can be composed from host-based routers initially, and use Active Networks (AN) to load code at dedicated routers, when available [7]. Xd provides a uses ANs programmability, and supplies a tool to manage distributed coordination and code deployment.

Xd also coordinates the sharing of resources at the local site, and provides a single interface to RSVP, to allow configuration of link reservations that span an entire overlay. This allows

xd to support resource reservation (e.g., RSVP [10]) without requiring implementation in each overlay's daemon. Xd also provides a user interface and API, which allows external programs or users to manually override and parameterize the overlays.

## **Benefits of the X-Bone**

The X-Bone provides isolation between overlay networks. It provides a partitioning of resources that allows experimental overlay networks to avoid interference with production services, and allows overlay testbeds to guarantee their own service for dedicated experiments. In this way, overlays can provide per-testbed virtual networks using both public production services and dedicated resources, such as the bandwidth resources of the vBNS backbone [9].

The xd tool provides coordination that allows overlay networks to be automatically deployed anywhere IP multicast can reach. Network managers can perform manual adjustment of this automatic infrastructure using its remote tunneling services. By encapsulating resource partitioning in the xd tool, authors of per-overlay daemons avoid replication of effort interfacing to reservation services.

For example, A TCP-NG-Bone can be deployed to test new TCP-NG capabilities without disrupting existing TCP. An emergency-services overlay backbone can be deployed, in which capacity is reserved on the tunnel links to ensure traffic priority. A military brigade can deploy a backbone for temporary private network service using existing encrypted link technology.

## **Research Issues**

Developing the X-Bone in the IPv4 (current) Internet would be challenging; however, by raising the basic services required for participation in the X-Bone to include multicast IP, first-phase X-Bone services can be deployed in the very short term (6 mos.). This would provide automatic tunnel connectivity of daemons that are deployed via script control of a network manager, using out-of-band mechanisms or emerging tunnel management protocols [6].

Second phase X-Bone tools (second year) will be interfaced to emerging distributed security infrastructure and resource reservation services. This phase provides resource partitioning and automated daemon deployment.

Third phase X-Bone (third year) will address advanced research issues, such as support for native infrastructure in addition to IP-encapsulation overlays, daemons that participate in multiple overlays (inter-overlay routers), and intelligent dynamic coordination of resource partitioning to optimize efficiency (rather than merely overbooking). These exploit the results of other research and development efforts, such as Argonne's MorphNet [2], which would provide reservation services to native infrastructure at the link and physical layers.

We believe the X-Bone offers the existing Internet infrastructure a path to providing rapidly deployable virtual infrastructure with the same ubiquity as current services enjoy. This infrastructure is critical to the effective development of new services that will bring the Internet into the 21st century.

## References

- [1] "6-Bone" Web home page, <http://www-6bone.lbl.gov/6bone/>
- [2] Aiken, R., Carlson, R., Foster, I., Kuhfuss, T., Stevens, R., and Winkler, L, "Architecture of the Multi-Modal Organizational Research and Production Heterogeneous Network (MORPHnet)" Tech. Report ANL-97/1, Argonne National Lab, IL., Jan. 1997.
- [3] Bowman, C.M., Danzig, P., Hardy, D., Manber, U., and Schwartz, M. "The Harvest Information Discovery and Access System," *Computer Networks and ISDN Systems* 28, 1995, pp. 119-125.
- [4] "Digital Libraries" Web home page, <http://www.dlib.org/>
- [5] Eriksson, H., "MBone: The Multicast Backbone," *Communications of the ACM*, Aug. 1994, Vol.37, pp.54-60.
- [6] Hamzeh, K., "Ascend Tunnel Management Protocol - ATMP," Internet RFC 2107, Ascend Communications, Feb. 1997.
- [7] Tennenhouse, D.L., Smith, J.M., Sincoskie, W.D., Wetherall, D.J., and Minden, G.J., "A survey of active network research," *IEEE Communications Magazine*, Jan. 1997, pp. 80-86.
- [8] Simpson, W., "IP in IP Tunneling," Internet RFC 1853, Daydreamer, Oct. 1995.
- [9] "vBNS" Web home page, <http://www.vbns.net/>
- [10] Zhang, L., Deering, S., Estrin, D., Shenker, S., and Zappala, D., "RSVP: A New Resource ReSerVation Protocol," *IEEE Network*, Sept. 1993.