Unmanaged Internet Protocol
Scalable Application-Friendly Internet Routing

Bryan Ford
Massachusetts Institute of Technology

August 29, 2003

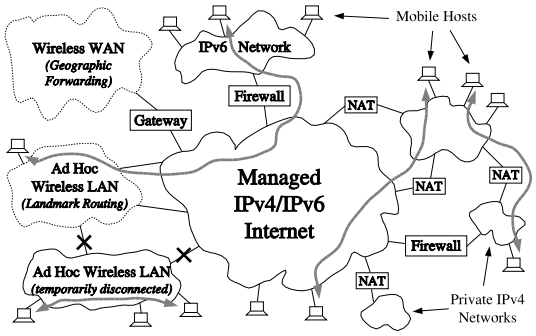

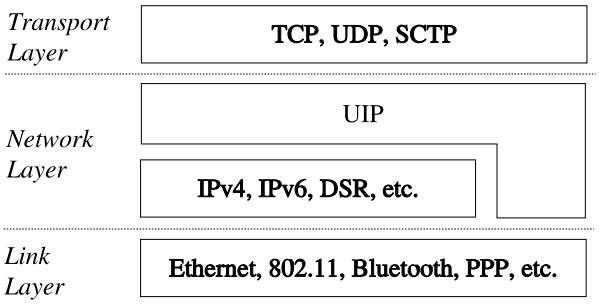

1 Introduction

Unmanaged Internet Protocol (UIP) re-implements the original ARPAnet vision of scalable, robust, any-to-any communication between all participating hosts, stitching together the currently discontinuous and fragmented Internet into a single logical network with uniform addressing and universal connectivity (above left). UIP provides decentralized, management-free routing among hosts on the existing IPv4 and IPv6 Internets, hosts on private networks behind firewalls and network address translators, mobile hosts with ephemeral IP addresses, and hosts on experimental or ad-hoc wireless networks with no IP-based connectivity. Applications address other UIP nodes using self-certifying cryptographic identities that can be created by anyone, require no centralized administration, and remain valid as long as desired even as nodes move [3]. All communication between UIP nodes is integrity- and privacy-protected by default, making application-layer security protocols such as TLS [1] unnecessary.

2 Routing

UIP acts as a new network sublayer on top of IP (above right), providing connectivity based on cryptographic identities while hiding topology-sensitive IP addresses from applications. UIP appears to IP as a new upper-level protocol, and to transport protocols and applications as a new address/protocol family. UIP can also run directly on top of link-layer protocols or other non-IP-based network layers such as ad-hoc wireless routing protocols.

UIP takes advantage of underlying link- and network-layer protocols for efficient packet forwarding whenever possible, but when underlying protocols fail to provide direct connectivity due to network address translation (NAT), incompatible address formats or routing technologies, or other transient or persistent failures, UIP uses its own management-free routing and forwarding mechanism to route around these failures automatically. UIP uses a scalable, self-organizing distributed hash table (DHT) similar to the one used by the Kademlia peer-to-peer system [2], both for maintaining strong network connectivity, for locating nodes based on their cryptographic identities, and for discovering efficient forwarding paths to those nodes when necessary.

References

[1] T. Dierks and C. Allen. The TLS protocol version 1.0, January 1999. RFC 2246.

[2] Petar Maymounkov and David Mazières. Kademlia: A peer-to-peer information system based on the XOR metric. In *Proceedings of the 1st International Workshop on Peer-to-Peer Systems*, March 2002.

[3] R. Moskowitz and P. Nikander. Host identity protocol architecture, April 2003. Internet-Draft (Work in Progress).