Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

SOME OBSERVATIONS ON THE ATM ADAPTATION LAYER SERVICES NEEDED TO SUPPORT TCP/IP.

Grenville Armitage
(gja@mullian.ee.mu.OZ.AU)

Keith Adams
(kma@mullian.ee.mu.OZ.AU)

SUMMARY

This paper examines four issues:

1. The datagram transport needs of the DARPA Internetworking Protocol,

2. The basic data transport facilities provided by CCITT's evolving ATM Adaptation Layer,

3. The relationship between these facilities and the needs of IP,

4. Encapsulation of IP datagrams, and the provision of dynamic Address Resolution.

Comparisons are also made with LAN support of IP services.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

INTRODUCTION

DARPA Internetworking Protocol (IP) forms the basic networking support structure for many academic and research communities.

- ‘Main frames’ use it,
- Standalone workstations use it,
- PC’s and clones use it,
- Runs over many varied physical media.

IP is not UNIX specific.

OSI may be on the way in,

but IP is here and now.

Giving commercial and government IP service providers the ability to use B-ISDN as an IP transport medium increases B-ISDN market.

This paper should give a brief overview of IP, and what it may demand of B-ISDN.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

THE IP WORLD

IP provides a shielding layer between physical transport media and higher reliable transport layers and application programs.

IP provides a virtual connectionless datagram network to higher layers.

- Non-hierarchical.
- Consistent address format.
- All IP interfaces appear ‘directly’ addressable.
- IP transfers data in variable sized ‘chunks’, or datagrams.
- Delivery of datagram to any given IP destination, at any given time, is not guaranteed.

The IP layer incorporates routing algorithms to correctly pass datagrams across different physical layers to reach the destination IP address.

To provide reliable data transport services the Transport Control Protocol (TCP) is layered on top of IP - this is beyond the scope of this presentation.

TWO VIEWS OF THE IP WORLD

To any one interface, all other interfaces are directly addressable.

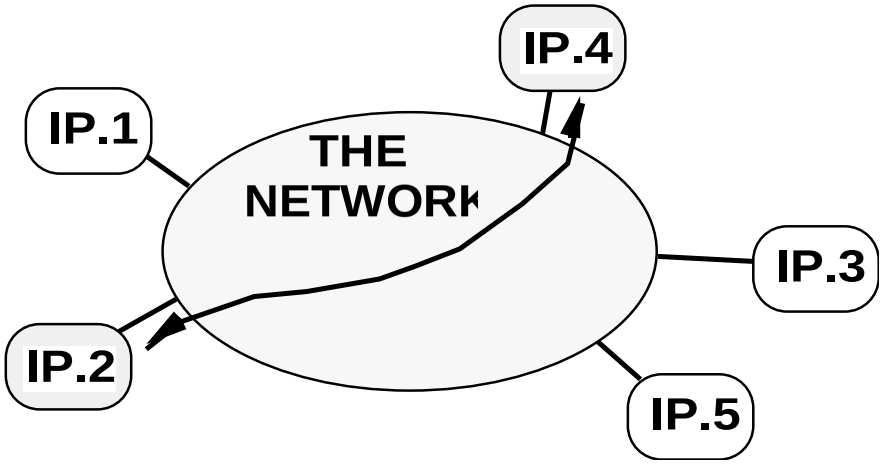

Inside the virtual network, physical networks are being traversed.

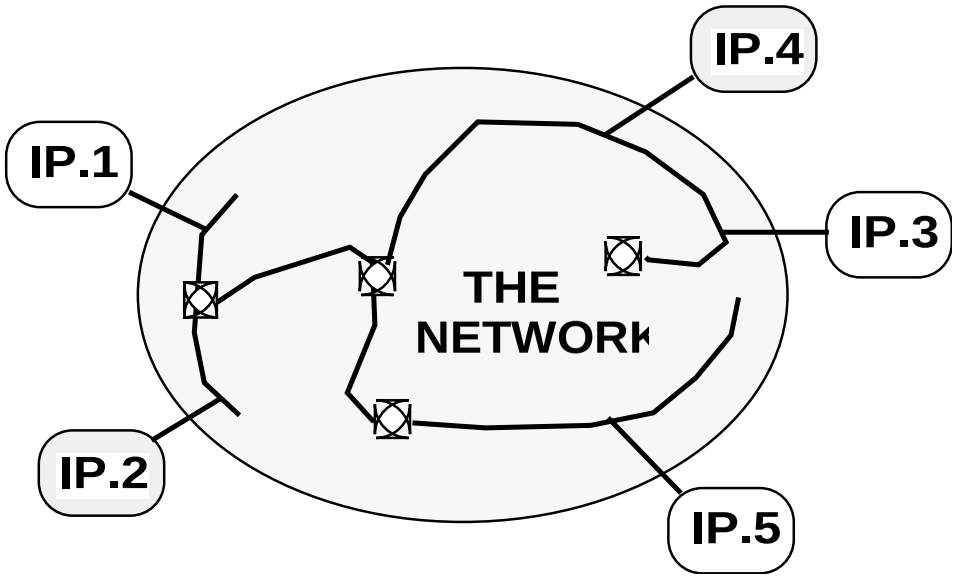

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

IP IN AUSTRALIA

Australian Academic Research Network (AARNet) is biggest IP provider in Australia.

Dedicated links between capitols feed and interlink state hubs.

State hubs feed academic and research institutions within each state.

Institutions feed individual departments using LANs.

Departments feed workstations using LANs or low speed lines.

Can a country wide B-ISDN service fit into this scene at any useful level?

- Start by serving the inter city links.

- Evolve towards B-ISDN wall-sockets replacing departmental LANs ?

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

PHYSICAL LAYERS

These provide the actual mechanism for transporting IP datagrams.

An upper layer sends to the IP layer

- A block of data
- A header containing the destination IP address and the source IP address.

IP layer needs to locate a physical layer destination address that corresponds to the destination IP address.

IP layer then sends to the physical layer

- The IP datagram
- A header containing the physical layer destination and source addresses.

Physical layer must transport IP datagrams with boundaries intact.

IP Routers/Gateways

- Nodes with multiple physical layer interfaces.

- Pass IP datagrams received on one interface to an interface ‘closer’ to the IP destination.

- May be dedicated machines, or part of user workstations.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

LOCAL AREA NETWORKS

Frequently used to support localised IP services within organisations or buildings. Inherently preserve boundaries of the datagram being transmitted.

Popular example is Ethernet (and close relative, IEEE 802.3).

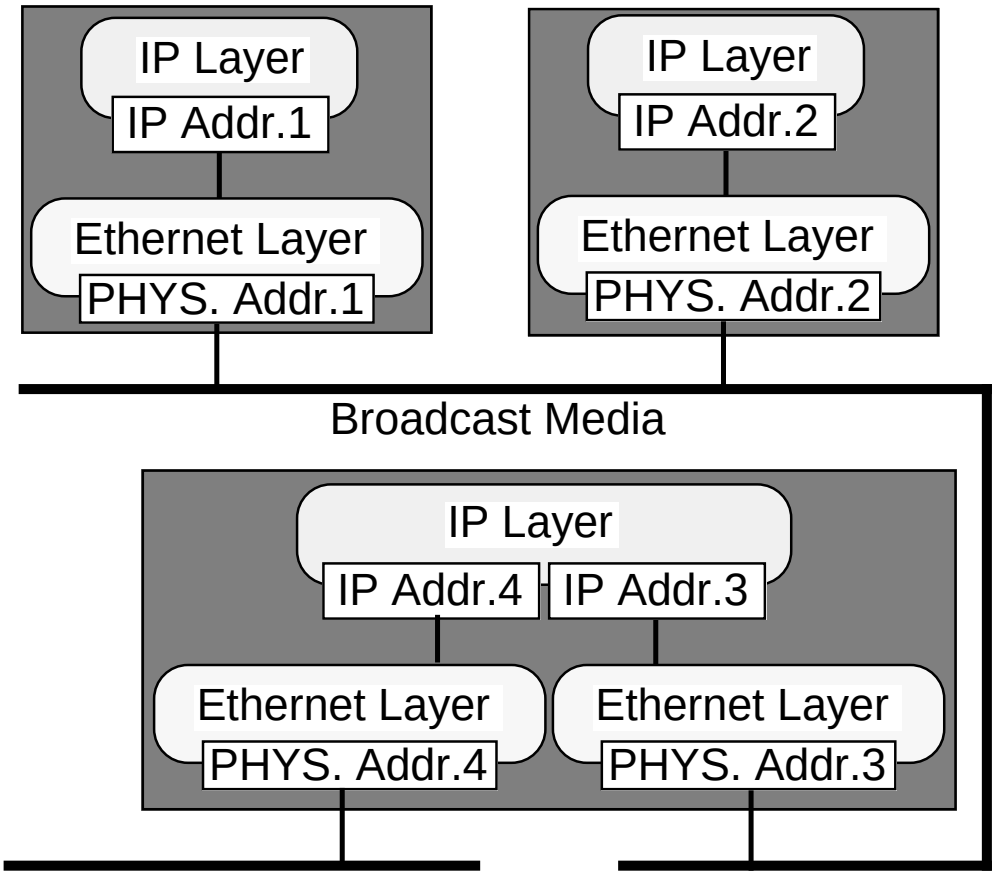

The immediate Ethernet destination is either the address of the destination IP entity, or that of an IP Router believed to be ‘closer’ to the destination IP entity.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

POINT TO POINT LINKS

Covers a variety of schemes

- Dialup modem connections,
- Dedicated serial lines,
- Permanent or semi-permanent virtual circuits.

Each link has only 2 ends. Each end may or may not have an associated Physical Layer ‘address’.

IP layer generally does not need to establish the physical layer identity of the destination on a per-datagram basis.

Byte or bit orientated schemes utilise encapsulation to preserve IP datagram boundaries.

- Serial Line IP (SLIP)
- Point to Point Protocol (PPP)
    - PPP also provides for multi-protocol support over a point to point link.

Virtual circuit services may by treated in a similar fashion, although it is possible to create schemes where the point to point circuit is established ‘on demand’ and removed when idle.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

SMDS
(and similar schemes)

Wide area LAN-ish service.

Transports packets or ‘frames’ in a connectionless manner.

IP entities will exist ‘on top’ of a given SMDS interface.

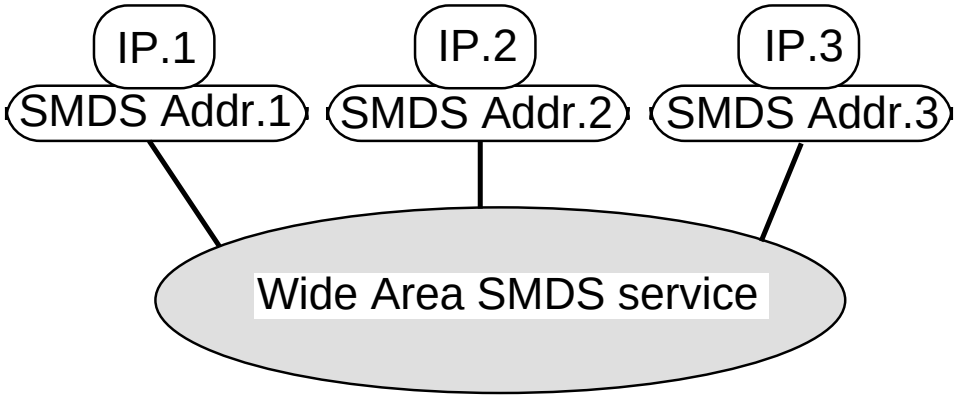

Support for multicast groups.

Physical layer broadcasts cannot be supported

- ‘Blind’ Address Resolution not easy,
- SMDS interface addresses need to be pre-configured.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

THE 3 IP ISSUES

Providing the IP ‘virtual network’ using different physical media boils down to three major issues

- Encapsulation of IP datagrams,

- Routing IP datagrams to the physical network on which the destination resides,

- Resolving IP addresses to physical layer addresses.

IP addressing scheme assumes a broadly hierarchical interconnection of physical layers into networks, sub-networks, and hosts.

Partial routing knowledge is extracted from IP addresses.

So what schemes are available under B-ISDN to solve these problems?

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

ADAPTATION LAYER SERVICES

Datagram transport has been “designed-in”

- Type 3 and 4 SAR/CS service
- Type 5 (?)
- Connectionless transport (layered over AAL4).

Transport links can be

- Virtual connections (Point to Point)
- Multicast (Point to multi-point)
- Connectionless (I.cls)

How much encapsulation do you really need?

- Simple is beautiful, need almost nothing.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

CONNECTIONLESS MODE

Relevant characteristics

- Preserves datagram boundaries,
- Carries source and destination E.164 addresses,
- Carries an indication of the higher layer protocol intended to process the datagram at the destination (HLPI field),
- “Best effort” delivery.

Vaguely analogous to LAN-ish transport service.

IP datagram can be supplied directly as the data field of the CLNAP-UNITDATA.request()

‘Broadcast’ transmissions are an unsupportable concept, tricky for dynamic address resolution of hosts.

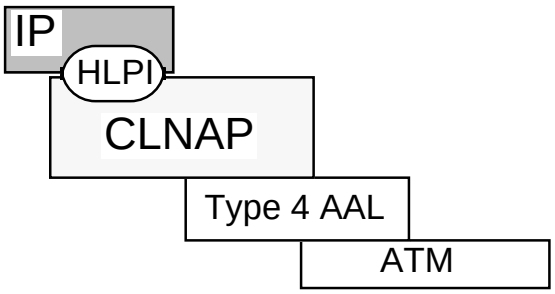

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

AAL3 / AAL4

Relevant characteristics

- Connection oriented,
- Preserves datagram boundaries,
- May or may not guarantee delivery.

Small is beautiful approach suggests

- No other encapsulation needed, IP datagram becomes the AAL-SDU.
- Service Specific Convergence Sublayer (SSCS) may be null when supporting IP.

Analogous to a Point to Point link.

Multicast is possible.

One VCI/VPI pair used per IP entity to IP entity link.

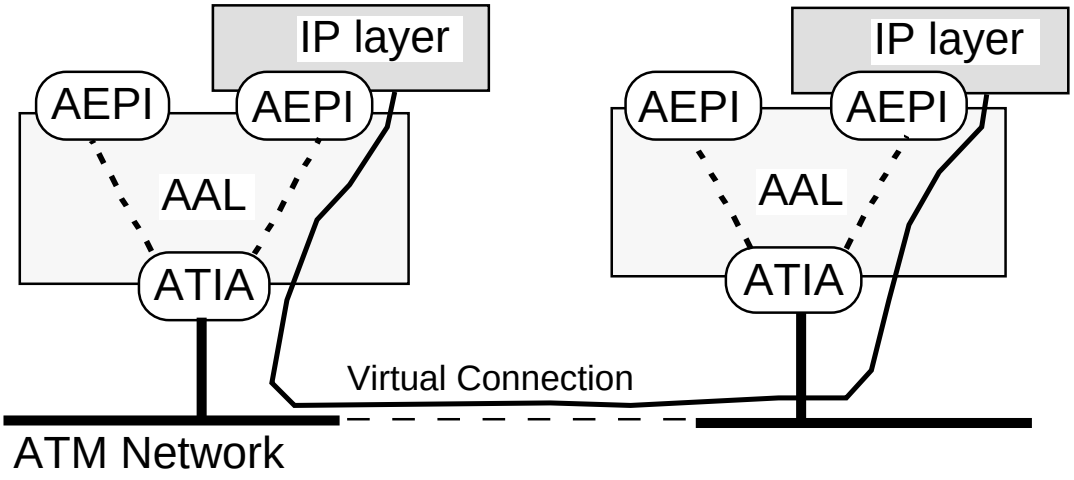

ATIA: ATM Terminal Interface Address (E.164 address)
AEPI: AAL Endpoint Protocol Identifier

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

IP TO E.164 ADDRESSES

IP routing methodologies can be applied if we emulate traditional physical layer topologies with B-ISDN.

Major IP routers linked by dedicated lines usually have statically defined mappings.

B-ISDN can interlink IP routers with little change.

Problem arises when hosts connect through B-ISDN.

- Destination host may be on B-ISDN. How do we locate its E.164 address?
- Destination host may be on another network "off the edge" of the B-ISDN. How do we locate the E.164 address of nearest router?

Seed hosts with E.164 address of "nearest" IP router.

Create Logical IP Subnets (LIS) - multicast groups of hosts and router(s).

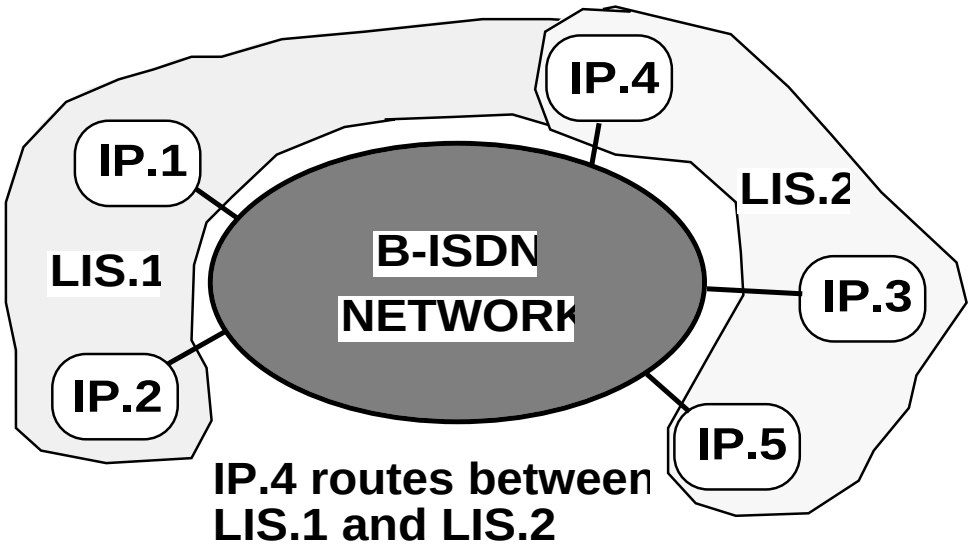

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

LOGICAL IP ROUTERS

Two problems with LIS

- Routers need to know IP addresses of members of the multicast groups it routes between.
- Inter-LIS traffic flows across B-ISDN more than necessary.

'Broadcast' probes to a multicast group may establish if a given IP address is attached.

Logical IP Router (LIR) can reduce inter-LIS traffic

- When destination IP address is not on local LIS, host queries an LIR.
- LIR returns E.164 address of destination, if known.
- If destination not known, LIR returns E.164 address of an LIR logically 'closer' to the destination. This new LIR is then queried.
- When destination is located, mapping is cached locally and IP datagram is sent.

Once the mapping is made, two hosts can converse using a direct B-ISDN connection.

If the destination is 'off the edge' of the B-ISDN the search returns the address of the closest 'real' router.

LIRs still need some mechanism to establish what other hosts and LIRs they 'know about'.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

THE LIR IN CLOSEUP

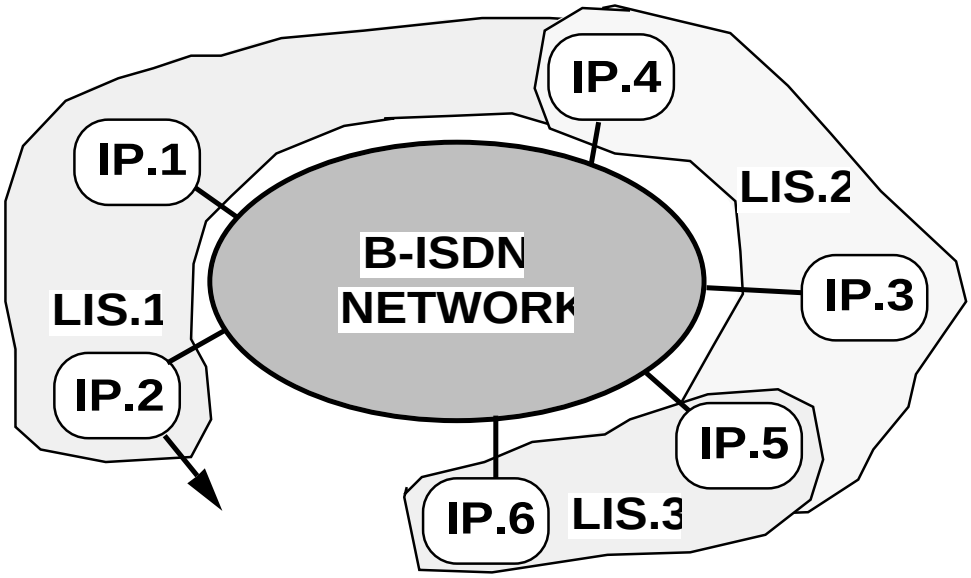

IP.2 is a normal router, connected to B-ISD and other physical networks

IP.4 and IP.5 are Logical IP Route

Closed groups

- Direct B-ISDN connection may be blocked for security, policy, or accounting reasons.

- Place an IP router node into the closed group, but accessible from rest of B-ISDN world.

- When the LIR probe reaches this router, it insists that traffic to the destination must pass through it.

AN OVERVIEW

Dynamic Address Resolution is going to be a difficult area in the evolution of IP over B-ISDN.

Logical IP Subnets using multicast groups will be useful for

- conceptual simplicity,
- accounting,
- security.

Routing between Logical IP Subnets can be minimised if a new ARP strategy is implemented using Logical IP Routers.

Both Connectionless and raw AAL3 or AAL4 services can transport IP datagrams without further encapsulation.

Possible use of MID field in AAL3 and AAL4 has not been looked at.

High speed call establishment would allow ‘on demand’ connections for datagram transport.

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.

AND AN ASIDE

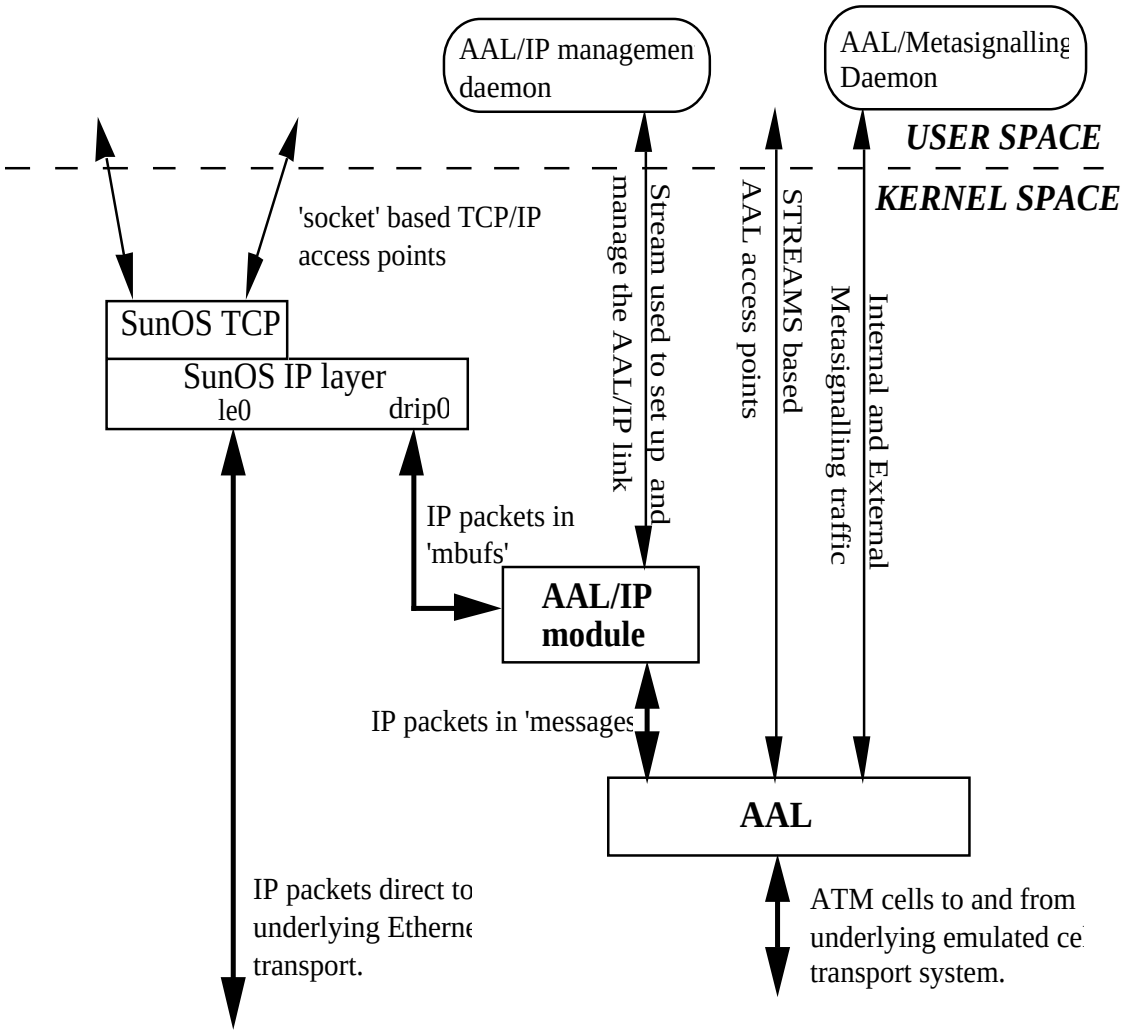

Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, pp.383-390. These presentation slides reconstructed Dec 2005.