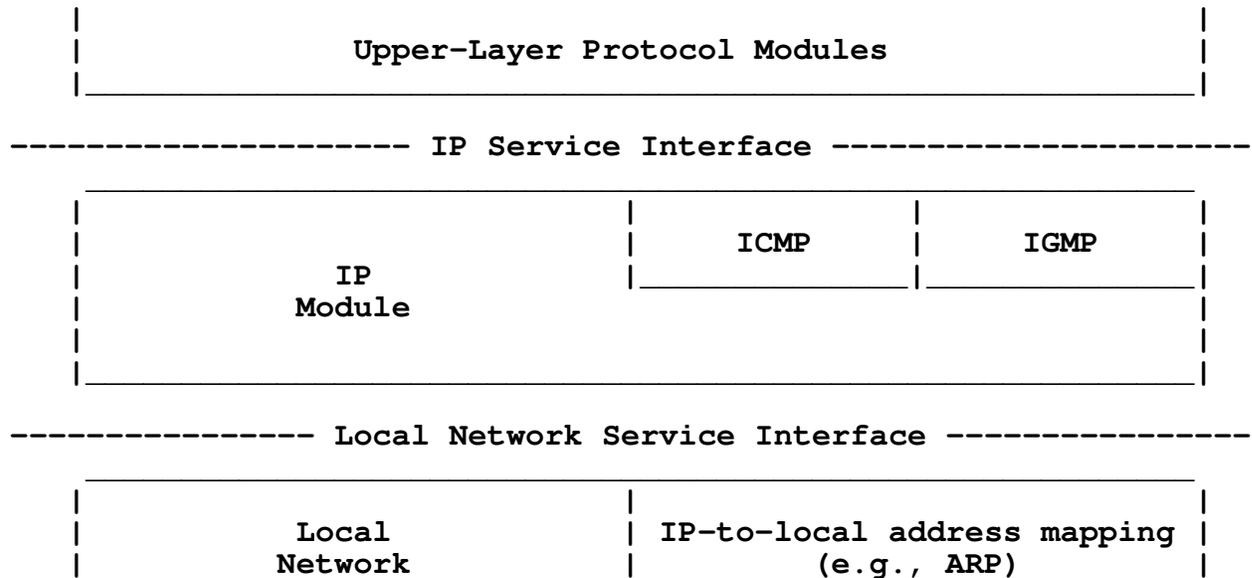


Chapter 13 – Internet Group Management Protocol

Model of host implementation of IGMP



7.1. Extensions to the IP Service Interface

Incoming multicast IP datagrams are received by upper-layer protocol modules using the same "Receive IP" operation as normal, unicast datagrams. Selection of a destination upper-layer protocol is based on the protocol field in the IP header, regardless of the destination IP address. However, before any datagrams destined to a particular group can be received, an upper-layer protocol must ask the IP module to join that group. Thus, the IP service interface must be extended to provide two new operations:

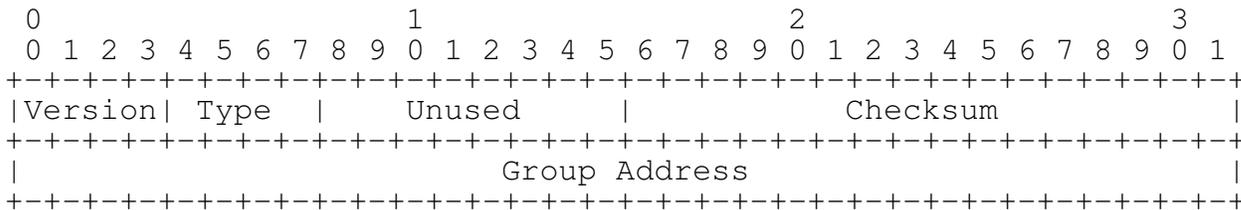
```
JoinHostGroup(group-address, interface )
LeaveHostGroup(group-address, interface )
```

7.3. Extensions to the Local Network Service Interface

Incoming local network multicast packets are delivered to the IP module using the same "Receive Local" operation as local network unicast packets. To allow the IP module to tell the local network module which multicast packets to accept, the local network service interface is extended to provide two new operations:

```
JoinLocalGroup ( group-address )
LeaveLocalGroup ( group-address )
```

IGMP Packet format



Version

This memo specifies version 1 of IGMP. Version 0 is specified in RFC-988 and is now obsolete.

Type

There are two types of IGMP message of concern to hosts:

- 1 = Host Membership Query
- 2 = Host Membership Report

Group Address

In a Host Membership Query message, the group address field is zeroed when sent, ignored when received.

In a Host Membership Report message, the group address field holds the IP host group address of the group being reported.

IGMP Implementation

Host sends an IGMP report when its first process joins an IGMP group
(The report is repeated after random delay ≤ 10 secs).

Hosts don't report group exits <---- (but our Solaris host did!)

Multicast routers periodically IGMP query their LAN

Query address is 224.0.0.1

Hosts reply

using the group address as the destination

after a random delay

==> if one host hears another's reply its own can be discarded.

Hosts send reports for each group which still has a member

IGMP within an isolated LAN supporting hardware multicast
No real need for its existence.

IGMP within multiple LAN's connected by an IGMP capable routers
Query / report technique described above will suffice
(Modulo loops in the topology).

IGMP within an AS
MOSPF is designed to provide optimal routing for multicasts, interact properly with OSPF
and other wonderful stuff

IGMP within the global Internet
The most commonly used protocol is DVMRP (Distance Vector Multicast Routing Protocol)
A distance vector protocol like RIP with some enhancements:
More meaningful costs than hops
Theoretically more robust with respect to routing loops

Characteristics of DVMRP

Routers uses special types of IGMP packets to exchange packets RIP like packets by which
distances to each destination are established.

Once this is done reverse path forwarding is used to perform the multicast.

If packet arrives on the link which is the shortest path back to the source it is
forwarded, else it is discarded.

Is implemented in the `mROUTED` routing daemon

Uses encapsulation to route across non-multicast routers

DVMRP is the basis of the MBONE

PIM (Protocol Independent Multicast)

Two operational modes:

Dense mode: Broadcast and prune

Sparse mode: Based upon rendezvous points