              Toward an Internet Standard Scheme for Subnetting

STATUS OF THIS MEMO

   This RFC discusses standardizing the protocol used in subnetted
   environments in the ARPA-Internet. Distribution of this memo is
   unlimited.

   The author of this RFC is the Gateway Algorithms and Data Structures
   (GADS) Task Force, chaired by David L. Mills.

INTRODUCTION

   Several sites now contain a complex of local links connected to the
   Internet via a gateway. The details of the internal connectivity are
   of little interest to the rest of the Internet.

   One way of organizing these local complexes of links is to use the
   same strategy as the Internet uses to organize networks, that is, to
   declare each link to be an entity (like a network) and to
   interconnect the links with devices that perform routing functions
   (like gateways). This general scheme is called subnetting, the
   individual links are called subnets, and the connecting devices are
   called subgateways (or bridges, or gateways).

   All hosts in the Internet must make a decision when sending a
   datagram, that is, they must answer the question "Is this datagram
   addressed to a host on a directly connected network, or must it be
   sent to a gateway?". In a subnetted environment, this question is
   extended to "Is this datagram addressed to a host on a directly
   connected subnet, or must it be sent to a (sub)gateway?". Let us
   call answering this question "making the routing decision".

   Because the hosts used in a subnetted environment must implement in
   their IP or network interface software procedures for making the
   routing decision, and because such hosts may be acquired from various
   sources, it is important that a standard subnetting scheme be
   identified so that different suppliers can provide compatible hosts
   (that is, hosts compatible with the complexes at different sites and
   each other). Without a designated standard for a subnetting scheme
   suppliers can not create compatible hosts.

   The potential problem is that if different subnetting schemes are
   developed by different suppliers a customer that installs hosts from
   two or more suppliers may find that they do not work together.

   This topic has been discussed in a set of RFCs [1,2,3,4] and in a
   flurry of messages in the Gateway Algorithms and Data Structures Task
   Force. It is strongly suggested that if subnetting is used at all,
   it be according this new standard scheme.

APPROACH

   An Internet address currently consists of a two-layer hierarchy, a
   'network' and a per-network 'rest' field. This subnet scheme adds an
   optional 'subnet' layer and field.

   The subnet field is created by stealing some bits from the rest (or
   host) field of the address. The details of the subnet field are site
   specific. All three classes (A, B, and C) of networks may be
   subnetted.

   The use of subnets is an optional local decision. The fact that a
   network has subnets is invisible outside that network, and the change
   is local and can be instituted at a site without any global Internet
   perturbations. A complex of links is assigned a single IP network
   number, and outside that complex it appears as a single network with
   that number. Only inside does local structure appear.

   However, while the decision to use subnets at a site is optional, any
   IP implementation which may possibly be used in a potentially
   subnetted environment, should provide for subnet field configuration
   as described above. Such an implementation will function properly in
   environments with or without subnetting. On the other hand,
   implementations lacking this provision will not function in a
   subnetted environment, and are thus potentially less useful.

   This specifications is not intended to require a particular
   implementation technique inside the host, but rather to define the
   external behavior of the host in a subnetted environment. It does
   not specify how routing is done or the details of host construction.
   Note that gateways are hosts, too.

   However, it seems easiest to explain the approach by describing one
   possible host implementation.

      Example Implementation:

         Let us use "subnet" to mean the locally attached transmission
         medium.

         The key decision to be made is "Is the destination IP address

on my subnet or not?". Once this decision is made the host
knows to whether to send the datagram directly to the
destination on the subnet or to send the datagram to a gateway.

The host uses a 32-bit mask, along with the host's own IP
address, to determine whether or not destination IP addresses
are on its subnet.

The mask can be configured at boot time as a static quantity or
distributed by a protocol that is beyond the scope of this
memo.

If the bitwise AND of the mask with the destination IP address
matches the bitwise AND of the mask with the host's own IP
address, the destination is assumed on its subnet; if not, the
destination is assumed on a subnet or network reachable only
via a gateway.

   Note: if the mask is all zeros, all destinations will appear
   to be on this subnet; while, if the mask is all ones, only
   the sending host itself will appear to be on this subnet.
   If the mask contains ones in the network field and zeros in
   the rest field, subnets are not in use.

The above procedure must be treated as a per interface
procedure for multihomed hosts.

   For further information on background and rationale, see RFC-917,
   "Internet Subnets" [1].

REFERENCES

   [1] Mogul, J., "Internet Subnets", RFC-917, Stanford University,
   October 1984.

   [2] Postel, J., "Multi-LAN Address Resolution", RFC-925,
   USC/Information Sciences Institute, October 1984.

   [3] Clark, D., "A Subnetwork Addressing Scheme", RFC-932, MIT LCS,
   January 1985.

   [4] Karels, M., "Another Internet Subnet Addressing Scheme",
   RFC-936, UC Berkeley, February 1985.