


From: Paul F. Tsuchiya, Bellcore  
 To: X3S3.3  
 Topic: Proposal to keep NEXT\_HOP in IDRP

At the December 1990 meeting of X3S3.3 in Boulder, the option of expanding the use of the NEXT\_HOP attribute to allow third-party routing information to be passed was discussed. The suggestion was made by someone not present through Yakov Rekhter. Neither Yakov nor anybody else at the meeting could see any use for the expanded function, and indeed Dave Oran pointed out the “lying” is always dangerous in routing, and so it was decided to not expand the use of NEXT\_HOP. In fact, it was decided to get rid of the NEXT\_HOP attribute altogether, since without the expanded use, it didn't provide any new information.

Since that discussion (in fact, the very next day at the IP over Large Public Data Networks working group meeting in IETF), a good reason for the NEXT\_HOP attribute was discovered. In a large public data network, there may be hundreds and thousands of BISs attached. Theoretically, each may want routing information from any of the others. However, it is not efficient for each to establish an IDRP relationship with each other. It is more efficient for a small number of “IDRP servers” to be established, to have each BIS establish an IDRP relationship with an IDRP server, and to have the IDRP server pass on to each BIS information it heard from the other BISs.



Of course, for this to happen, the IDRP Server (which is really nothing more than a BIS itself, except that it does nothing but pass on IDRP information) must be able to advertise a next hop other than itself. For instance, a BIS in Domain X with NSAP Address x would advertise to the IDRP Server that addresses in X were reachable through it, and that the NEXT\_HOP is itself (x). The Route Server would then tell another BIS, say in Domain Y, that X was reachable through NEXT\_HOP x.

Of course, for Y to know how to send a packet to NSAP Address x, it needs to know the SNPA (public network) address of x as well. While this could be learned using 10030, it is convenient for the IDRP Server to tell the BIS the subnetwork address at the same time it tells the BIS the NSAP Address. Indeed, the BIS actually never needs to know the NSAP address of another BIS that it hears about through an IDRP Server.

Therefore, I propose that we keep the NEXT\_HOP Attribute, and expand its use so that any BIS can advertise third party information. Of course, for this to work, the BIS advertised in the NEXT\_HOP attribute must be reachable directly over a subnetwork by the BIS receiving the advertisement. This must be configured properly, since it is not always possible to verify that an address is so reachable. Also, since most BISs will not be giving out NEXT\_HOP information, one would probably want to configure whether a neighbor can legally advertise a NEXT\_HOP at all.

I suggest that the NEXT\_HOP attribute work as follows. If no NEXT\_HOP is being advertised (in other words, the next hop is the BIS sending the advertisement), then the next hop field has a length of 0, and no address is given. If a NEXT\_HOP is being advertised, then the attribute has two parts, the NSAP address and the SNPA address. It should be encoded as follows:

Length of NET (1 octet)
NET of next hop (variable)
Length of SNPA (1 octet)
SNPA of next hop (variable)