Title: Comments on ISO 8348: 1987/PDAM5 (Group NSAP Addressing)

Source: IBM

Reference: SC6 N7551 (X3S3.3/92-356)

IBM is concerned that the text of this PDAM, as written, is founded upon a number of misconceptions, and contains numerous technical errors. Because these errors and misconceptions are spread throughout the PDAM text, the PDAM presents a very confusing image. In the area of addressing, it is vitally important to use precise language, and this has not been done in the PDAM text. For example, there is language implying that a single AFI can be expressed in either of two different abstract syntaxes, which clearly is not the case. There in no definition given for ″unicast AFI″ or ″multicast AFI″. There is no table for multicast AFIs that presents the type of information given in Table 2 for unicast AFIs, etc.

We offer the following specific comments, and suggested resolutions for each of them:

1. Inappropriate References (E): All clause references in the PDAM text should refer to the final consolidated text of ISO 8348, which is contained in SC6 N7558 (X3S3.3/92-391).

   Note: However, in order to maintain continuity of our comments, we continue to refer to clauses in ISO 8348/Add.2, since this is the method used in SC6 N7551.

2. Clause 0 (M): Although the third paragraph claims that the amendment provides a ″new abstract syntax for every existing AFI″, this clearly is not the case. The existing AFIs are not changed in any way. What the proposed amendment does is to define a new set of AFIs for group addresses: the new AFIs will be expressed in the binary abstract syntax, and constraints will be placed on the allowable binary values that can be used for the new AFIs. We also note that it is not the use of the binary abstract syntax that makes a group NSAP ″easily distinguishable″: it is the fact that the preferred binary encoding of the proposed AFIs is guaranteed to produce different encoded bit strings than will be obtained from the preferred binary encoding of the AFIs of ″conventional NSAPs″.

   Therefore, we suggest that the last paragraph of clause 0 be replaced in its entirety with the following text:

   The AFIs of NSAP addresses that are presently allocated or reserved for allocation by ISO 8348/Addendum 2 are expressed in the decimal abstract syntax. These AFIs are all unicast AFIs.

   This amendment defines a new set of AFIs to be used for group NSAP addresses, referred to as multicast AFIs. Multicast AFIs are expressed in the the abstract binary syntax. Allowed values for multicast AFIs are limited to the values shown in Table X in order to guarantee that their preferred binary encoding will differ from the preferred binary encodings of the unicast AFIs. Each multicast AFI is associated with a single unicast AFI, as shown in Table X, and the pair of AFIs is administered by a single addressing authority.

3. Clause 3 (M): A group NSAP address that identifies zero NSAPs is uninteresting, and a group NSAP that identifies exactly one NSAP is in fact a unicast NSAP, as defined today in ISO 8348/Add.2. The definition in the text also confuses the constraints on usage of group NSAPs (lim-

iting them to destination addresses only) with their actual semantics. Hence, we recommend that the definition be replaced in its entirety with the following new definition. We also recommend that a definition be provided for a ″unicast NSAP″, since this new term is used in the proposed changes to clause 8.2.1:

> ***group network address:*** *a single network layer address that identifies two or more Network entities.*
>
> ***unicast network address:*** *a single network layer address that identifies a single Network entity.*

To provide an unambiguous way to refer to already existing AFIs and the new AFIs proposed by this amendment, we also recommend adding two more definitions to 6.1.2:

> ***unicast AFI:*** *an AFI that is used to identify a unicast NSAP address.*
>
> ***multicast AFI:*** *an AFI that is used to identify a multicast NSAP address.*

4. Clause 6.1.2 (E): The proposed text to be added as a new paragraph at the end of 6.1.2 is very muddled. It seems to give a defintion of group NSAPs (which is already handled in clause 3), and then places a constraint on where Group NSAP addresses should not be used. Furthermore, the positioning as a new last paragraph seems illogical. The key points seem to be that group NSAP addresses can be used as the destination address of the N-UNITDATA primitive, but not as its source address. Having fixed up the ″definitions″ clause as above, we suggest that the critical information can be conveyed more clearly by replacing the suggested changes with the following:

   Replace the words ″NSAP addresses″ at the end of the first sentence of the third paragraph of 6.1.2 (The paragraph begins with *The NSAP address is the information...*) with the following new text:

   > *either unicast NSAP addresses or group NSAP addresses. Group NSAP addresses can not be used in the ″calling address″ or ″responding address″ parameters of the N-CONNECT primitive, nor can they be used in the ″source address″ parameter of the N-UNITDATA primitive.*

5. Clause 8.2.1.1 (M): This clause is technically incorrect: it states that each AFI has two syntaxes, when in fact a given AFI has one and only one abstract syntax. The text obscures the pairing of a unicast AFI and a mulitcast AFI, by referring to a single AFI with two different syntaxes. The concept of pairing is critical, since without it there would be need for an entirely new set of address administration authorities. To clear up this confusion, a number of changes need to be made:

   a. Table 1 of ISO 8348/Add.2 should be renamed ″Unicast AFI Allocations″

   b. A new Table 1.X should be added, entitled ″Multicast AFI Allocations″, which will list the corresponding information for multicast AFIs.

   c. Table X of the proposed text should be renamed ″Corresponding Pairs of Unicast and Multicast AFIs″, and a note should be added to the table saying that each pair of <unicast AFI, multicast AFI> is administered by the same addressing authority.

   d. Replace the suggested text changes for 8.2.1.1 with the following new text:

      > *There are two types of AFI: unicast AFIs and multicast AFIs. The abstract syntax of unciast AFIs is 2 decimal digits; the abstract syntax of multicast AFIs is one binary octet. Unicast AFI allocations are shown in Table 1, and multicast AFI allocations are shown in Table 1.X. Each pair of unicast and multicast AFIs shown in Table X is administered by the same addressing authority.*

e. Table 2 needs to be retitled ″Allocated Unicast AFI Values″, and a new Table 2.X needs to be added to provide the equivalent information for ″Allocated Multicast AFI Values″.

f. The proposed text to be added to 8.2.1.2 is no longer necessary, given that we create Table 2.X. All that needs to be done is to add a reference to Table 2.X in the very first sentence: ″...as summarized in table 2 *and table 2.x.*″

6. Clause 8.3 (e): The proposed amendment offers no change to this clause. However, point ″a″ in fact needs to be expanded to include the multicast AFIs:

   a) the AFI, with an abstract syntax of two decimal digits for unicast AFIs, *and with an abstract syntax of one binary octet for multicast AFIs;*

7. Clause 8.3.1 (M): The suggested changes negate the main purpose of this PDAM—which is to establish a set of AFIs that unambiguously mark an NSAP address as ″multicast″. The language ″when the AFI is to be used for a group address″ implies that an AFI could be used either as a unicast or a multicast AFI. As noted in other comments, this is simply not the case: unicast AFIs can only be used in unicast NSAP addresses, and multicast AFIs can only be used in multicast NSAP addresses. We recommend scrapping the proposed changes to 8.3.1 and replacing them with the following:

   a. In item ″a″, replace the words ″the AFI″ with ″a unicast AFI″

   b. Add a new item ″b″ and renumber the remainder of the items:

      b) using a single octet to represent directly the binary octet of a multicast AFI;