ATNP/WG2-16
WP
28 September 1998

AERONAUTICAL TELECOMMUNICATIONS NETWORK PANEL

Working Group 2

16th Meeting

Bordeaux, France

30 September - 2 October 1998

Issues Raised on Internet Communications Service SARPs

Working Paper

Presented by Klaus-Peter Graf

Summary

For information of the WG2 members, this paper compiles a number of technical issues which have been raised on the current specification of the ICS SARPs but which have not been officially submitted as defects.

The working group is invited to note and review these issues and to agree on amendment text to the current ICS SARPs or Guidance Material if deemed appropriate.

1 Introduction

This paper compiles a number of technical issues concerning the current text of the ICS SARPs for review and discussion by the working group. These issues have been raised by several sources but have not been officially submitted as defects against the ICS SARPs at this stage.

2 Octet-Aligned NSAP Prefixes

2.1 Reported Issue

The ICS SARPs section 5.8.3.2.9.2 specifies that prefixes for routes must be octet-aligned when originating or reducing routes. The intent of this requirement being to simplify prefix matching as the note 4 in the same section states. Note 3 of section 5.8.3.2.9.2 states that an implementation must still be able to handle bit-aligned prefixes. The problem here is that if a router design caters for bit-aligned prefix, then the prefix matching must also be bit-aligned. Therefore losing any of the savings intended by the SARPs.

Prefix matching is of course a key event. It will be carried out multiple times for every CLNP packet forwarded. So it is a performance issue. Currently the only possibility of a bit-aligned prefix being generated is from NLRI aggregation. It seems an unlikely event that two routes for aggregation will only differ in the last bit.

The question boils down to whether the ICS SARPs can mandate that all prefixes in the ATN internet must be octet-aligned in order to reap the benefits of this restriction ?

2.2 Discussion

The requirement for bit-aligned prefix handling by ATN routers has been carried forward from the ISO/IEC 10747 standard. Explicitly removing this requirement from the ICS SARPs for the performance reason mentioned above would make ATN IDRP implementations non-compliant with the ISO/IEC 10747 base standard. This is considered to be not appropriate. Therefore, no action on this issue is proposed.

3 End-of-Block Code for Deflate Data Blocks

3.1 Reported Issue

Section 5.7.6.5.4.1.3 does not include the value 256 (corresponding to the end-of-block code) in the set of used Huffman encoded values. Note 2 of this section justifies this omission by the fact the end-of-block code has not to be used, as each Deflate Data Block is delimited by ISO/IEC 8208 packet boundaries.

However, this implies that if the encoding of an NPDU uses a Deflate Data Block which is compressed with fixed Huffman codes, that this data block must be the last, and that the compression function is not free to switch at any time between blocks of the available three types:
- uncompressed
- compressed with fixed Huffman codes
- compressed with dynamic Huffman codes.

This seems to contradict the note following 5.7.6.5.2.3 which explains that "the compressor may decide to switch between any one of these strategies at any time and not just at an NPDU boundary".

5.7.6.5.4.1.3 also seems to contradict the Note in 5.7.6.5.4.2.5.3 which explains that removal of trailing zero octet is highly likely to occur because padding after end-of-block marker with zero bits will frequently yield a trailing zero octet.

3.2 Discussion

It is possible to have more than one Deflate Data block in the same NPDU. If one applies the rule of Note 2 to intermediate blocks an error will result as the Deflate format relies on the use of the end-of-block marker to separate strings of the three types mentioned above. Therefore, the requirement should only apply to the last Deflate block in an NPDU and not to every block.

3.3 Proposal

The following modifications are proposed:

1. In section 5.7.6.5.4.1.3 replace „255“ by „256“
2. Rephrase Note 2 of section 5.7.6.5.4.1.3 to read: „*Note 2.—In IETF RFC 1951, the value 256 indicates end-of-block. An octet containing this value may be removed from a Deflate Data Block, if this data block is the last one in an NPDU. In this case the block is delimited by the NPDU boundary and by not using this value, the size of the compressed data stream is reduced.*“

4 Removal of Trailing Zero Octet of Deflate Block

4.1 Reported Issue

Section 5.7.6.5.2.9 states that the last deflate data block should be right-padded with zero-bits to fill octet. This seems to apply to blocks compressed using fixed or dynamic Huffman codes. However trailing zero octets are only stripped if the final Deflate Data Block in a compressed NPDU uses Fixed Huffman codes. Likewise on decoding a zero octet is only appended to the final block if it used Fixed Huffman codes. Is it correct to assume that this removal/addition of the zero octet does not apply to dynamic Huffman codes ?

4.2 Discussion

The requirement to remove the trailing zero octet came from inspection of the results of the compression algorithm used for testing the applicability of the Deflate algorithm for the ATN. In almost all cases, the last octet was zero and conveyed no useful information. This is because the code for end of stream is binary 256 and you have to right pad with zero bits to the next octet boundary. It made sense to delete this octet as it could always be assumed by the decompressor if reached the last octet without finding an end of stream code.

Deletion of the final zero octet only applies to fixed Huffman codes. When the codes are dynamically determined, one cannot assume that the end of stream code is encoded as binary 256 and consequently may be removed. A note may be added

4.3 Proposal

An appropriate note should be added to the SARPs to clarrify this issue.

5 Backwards Window Size

5.1 Reported Issue

Section 5.7.6.5.7.1.3 recommends use of full 32Kb range for distance values, while Note 1 states that this permits an implementation to limit the size of the backwards window (implying less than 32kb). Section 5.7.6.5.7.2.3 also states that the last 32Kb be used for matching octet sequences.

Is an implementation required to use the full 32Kb on encoding or can it decide to use a smaller backwards window to save on memory. If so is there a minimum backwards window it should try to use ?
Is an implementation required to support the use of the full 32kb backward window for Decoding ?

5.2 Discussion

There is no mechanism to communicate the window size used by the compressor to the decompressor, hence the decompressor has to assume a 32KB window. On the other hand, the compressor can still use a smaller window without notifying the decompressor or affecting this. The memory requirements of the compressor are much larger than the decompressor, as it must maintain large hash tables in order to perform the compression, whilst the decompressor needs only to maintain a buffer holding the last 32KB decompressed, hence there may well be benefit in decreasing the window size by the compressor.

The analysis of the Deflate memory requirements reported in ATNP WP xxx ("Data Link Compression Evaluation Report (2.1)" came to the conclusion to recommend an 8KB window for initial ATN use.

5.3 Proposal

The following modification to Note 1 of section 5.7.6.5.7.1.3 is proposed:
Note 1.—This permits an implementation of the compressor to autonomously limit the size of the backwards window used to compress data in order to optimise the use of memory resources. However, the result will be a poorer compression ratio. On the other hand, the decompressor must always be able to accept any valid distance value, i.e. must maintain a 32KB buffer.

6 Aggregation of routes with same Security RIB-Att

6.1 Reported Issue

Consider BIS A and BIS B that have established an IDRP connection with each other. Suppose BIS A wishes to advertise to BIS B two routes to a third BIS, say BIS C. Furthermore, suppose that these two routes have the same NLRI and (being ATN routers) the same RIB-Att (namely the Security attribute), but these two routes differ only in the values of the Security attribute. Is the BIS allowed to treat these as two separate routes? More specifically, can BIS A place each of these routes in a separate Adj-RIB-Out?

(Clearly, we want the answer to this question to be 'yes', but I can't find the justification in ISO/IEC 10747. Indeed, on page 33 of that standard it is stated: "Within the set of Adj-RIBs-Out associated with a given neighbor BIS, no two shall have the same RIB-Att". Now, in the scenario mentioned above, these two routes have the same RIB-Att, namely the set of attributes comprising only the Security attribute, but they have different values of this type-value specific attribute).

6.2 Discussion

The received routes described above are subject to aggregation before insertion into the appropriate Adj-RIB-Out. Section 5.8.3.2.6.4 and sub-ordinated paragraphs specify the relevant aggregation rules. These rules use the value of the security path attribute information field as input parameter. For example, if one route is an ATSC route and the other one a non-ATSC route and both routes have dissimilar NLRI, then they shall not be aggregated. On the other hand, if both routes are ATSC routes, contain an ATSC Class Security Tag, and have dissimilar NLRI, then these routes shall be aggregated and the ATSC Class of the aggregated route shall be the lowest ATSC Class of both routes.

The specification in chapter 5.8.3.2.6 is considered clear and detailed enough, in order to stay as it is.

7 Advertisement of changes in mobile subnetwork connectivity

7.1 Reported Issue

Consider BIS A and BIS B where BIS A is an airborne router and BIS B is an A/G router on the ground. Suppose that these two BISs have established a BIS-BIS connection via the route initiation process. Furthermore, suppose that this BIS-BIS connection currently utilizes a subnetwork that allows ATSC traffic of Classes D or below (i.e. Classes D, E, F, G, H, and No Preference). Let us designate this subnetwork as subnetwork 1. Now, suppose, an additional subnetwork, say subnetwork 2 becomes available between these two BISs, and suppose that subnetwork 2 allows ATSC traffic of Classes B or below. I would make the argument that BIS B should advertise the fact that it now has an additional route to BIS A with a different class to other routers on the ground, because there is now increased flexibility in terms of the classes of traffic that can be sent, and it seems reasonable that we would want other routers to be informed of this increased flexibility. Am I correct

7.2 Discussion

Section 5.8.3.2.4.1.1 specifies that the A/G router has to update the security path attribute to reflect the current availability of mobile subnetworks before advertising the route further downstream on the ground. Therefore, this is not considered to be an issue and no amendment to the SARPs is proposed at this stage.

8 Frequently changing subnetwork connectivity

8.1 Reported Issue

Recently some scenarios were reported that had been encountered while working on FANS-1. These scenarios may impact the specification on updating the security information by an A/G BIS prior to advertisement to an airborne BIS in the current ICS SARPs.

In the encountered scenario an aircraft was flying along the coast of Australia, just on the fringe of VDL coverage, with a Satellite connection. The aircraft was constantly coming in and out VDL coverage. A "media advisory" message was generated by the ground station every time the VDL connection changed. Each message to be sent to the aircraft. These messages built up until a constant subnetwork connection was established. They then where all transmitted.

A similar scenario could arise in the ATN. A new update would be generated every time the VDL connection changed. Given the slow satellite connections a list of Updates could build awaiting transmittal, wasting bandwidth.

8.2 Discussion

The requirement to re-send the routes every time the subnetwork connectivity changes is in section 5.3.5.2.10.5 and was added in an attempt to inform the airborne router of any changes to the ATSC Class. Consequently this issue is closely related with the issue raised in PDR 98060006 and should be considered when developing a solution for this PDR.