International Civil Aviation Organization

THE FIRST MEETING OF AERONAUTCAL COMMUNICATION SERVICE (ACS) IMPLEMENTATION CO-ORDINATION GROUP OF APANPIRG (ACSICG/1)

Seoul, Republic of Korea, 13 - 16 May 2014

Agenda Item 11: **Any Other Business**

GUIDANCE CONCERNING THE USE OF P1 ASSOCIATIONS BETWEEN MTAS

(Presented by United States of America)

SUMMARY

This information paper presents a discussion AFSG regarding several options for operating P1 associations. There is a particular focus on the question of whether or not an MTA should operate multiple simultaneous associations to neighbor MTA.

1. Introduction

1.1 The attached paper, presented at the Aeronautical Fixed Services Group (AFSG) meeting in April 2014, discusses issues to be considered when configuring and operating P1 associations between MTAs. Several options are considered, as well as the implications and advantages/disadvantages of each. While there are references to networks within the European Region, the concepts are transferable to all regions.

2. Discussion

2.1 See Attachment A.

3. Action by the Meeting

3.1 The meeting is invited to:

a) note the information contained in this paper; and

b) discuss any relevant matters as appropriate.

– – – – – – – – – – – – – –

European and North Atlantic Office

AFSG/18 – IP/05
19/03/2014

AERONAUTICAL FIXED SERVICES GROUP (AFSG)

of the European Air Navigation Planning Group (EANPG)

EIGHTEENTH MEETING

(Paris, 7 - 11 April 2014)

Agenda Item 4: Technical/Documentation matters

GUIDANCE CONCERNING THE USE OF P1 ASSOCIATIONS BETWEEN MTAS

(Presented by Germany)

SUMMARY

This information paper discusses several options to operate P1 associations and gives guidance for the European AMHS network. There is a particular focus on the question, whether an MTA should operate multiple simultaneous associations to a neighbour MTA.

References

[1] ICAO EUR Doc 020 - EUR AMHS Manual, Version 8.0, 25/04/13 - Appendix B, Annex J (normative) - Requirements of OSI upper layers for AMHS - Annex B (RTSE).

[2] ISO/IEC 10021-6:2003 / ITU-T Recommendation X.419 (06/1999) - Message Handling Systems (MHS): Protocol Specifications.

[3] ISO/IEC 10021-4:2003 / ITU-T Recommendation X.411 (06/1999) - Message Handling Systems (MHS) - Message Transfer System: Abstract Service Definition and Procedures.

1. Introduction

1.1 There are several options to operate P1 associations (i.e. logical connections) between two adjacent MTAs:

a) An association can be monologue or two-way-alternate.

b) An association can be permanent or dynamic (i.e. on-demand).

c) An MTA can operate a single association or multiple simultaneous associations to a neighbour MTA.

(5 pages)

1.2 Before an association can be set up, partner organisations must agree on the (maximum) number and the type of associations they want to use between their MTAs. From a technical point of view, the MTA that opens an association (the initiator) is responsible to determine whether the association is only temporary (dynamic) or permanent and proposes to the peer MTA either the monologue or two-way-alternate mode. The peer MTA (the responder) needs to accept the association and in particular the mode of operation.

1.3 No clear instructions are given in the ISO/IEC standards or ITU-T Recommendations, on which option or which combination of options to use. There are also a number of open questions remaining from specifications in ICAO Doc 9880 Part II and the EUR AMHS Manual. Therefore, conformant AMHS implementations usually support all options leaving the decision how to configure and operate the AMHS system up to system administrators and operators.

1.4 This working paper briefly explains the characteristics of the above mentioned options and gives some guidance how to use P1 associations in a (fully) meshed European AMHS network.

2. Discussion

Monologue versus two-way alternate

2.1 When an MTA initiates an association in the two-way alternate mode and the responding MTA accepts this mode, both MTAs can send and receive messages, reports and probes over this (single) association. In contrast, when an MTA initiates an association in the monologue mode, it can only send but not receive and, therefore, the peer MTA needs to initiate a second association for bidirectional information exchange.

2.2 When two-way-alternate mode is used, there is a need to exchange a data token between the MTAs. Only the MTA, which is in possession of the token, has the right to send and the token needs to be passed to the other side, whenever the transfer direction changes. This procedure causes delay in case of bidirectional communications, which can be significant in high traffic situations.

2.3 There is no substantial value in having two-way-alternate mode rather than two monologue associations. A possible reason for having two-way-alternate mode could be to generally allow only one side to initiate associations, for example for billing purposes or because of a security policy.

2.4 *Recommendation*: To avoid delay induced by token exchanges a general preference should be given to operate P1 associations in the monologue mode.

2.5 Note that the support of the monologue mode is mandatory for AMHS implementations which conform to the European ATS Messaging Service Profile [1].

Permanent versus dynamic associations

2.6 A permanent association is opened once (e.g. at system initialization or by operator’s action) and remains open until explicitly closed (e.g. disabled by operator’s action or system shutdown) or aborted by interruption of the network service. In contrast, a dynamic association (also called on-demand association) is opened only when a message needs to be sent and then closed again.

2.7 Opening an association involves a protocol handshake consisting of sending one association request (AARQ) packet containing the MTA Bind-Argument for the MTA’s authentication and one association response (AARE) packet containing the MTA Bind-Result. Normal closure of an association involves sending one association release request (RLRQ) and one association release response (RLRE) packet. The protocol handshake prior to sending the first message typically needs 100 to 200 milliseconds, but can also take 500 milliseconds or more, for example in satellite connections, where signal propagation time is significant.

2.8 To avoid too many handshakes, a dynamic association is usually configured with a connection hold time. As long as the connection remains open, messages can be transferred without repetition of handshakes. With 5 minutes connection hold time, for example, a new handshake will be necessary only after 5 minutes have passed without any traffic.

2.9 Compared to dynamic associations, the operation of permanent associations consumes more system resources and may even need additional operational support for monitoring and control. Particular tasks of analysing the actual connection status or disabling/enabling a connection may become necessary, for example, when network service problems affect an open (permanent) association.

2.10 Considering a fully meshed European network, where each MTA may have up to 50 adjacent MTAs, the total number of associations that an MTA must support is considerable. Up to 100 simultaneous associations (50 inbound and 50 outbound) can be required to be maintained in open state, when permanent associations are operated in monologue mode. This number increases when an MTA initiates multiple outgoing associations to an adjacent MTA or priority based associations are used (as explained further below).

2.11 Permanent associations can be maintained to neighbour MTAs as long as the total number of associations does not exceed system limitations.

2.12 Dynamic (on-demand) associations can be useful in test environments (as long as delay is not an issue) or in very low traffic situations, for example, when two MTAs exchange messages occasionally, but not every day or every hour.

2.13 *Recommendation*: In an operational AMHS environment, a general preference should be given to operate permanent P1 associations to avoid messages being delayed due to protocol handshakes required for opening an association. Dynamic (on-demand) associations may be used in test environments or for the operation of very low traffic connections.

Multiple simultaneous associations between a pair of MTAs

2.14 Note: The following statements basically assume that permanent associations are operated and the monologue mode is used.

2.15 In a meshed network it is obvious that an MTA must be able to handle multiple associations simultaneously, for example one incoming and one outgoing to each adjacent MTA. But what is the reason or advantage of initiating a second or even maintaining multiple outgoing associations to a single adjacent MTA? An answer can be given when looking at protocol operations in a particular traffic situation.

2.16 If only a single outgoing association is available, a blocking situation can occur, when an MTA needs to send a second message while a first is still in transfer. Consider a scenario, where an MTA receives a long message from one user (or direction) immediately followed by a short message from another user (or direction) both to be forwarded to the same next hop MTA. On a single association the MTA cannot start sending the second message as long as the previous message is in transfer. Instead, the second message has to wait until the transfer activity for the first is completed (Activity End SPDU[1] is sent) and acknowledgement is received (Activity End Ack SPDU is received) from the peer MTA. The resulting delay for the second (short) message can thus become significant or even unacceptable. On 64 kbps links, for

[1] Session Protocol Data Unit

example, a 64 kilobytes message in transfer can delay the next following message up to 10 seconds and a message of 2 megabytes can cause delay for the next following of even 5 minutes or more.

2.17 With limited bandwidth, high transit times have to be tolerated for long messages, but they could be unacceptable for short messages, and in particular for high-priority (SS) messages.

2.18 If an MTA can use a second outgoing association, two messages from the outgoing queue can be transferred to the next hop MTA in parallel, so that a single long message in transfer would not hold back a second message until this transfer is completed and acknowledged. A short message, for example, could overtake a long message in transfer.

2.19 *Recommendation*: The following points should be considered when defining the (maximum) number of associations between two adjacent MTAs:

a) When the monologue mode is used, an MTA must operate at least one outgoing and one incoming association to allow bidirectional information exchange with an adjacent MTA.

b) If an MTA has a restriction to use only one outgoing association, a long message in transfer can delay the next following (short) message significantly.

c) If an MTA has two or more outgoing associations available, messages can be sent in parallel and short messages, for example, can overtake long messages.

d) In a (fully) meshed European AMHS network the total number of associations to be maintained by an MTA in open state to its neighbours may put constraints on the number of simultaneous associations the MTA can initiate to one adjacent MTA.

An approach to priority based associations

2.20 There is additional benefit from using multiple associations, when the MTA implementation has the capability to assign a minimum priority to an outgoing association. This capability is not a standard requirement, but an option described in section 11.4 of the Message Transfer Protocol Specification - see ISO/IEC 10021-6 or ITU-T Rec. X.419 [2]. In such an implementation[2], an MTA is able to maintain one dedicated outgoing association only for sending high priority messages, while other outgoing associations could be used for lower priority messages.

2.21 An example with three configured outgoing (permanent) P1 associations is illustrated in the following table. The associations have a minimum priority assigned corresponding to the message priority argument (urgent, normal, non-urgent) as defined in ISO/IEC 10021-4 or ITU-T Rec. X.411 [3].

Outgoing P1 association	Minimum Priority configured	Possible traffic
A1	urgent (2)	only urgent messages can be sent over this association
A2	normal (0)	urgent and normal messages, reports and probes can be sent over this association
A3	non-urgent (1) or none	all messages can be sent over this association

2.22 In this example, only the non-urgent traffic is restricted to a single association and hence exposed to the risk of blocking situations as explained above. Normal priority messages, however, have two and urgent messages three associations that they can use.

[2] For example, AMHS systems with an integrated Isode M-Switch (Release R14 or later) support this capability.

2.23 Such a priority based association configuration can also be used in peek traffic or network overload situations to restrict outgoing traffic to higher priority messages, only. For example, an operator could disable association A3 to block low priority messages and only allow normal and higher priority messages to be transferred or, when disabling associations A2 and A3, even restrict outgoing traffic to urgent messages, only.

3. Action by the Meeting

3.1 The AFSG is invited to note the contents of this information paper.

– END –