Network Working Group
INTERNET-DRAFT

Brian Bidulock
OpenSS7 Corporation

Expires in six months

January 10, 2002

Correlation Id and Hearbeat Procedures (CORID)
Supporting Lossless Fail-Over between SCTP Associations
for
Signalling User Adaptation Layers
<draft-bidulock-sigtran-corid-00.ps>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of Section 10 or RFC 2026. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as 'work in progress'.

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

To learn the current status of any Internet-Draft, please check the Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ftp.ietf.org (US East Coast), or ftp.isi.edu (US West Coast).

Abstract

This Internet-Draft describes Correlation Id and Heartbeat procedures to support lossless fail-over between SCTP [RFC 2960] associations for Signalling User Adaptation Protocols [M3UA, SUA, TUA] above MTP3 [Q.704] supporting the concept of a *Routing Context*. These procedures permit lossless fail-over between Application Server Processes (ASPs) at a Signalling Gateway (SG) and fail-over between Signalling Gateway Processes (SGPs) and Signalling Gateways (SGs) at an Application Server Process (ASP). Lossless fail-over permits these fail-overs to occur without loss or duplication of UA-User messages.

1. Introduction

1.1. Scope

This Internet-Draft describes Correlation Id and Heartbeat (CORID) procedures to support lossless fail-over between SCTP [RFC 2960] associations for Signalling User Adaptation Protocols [M3UA, SUA, TUA] above MTP3 [Q.704] supporting the concept of a *Routing Context*. These procedures permit lossless fail-over between Application Server Processes (ASPs) at a Signalling Gateway (SG) and fail-over between Signalling Gateway Processes (SGPs) and Signalling Gateways (SGs) at an Application Server Process (ASP). CORID permits these fail-overs to occur without loss or duplication of UA-User messages.

UA implementations with **CORID** are intended to be compatible with UA implementations not supporting this configuration; however, the full benefits acheived by the **CORID** procedures will not be realized unless implementations at both endpoints implement **CORID**.

1.2. Terminology

CORID supplements the terminology used in the UA documents [M2UA, M3UA, SUA, TUA] by adding the following terms:

Changeback – the MTP3 [Q.704] procedure for redirecting signalling traffic back to a primary linkset from an alternate linkset.

Changeover – the MTP3 [Q.704] procedure for diverting signalling traffic from a failed primary linkset to an alternate linkset.

Lossless Fail-Over – is mechanism for fail-over between SCTP [RFC 2960] associations (i.e, between ASPs, IPSPs, SGPs or SGs) that provides for the elminitation of duplication or loss of UA-User messages between SG and AS.

Message Duplication – a situation where multiple copies of a UA-User message arrives at a Signalling Endpoint.

Message Loss – a situation where instances of a UA-User message is lost in transit between Signalling Endpoints.

Message Mis-sequencing – a situation where UA-User messages that are intended to arrive in sequence, arrive at a terminating Signalling Endpoint in an order other than the order in which the messages were transmitted at the originating Signalling Endpoint.

Signalling Endpoint (SEP) – in this document, a *Signalling Enpoint* is an SS7 SEP [Q.700] or an Application Server.

Signalling Peer Process (SPP) – refers to an ASP, SGP or IPSP.

Signalling User Adaptation Layer (UA) – one or more of the Stream Control Transmission Protocol (SCTP) [RFC 2960] SS7 Signalling User Adaptation Layers [M2UA, M3UA, SUA, TUA] supporting the *Correlation Id* parameter and the **BEAT** message.

Time-controlled Changeover – the MTP3 [Q.704] procedure for diverting signalling traffic from a failed primary linkset to an alternate linkset where sequence number information cannot be exchanged between signalling points or where it is undesirable to use the normal changeover procedures.

1.3. Overview

The existing UA [M3UA, SUA, TUA] procedures do not include procedures to avoid loss or duplication of messages when a UA peer must fail-over between SCTP [RFC 2960] associations between diverse Application Server Processes (ASPs), Signalling Gateway Processes (SGPs), Signalling Gateways (SGs), and IP Signalling Processes (IPSPs).

CORID provides procedures to eliminate message loss, duplication or mis-sequencing under all failure, deactivation, recovery and activation scenarios. **CORID** provides the following capabilities that are not provided for in the existing UA specifications:

- Support for elimintating mis-sequencing of UA-User messages at signalling endpoints (Application Servers or SS7 SEPs) when diverting messages between ASPs, SGPs, SGs, or IPSPs by supporting **BEAT** procedures analogous to the MTP3 [Q.704] Changeback procedure.

- Support for eliminating duplication of UA-User messages at signalling endpoints (Application Servers or SS7 SEPs) or SS7 endpoints across fail-over between ASPs, SGPs, SGs, or IPSPs.

- Support for elimination of message loss of UA-User messages between Signalling Gateways (SGs) and Application Servers (ASs) across fail-over between ASPs, SGPs, SGs, or IPSPs.

1.3.1. Configuration

For carrier-class operation, the SS7 Signalling User Adaptation Layers recommend that Signalling Gateways and Application Servers be configured such that there is no single point of failure within the SG/AS architecture or in the intervening network. The SS7 UAs also recommend that no Application Server be configured for less than two (2) Application Server Processes.

All of the UAs describe an override, loadsharing and broadcast traffic mode. The UA protocols place no restrictions on the distribution algorithm which is used for distributing traffic over multiple Signalling Processes. Additional traffic distribution proposals have been put forward for Load Selection [LOADSEL] and Load Grouping [LOADGRP]

Fail-over between Application Server Processes (ASPs) and Signalling Gateway Processes (SGPs) is not detailed in the UA protocols [M3UA, SUA, TUA], but it is clear that when an SCTP association fails and the ASP transitions to the ASP-DOWN state from the perspective of the SGP peer, the traffic which the associated ASP was previously responsible needs to be diverted to an alternative ASP (if available) in the same Application Server pool.

1.3.2. Conditions at Fail-Over

The details of this diversion of traffic is not specified, however, a dichotomy exists when such fail-over occurs as a result of the loss of an SCTP association between these Signalling Peer Processes (SPPs). When an SPP loses its SCTP association with another SPP, and diverts traffic towards another SPP, there exists the possibility that messages previously destined to the peer SPP exist in several categories, as follows:

Category (1) – Queued in the sending SPP process,

Category (2) – queued for transmission, but not yet transmitted by the transport provider (SCTP),

Category (3) – queued for retransmission, but not yet acknowledged by the peer transport provider (SCTP), and,

Category (4) – acknowledged by the peer transport provider (SCTP) and deleted from the sending transport provider's (SCTP's) retransmission queue.

These categories are illustrated in *Figure 1*. Note that to retransmit categories (2) and (3) (and perhaps categories (1)) on another link requires sent data acknowledgment or buffer retrieval capability by the underlying transport provider.

As there is no SPP peer-to-peer acknowledgement, messages in Categories (3) and (4), the message might or might not have been delivered to the peer SPP. Therefore, at the time of failure of an SCTP association between two Signalling Peer Processes (SPPs), it is not possible for either SPP to determine which of the messages in categories (3) and (4) above (transmitted, but not yet acknowledged; transmitted and acknowledged) were successfully received by the peer before failure. Without information concerning which messages in this category were successfully received by the peer, the SPP either risks message loss or message duplication when it diverts traffic from the failed association.

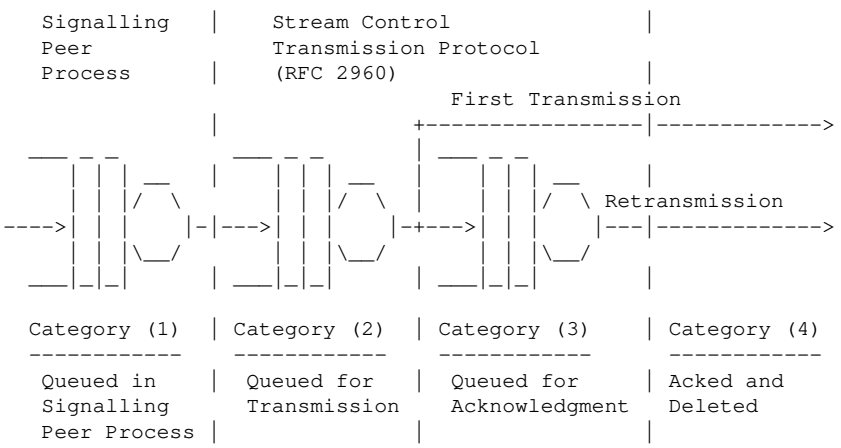

Figure 1. Buffer Categories at SCTP Association Failure

1.3.3. Sources of Message Loss and Duplication

If the messages from category (3) or (4) are retransmitted on an alternative association, the SPP diverting the traffic risks message duplication. This is because some messages of the category might possibly have been successfully received by the peer before fail-over.

If the messages from category (3) and (4) are discarded before diverting messages from categories (1) and (2) and then new traffic on an alternative association, the SPP risks message loss. This is because some of the messages in category (3) and (4) might possibly have *not* been received by the peer SPP before the association failed.

This is the dychotomy: regardless of the nature of a policy concerning the disposition of messages at an SPP experiencing failure to its peer, without information concerning messages successfully received by the peer, the SPP risks message loss or duplication.

It should be possible to induce such a system to demonstrate message loss or duplication.

Because SS7 performance requirements [Q.706] have more stringent requirements against duplication of messages than loss of messages, the only policy is to discard messages in category (3).

To avoid loss of messages to meet SS7 performance requirements [Q.706] in consideration of this dichotomy, implementation cost may be driven higher than would be the case if a procedure were established to exchange information between the Signalling Processes on either side of a failed association.

This Internet-Draft provides Correlation Id and Heartbeat procedures for fail-over for the SS7 signalling UAs which will remove the possibility of message loss or duplication in the event that an SCTP association failure while communication between the Application Server and Signalling Gateway is still possible.

1.3.4. Conditions at Recovery

Figure 2 illustrates an example (A) configuration of ASPs and SGPs. In this example, the ASP and SGP are interconnected with a full-mesh arrangement of SCTP Associations. Each ASP is interconnected to each SGP by an association.

When a failure of the SCTP assocation occurs, it is, for example, between 'SGP1' and 'ASP1' as indicated by the (X) in the *Figure 2*. When a recovery occurs, it is also the SCTP association between 'SGP1' and 'ASP1' that recovers.

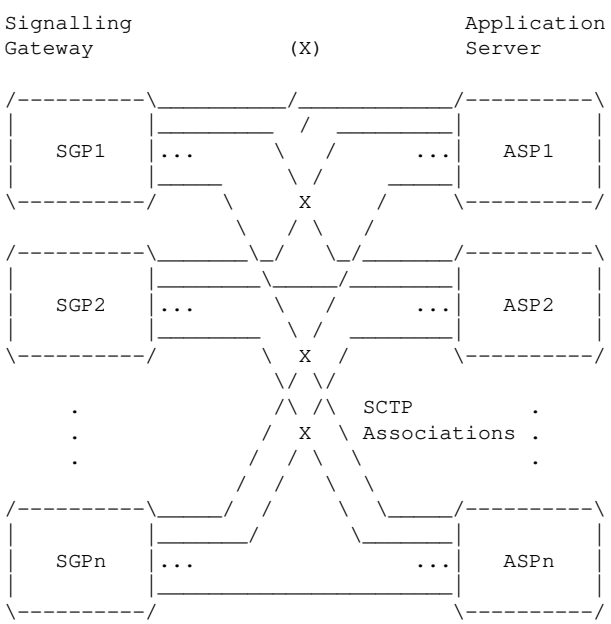

Figure 2. Example (A) Configuration of ASPs and SGPs

The normal procedure for dealing with such a failure[1] is for SGPn to mark ASPn in the ASP-DOWN state and to redirect traffic over the remaining ASPs in the Application Server[2].

When the SCTP association between ASP1 and SGP1 recovers and ASP1 succesfully activates for the AS using the ASP Active Procedures[3], once ASP1 has entered the ASP-ACTIVE state for the AS, message mis-sequencing can occur if traffic is immediately applied on the newly active association.

The UA procedures[3] provide no detail concerning the restarting of traffic to recovering ASPs in the AS.

1.3.5. Sources of Message Mis-Sequencing

Because the SGPs can be experiencing different loads and other local factors each SGP may differ, restoring a traffic flow to a newly active SGP without first ensuring that messages are purged through the old path before the diversion can result in message mis-sequencing. This is illustrated in *Figure 3*.

Before switching traffic back to SGP1 from SGPx, SGPx is queueing traffic from ASP1 to the SS7 network. If the traffic flow from ASP1 is switched rapidly to SGP1, a race condition exists between messages in SGPx's queue and messages in SGP1's queue. A rapid switch can result in mis-sequencing.

As SGPx and SGP1 do not necessarily have to belong to the same SG, close queue synchornization between SGPx and SGP1 cannot be expected.

CORID provides both a time-controlled and a Heartbeat procedure for restoration of traffic to avoid mis-sequencing during restoration.

1.4. Functional Areas

The **CORID** procedures to avoid message loss, duplication and mis-sequencing under these types of scenarios requires a protocol parameters that provide a clear identification of the independent traffic flows involved. Then, procedures are required to control the fail-over and restoration of the identified traffic flows to avoid message loss.

The SS7 MTP3 [Q.704] provide an excellent example of the types of procedures that can be applied to the problem of switching traffic flows across redundant processes[4].

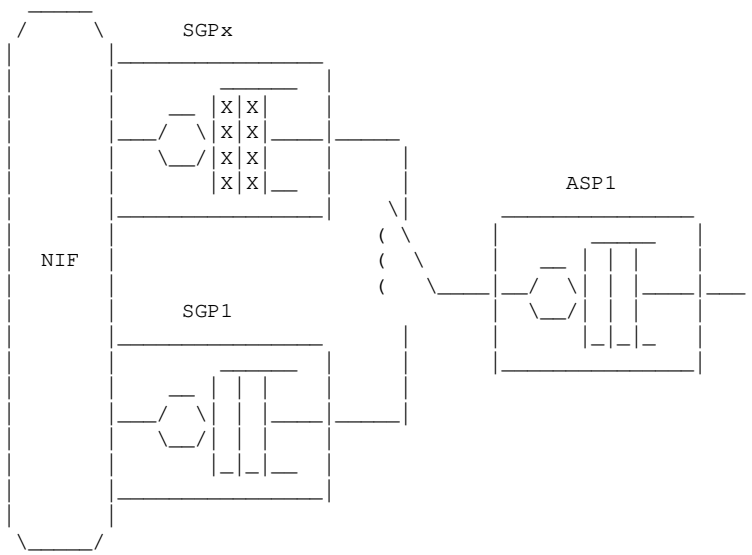

Figure 3. Restoration of a Traffic Flow

1.4.1. Identification of Traffic Flows

Traffic flows between Server Processes in the UAs are managed on the basis of the Application Server to which the traffic flows correspond. Traffic flows from SG to AS are identified by the Routing Key or Routing Context to which they correspond[5].

An Application Server Process can be active and handling traffic for any number or combination of traffic flows. That is, the ASP can be actively handling traffic for any number of Application Servers.

When an SCTP [RFC 2960] association fails, it is necessary to identify both the sequence of the last message successfully received and processed by the Signalling Process, as well as the traffic flow within which that sequence applies.

Therefore, this document identifies a message in a traffic flow by the Routing Context, Load Id, Stream Id and the sequence number within that flow as identified by the Correlation Identifier. The Correlation Identifier is a combination of traffic flow identifier and correlation number which is applied to all divertable traffic.

For details on the assignment of Traffic Flow Identifiers and Correlation numbers, see Section 4.1.2 "Correlation".

1.4.1.1. SGP Starting New SGP-to-ASP Traffic

When traffic is orignally started for a traffic flow the first divertable message in the traffic flow is assigned a correlation number of one (1) by the sending Signalling Process. Subsequent divertable messages within the routing context are given the Correlation Id number of two (2), three (3), and so on.

Because SCTP is a sequenced reliable transport [RFC 2960], it is only necessary to communicate this Correlation Id number between SPP peers for the intial message which is sent to the peer. Each Signalling Peer Process **MUST** be capable of counting the messages which have been sent or received on the SCTP association, and assigning each subsequent message the next sequential Correlation Id number.

1.4.1.2. SPP Diverting peer SPP Traffic

Should, for example, the association fail between the SGP and the ASP, the SGP will recover whatever buffers from categories (1), (2), (3) and (4), and immediately restart traffic, in sequence, on another active ASP within the AS. When the SGP restarts traffic on this alternate ASP, if the messages belong to Category (4) or (3) (i.e, they were transmitted on but not acknowledged by the underlying transport, or transmitted and acknowledged), it will label the initial message sent with the Correlation Id of the message at the time that it was originally sent. When the SGP sends tmessages from Category (2), (1) and newly arriving traffic, the SGP will not tag the messages with a Correlation Id, but instead will label them internally with the next sequential correlation numbers for the traffic flow.

Thus, the alternate Signalling Peer Process which is receiving diverted traffic will be able to distinguish the problematic Category (3) and (4) messages from those which follow. When an tagged message is received, the Signalling Peer Process is now aware that the messages were previously sent to the primary SPP to which the SCTP association was lost. When an untagged message arrives, the receiving Signalling Peer Process is aware that this and subsequent messages within the traffic flow represent previously unsent traffic.

(Detailed procedures for the tagging of messages are described in Section 4.1.3 and 4.1.5.2.1; for diversion, Sections 4.2.2, 4.2.3 and 4.1.6.)

1.4.1.3. SPP Receiving Diverted Traffic

At the Signalling Process receiving the diverted traffic for the Routing Context, three actions are possible (or, variations on the three):

(1) Ignore the Correlation Id and process the messages blind at the risk that message duplication will occur,

(2) discard all messages tagged with a Correlation Id at the risk of increased message loss, or,

(3) perform the procedures described in Section 4.1.5.2.2 minimizing the message duplication and loss resulting from the diversion.

1.4.1.4. SPP Restoring Traffic

Should, for example, the association recover between the SGP and ASP, the ASP will need to rebalance the load across the available SGPs and the newly available SGP. As discussed, if the ASP switches traffic immediately, message mis-sequencing can occur. Two procedures are provided by **CORID** for restoring traffic without message mis-sequencing: a Heartbeat procedure and a timer procedure.

The Heartbeat procedure withholds traffic from the SGP currently active for the traffic flow and sends a **BEAT** message on the flow. Once the **BEAT ACK** is received by the ASP, the ASP is assured that there is no traffic pending on the SGP and

the traffic flow can be switched to the recovered SGP. The Heartbeat procedure is applicable to recovery between SGPs in the same SG.

The Timer procedure witholds traffic from the SGP currently active for the traffic flow and waits until a timer expires. Once the timer expires, the ASP is resonably assured that there is no traffic pending on the SGP and the traffic flow can be switched to the recovered SGP. The Timer procedure is applicable to recovery between SGPs in different SGs.

Restoration of traffic is described in detail in Sections 4.2.3 and 4.1.6.

1.5. Sample Configurations

A typical Example (B) configuration multiple Signalling Gateways is illustrated in *Figure 4*. In this configuration a number of Application Server Processes (ASPs) serving a number of Application Servers (ASs) are connected to two Signalling Gateways (SGs). The SGs appear as mated SS7 Signalling Transfer Points (STPs) [Q.705] to the SS7 Network. Traffic originating at Signalling Endpoints (SEP) in the SS7 network and directed toward SEP in the IP network (i.e., Application Servers) is loadshared over the STPs by the Signalling Link Selection (SLS) [Q.704] value associated with each message. Traffic originating at the SEP in the IP network (i.e, AS) is loadshared over the SGs in the same fashion.

2. Conventions

The keywords **MUST**, **MUST NOT**, **REQUIRED**, **SHALL**, **SHALL NOT, SHOULD**, **SHOULD NOT**, **RECOMMENDED**, **NOT RECOMMENDED**, **MAY**, and **OPTIONAL**, when they appear in this document, are to be interpreted as described in [RFC 2119].

3. Protocol Elements

The following protcol element definitions are provided by **CORID** in extension to the existing protocol element definitions for the UAs [M3UA, SUA, TUA].

3.1. Parameters

The following subsections describe the parameters used for **CORID**, their format and the message in which they are used.

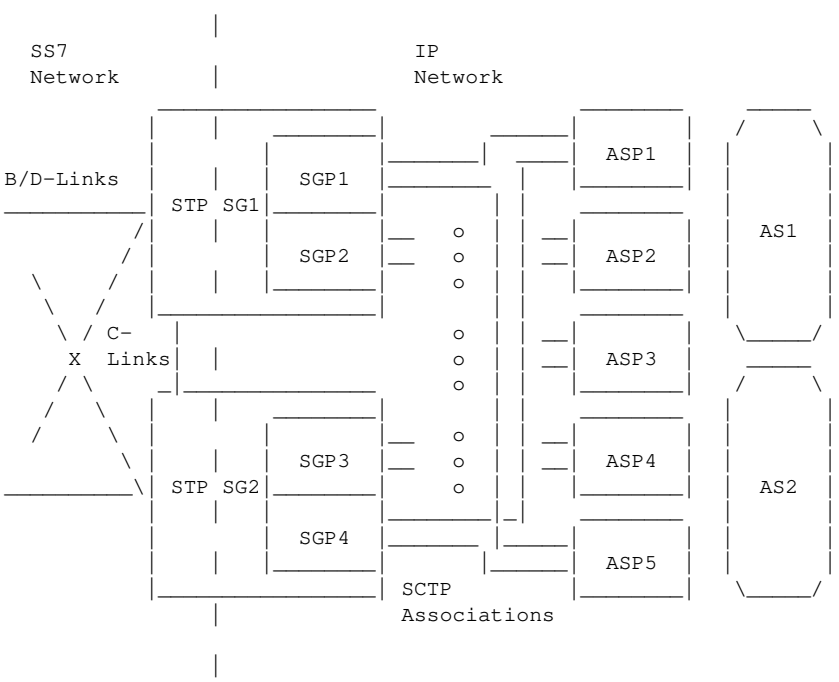

Figure 4. Example (B) Sample Multiple-SG Configuration

3.1.1. Correlation Id

The *Correlation Id* parameter is used in the **ASPAC**, **ASPAC ACK**, and UA-User data messages. It is used to identify data messages sent to a peer SPP.

The *Correlation Id* parameter is formatted as follows:

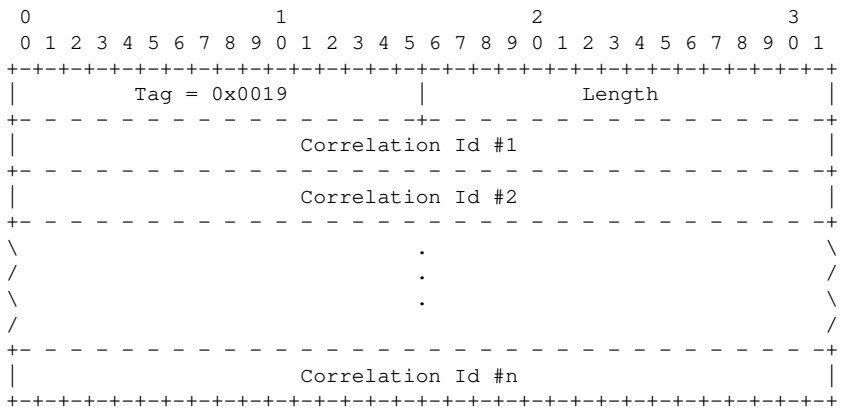

The *Correlation Id* parameter contains one or more of the following field:

Correlation Id field: 8-bytes

The *Correlation Id* field is formatted as follows:

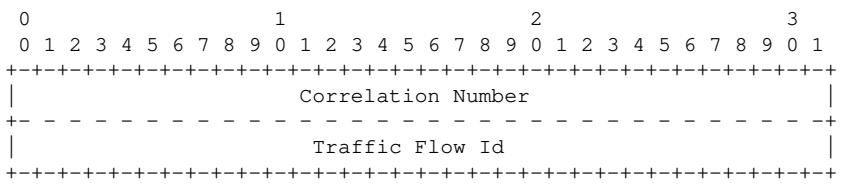

Correlation Number field: 32-bits (unsigned integer)

The *Correlation Number* field identifies a particular message within a traffic flow. When the *Correlation Id* parameter is included in the **ASPAC (ACK)** message, this field identifies the last sent message for the indicated traffic flow. When the *Correlation Id* parameter is included a UA-User data message, this field identifies the correlation number of the message in which it is contained.

Traffic Flow Id field: 32-bits (unsigned integer)

The *Traffic Flow Id* field identifies a particular indepdently sequenced traffic flow to which the *Correlation Number* field value applies. For details on *Traffic Flow Id* assignment, see Section 4.1.2.2. This field is formatted as follows:

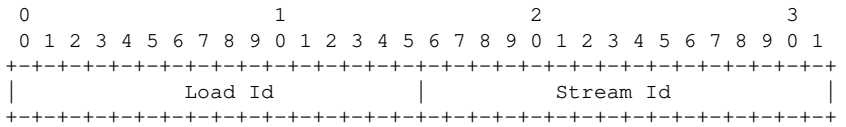

Load Id field: 16-bits (unsigned integer)

The *Load Id* field identifies a load range associated with an SPP. When used for tagging messages or in the **ASPAC (ACK)** message for an Load Selection [LOADSEL], Loadshare AS or Load Group [LOADGRP], the *Load Id* field must identify an SPP (and Load Selector) within an Application Server. For an Override or Broadcast AS, the Load Id is not required and **SHOULD** be coded zero (0). For details on the assignment of *Load Ids*, see Section 4.1.2.2.

Stream Id field: 16-bits (unsigned integer)

The *Stream Id* field contains the SCTP Stream Identifier [RFC 2960] on which the message or referenced message was sent.

When the *Correlation Id* parameter is included in the **ASPAC**, **ASPAC ACK**, and UA-User data messages, only one Routing Context representing a single Application Server **MUST** be associated (specified or implied) with the message.

3.2. Messages

3.2.1. ASP Active (ASPAC)

CORID supplements the **ASPAC** mesage by permitting the following optional parameters to be included in the message:

Extension Parameters	
Correlation Id	Mandatory

The format of the resulting **ASPAC** message is as follows:

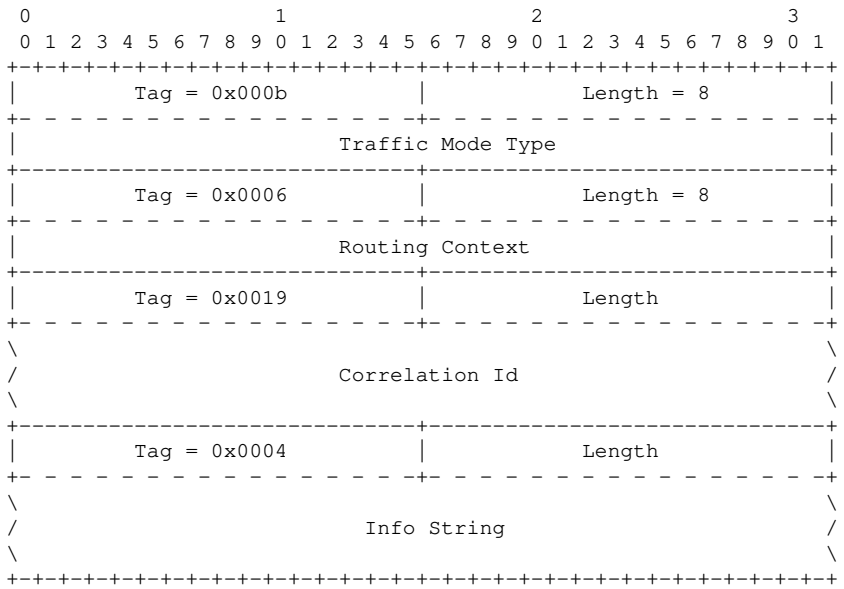

No other changes to the **ASPAC** message format are provided by this extension.

The *Correlation Id* parameter is used by the ASP in the **ASPAC** message to indicate the correlation identifier for the first UA-User message to be transmitted in each traffic flow from the Application Server being activated to the Signalling Gateway. The Application Servers for which the *Correlation Id* apply is either indicated in the **ASPAC** message by providing the associated *Routing Contexts*, or, if there is no *Routing Context* parameter in the message, the associated Application Servers are implied by the SGP and ASP configuration data.

When the *Correlation Id* parameter is present in the **ASPAC** message, the message **SHOULD** only contain one *Routing Context* in the *Routing Context* parameter. When the *Correlation Id* parameter is not present, but required by the SGP, the value of the correlation id is assumed to be zero (0).

The **ASPAC** message **MAY** contain additional extension parameters provided for by other extensions.

3.2.2. ASP Active Acknowledgement (ASPAC ACK)

CORID supplements the **ASPAC ACK** mesage by permitting the following optional parameters to be included in the message:

Extension Parameters	
Correlation Id	Mandatory

The format of the resulting **ASPAC ACK** message is as follows:

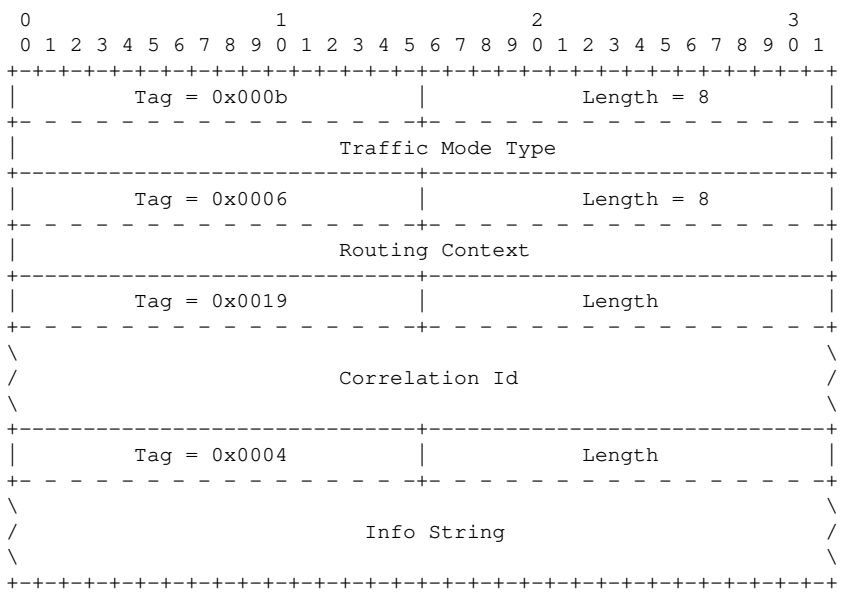

No other changes to the **ASPAC ACK** message format are provided by this extension.

The *Correlation Id* parameter is used by the SGP in the **ASPAC ACK** message to indicate the correlation identifier for the first UA-User message to be transmitted from the Signalling Gateway to the Application Server being activated for each traffic flow. The Application Servers for which the *Correlation Id* apply is either indicated in the **ASPAC ACK** message by providing the associated *Routing Contexts*, or, if there is no *Routing Context* parameter in the message, the associated Application Servers are implied by the SGP and ASP configuration data.

When the *Correlation Id* parameter is present in the **ASPAC ACK** message, the message **SHOULD** only contain one *Routing Context* in the *Routing Context* parameter. When the *Correlation Id* parameter is not present, but required by the ASP, the value of the correlation id is assumed to be zero (0).

The **ASPAC ACK** message **MAY** contain additional extension parameters provided for by other extensions.

4. Procedures

4.1. Traffic Handling

4.1.1. Classification

Divertable messages are any UA-User messages destined for an Application Server. Divertable messages are UA-User data and some management (non-ASP management) messages that have an explicit or implied Routing Context and have strict requirements preventing loss, duplication or mis-sequencing. SNMM messages do not quality as divertable because they can apply to more than on Application Server. UA-User messages that qualify as divertable messages are listed in *Table 1*. Although some messages in some message classes might be considered as non-divertable, all messages in the message classes listed in *Table 1* **SHALL** be treated as divertable.

Table 1. Divertable Messages by UA

UA	Class	Msg	Description
M3UA	Transfer	DATA	
SUA	Connection-less	CLDT CLDR	Marked Return on Error
	Connection-Oriented	CORE COAK CODT RESRE RESCO RELRE	
TUA	Dialogue Handling	TUNI TQRY TCNV TRSP	w/o components or marked Return on Error
		TUAB TPAB TNOT	
	Component Handling	CINV CRES	Operation Class 1, 2 or 3
		CERR CREJ CCAN	

4.1.2. Correlation

Each independent traffic flow for a given Application Server as identified by a Routing Context **MUST** be correlated using a Correlation Id. The Correlation Id consists of a correlation number and a traffic flow identifier. The correlation number is used to number each message within the given traffic flow.

4.1.2.1. Assignment of Correlation Ids

To accomodate all combinations of traffic modes at AS and SG, divertable messages are correlated by independent traffic flow. That is, each sent divertable message is labelled with a traffic flow identifier and a correlation number for the AS that is incremented for each message sent for the traffic flow. In the same fashion, each received divertable message is labelled with the identity of the traffic flow on which it was received and a correlation number for the AS that is incremented for each message received on that traffic flow.

An SPP maintains two correlation counters for each traffic flow for each AS: for each traffic flow, one counter tracks the correlation number of messages sent to the AS and the other tracks the correlation number of messages received from the AS. Before traffic is started for an AS on a traffic flow, these counters are set to zero (0). The first divertable message for the AS on the flow **MUST** then be assigned a coorrelation number of one (1); and subsequent divertable messages, the correlation number of two (2), three (3), and so forth.

Whenever traffic is started for the AS (using the ASP Active Procedures), the correlation counters **SHALL** be synchronized by exchanging correlation numbers and traffic flow identifiers in the **Correlation Id** parameter in the **ASPAC** and **ASPAC ACK** messages. For new traffic, the correlation number **MUST** zero (0); for restarting traffic, it is **SHOULD** be the correlation number of the last message transferred. (See Section 4.2.3.)

4.1.2.2. Assignment of Traffic Flow Ids

Traffic flow identifiers **SHALL** consist of two components:

(i) A *Load Id* component that identifies a switchable traffic load pattern within an Application Server. This component **SHALL** be assigned by the peer SPP.

(ii) A *Stream Id* component that identifies the SCTP stream upon which a message is sent or received. This component **SHALL** be assigned by the sending SPP and **MUST** correspond to the SCTP Stream upon which the message was sent.

Load Ids **SHALL** assigned by an SPP and **MUST** be communicated to the peer SPP in an **ASPAC** or **ASPAC ACK** message. For traffic distributions that do not loadshare (i.e, Override and Broadcast), the load flow identifier is not required and **MAY** be set to zero (0). Following are rules for the assignment of load identifiers at an SPP:

(i) If an SPP belongs to a regular Override or Broadcast AS, no Load Id need be assigned or included by the SPP in the *Correlation Id* parameter.

(ii) If an SPP belongs to a regular Loadshare AS, a Load Id is assigned and included in the *Correlation Id* parameter. The Load Id assigned **MUST** unambiguously identify the SPP within the AS.

(iii) If an SPP belongs to a Load Selector [LOADSEL], a Load Id is assigned and included in the *Correlation Id* parameter regardless of the Traffic Mode Type of the AS. The Load Id assigned **MUST** unambiguously identify the SPP and the Load Selector within the AS.

(iv) If an SPP belongs to a Load Group [LOADGRP], a Load Id is assigned and included in the *Correlation Id* for a Loadshare AS or Load Group. An assigned Load Id **MUST** unambiguously identify the SPP and the Load Selector within the AS. For a non-loadshare AS and Load Group, no Load Id need be assigned or included in the *Correlation Id* parameter.

4.1.3. Tagging

Each sent or received message for an AS is labelled when it is first sent or received. The message is labelled with the traffic flow id associated with the SPP to or from which the message was sent or received, and the correlation number assigned within the traffic flow (see Section 4.1.2.1).

Tagged messages contain a *Correlation Id* parameter: an untagged message is tagged by adding a *Correlation Id* parameter to the message. When a message is tagged, it **SHALL** be tagged with the same values of the traffic flow id (if required) and correlation number with which it was originally labelled.

Although each message is labelled with a traffic flow id and correlation number, the message is not necessarily tagged with the *Correlation Id* parameter when the message is sent. Messages for an AS that are sent for the first time **MUST NOT** be tagged. Messages retransmitted **MUST** be tagged.

4.1.4. Buffering

4.1.4.1. SPP witholding unsent messages

CORID procedures require that an SPP at times withhold AS traffic. To perform this, the SPP allocates a diversion buffer and places in the buffer all subsequent messages that would otherwise be sent to the SPP for the AS into the buffer.

4.1.4.2. Local copies of sent messages

To reduce loss of messages, an SPP **SHOULD** buffer messages until it can be assured that the peer SPP has received and processed the message. When a message is sent to an SPP supporting **CORID** a local copy of the message **MUST** be kept until it is discarded in accordance with a **CORID** procedure.[6]

(i) A local copy **SHOULD NOT** be discarded when it is acknowledged by the peer SCTP.

(ii) a local copy **SHOULD NOT** be discarded until the sending SPP is confident that the peer SPP has received and processed the message.

(iii) To ensure that stale messages do not propagate through the system, an SPP **SHOULD NOT** keep local copies of sent messages for longer than a maximum lifetime T(lifetime). Any local coppies of sent messages that are older (measured from the moment at which they were sent to the peer SPP) than T(lifetime) **SHOULD** be discarded.

4.1.5. Message Handling

4.1.5.1. Untagged Messages

4.1.5.1.1. SPP sending untagged messages

An SPP sends untagged messages to a peer SPP whenever the message is being sent for an Application Server for the first time. All divertable messages which have been transmitted for the first time **MUST NOT** be sent tagged.

Local copies of untagged messages awaiting acknowledgement or expiry are labelled with the Routing Context for the Application Server to which they were sent, the traffic flow id of the SPP to which they were sent, and the correlation number of the message. The correlation number with which a message is labelled **MUST** be the next sequential correlation number for the AS and traffic flow. These labels can be used later to tag a message that is marked for diversion.

4.1.5.1.2. SPP receiving untagged messages

When an SPP receives an untagged message, it associated with the message the next sequential correlation number for the Routing Context and traffic flow id for which the message was received. Untagged messages are received in order and **MAY** be processed when received. The SPP **SHOULD** keep track of the Correlation Ids that have been processed for the AS.

4.1.5.2. Tagged Messages

4.1.5.2.1. SPP sending tagged messages

An SPP sends tagged traffic whenever it sends traffic that is marked for diversion. That is, whenever an SPP sends divertable messages to an SPP other than the original SPP for which those messages were labelled, the SPP **MUST** tag the message with the *Correlation Id* parameter that contains the labelled traffic flow id (if required) and correlation number.

In addition, when a ASP becomes active for a Broadcast AS, an SGP **MUST** tag the first message in each traffic flow towards the ASP to allow the ASP to synchronize its entry into the Broadcast AS.

4.1.5.2.2. SPP receiving tagged messages

The handling of tagged messages is the mechanism that provides for the reduction of message loss, duplication and missequencing. An SPP receiving divertable messages containing a *Correlation Id* parameter **SHALL** perform the following actions:

(i) The SPP determines (by implementation-dependent means [7]) whether the message has already been processed for the AS.

(ii) If the message has not already been processed for the AS, it is processed as normal.

(iii) If the message has already been processed for the AS, it is discarded.

(iv) If, as a result of some failure, the SPP cannot determine with any certainly whether the tagged message has been processed for the AS, or not, the SPP **MUST** discard the message[8].

4.1.6. Diversion

4.1.6.1. SPP diverting traffic from a failed, deactivated or overridden peer SPP

4.1.6.1.1. Alternate SPP in same AS or SG, or No Alternate SPP

When an SPP diverts AS traffic away from a failed, deactivated or overridden peer SPP to an alternate peer SPP in the same AS or SG, the SPP **SHALL** perform the following actions:

(i) The SPP tags (see Section 4.1.3) each untagged message that is marked for diversion.

(ii) If an alternate SPP is available (active for the AS), the SPP sends the messages marked for divertion to the alternate SPP.

(iii) If no alternate SPP exists (the AS is AS-PENDING), the SPP buffers the marked messages in a buffer used for buffering messages while the AS is in the AS-PENDING state.

(iv) The SPP then diverts AS traffic, beginning with traffic withheld for the AS, to the alternate SPP or AS-PENDING buffer.

This procedure corresponds to the Sequenced Changeover procedure used by the SS7 MTP [Q.704].

4.1.6.1.2. Alternate SPP in different AS or SG

When an SPP diverts AS traffic away from a failed or deactivated peer SPP to an alternate peer SPP in a different AS or SG, the SPP **SHALL** perform the following actions:

(i) The SPP starts timer T(divert) and continues buffering AS traffic until the timer expires.

(ii) When T(divert) expires, and the failed or deactivated SPP has not recovered, the SPP continues with the following actions:

(iii) The SPP discards all tagged messages and messages marked for diversion.

(iv) The SPP starts AS traffic, beginning with the contents of the diversion buffer, to the alternate SPP.

This procedure corresponds to the Time-Controlled Changeover procedure used by the SS7 MTP [Q.704].

4.1.6.2. SPP diverting traffic from an active peer SPP

When an SPP wishes to divert AS traffic away from an active peer SPP, the SPP **SHALL** perform the following actions:

(i) The SPP witholds and buffers AS traffic for the SPP from which the traffic is being diverted.

(ii) The SPP sends a **BEAT** message with a unique identifier[9] in the *Heartbeat Data* parameter on the same SCTP stream(s) on which the traffic being withheld for diversion was previously sent.

(iii) The SPP starts a timer T(restore).

(iv) If the SPP receives the **BEAT ACK** message(s) that contain the unique identifier(s) in the *Heartbeat Data* parameter before timer T(restore) expires, the SPP diverts the traffic, beginning with the withheld traffic, to the target SPP and cancels the T(restore) timer.

(v) If the timer T(restore) expires, the diverting SPP diverts traffic, beginning with the withheld traffic, to the target SPP.

(vi) If an SPP receives a **BEAT ACK** message containing a unique identifier for which the timer T(restore) has already expired, the SPP ignores the message.

The purpose of this **BEAT** procedure is to avoid mis-sequencing by ensuring that all messages sent for the AS to the old SPP arrives before messages are sent to the new SPP. This avoids races between (and possible mis-sequencing of) messages sent on the old SPP and messages sent on the new SPP.

This procedure corresponds to the Changeback procedure used by the SS7 MTP [Q.704].

4.2. ASP Management Procedures

4.2.1. ASP Down Procedures

4.2.1.1. SPP detecting loss of SCTP association

When an SPP receives an SCTP COMMUNICATION LOST or RESTART indication and there are Application Servers active for the association, the SPP **SHALL** perform the following actions with regard to active AS traffic for the association:

(i) The SPP witholds AS traffic for the peer SPP in a diversion buffer.

(ii) The SPP marks for diversion all local copies of AS messages already sent to the peer SPP.

(iii) The SPP then **SHALL** perform the actions described in Section 4.1.7.1.

4.2.1.2. ASP sending ASPDN

An ASP **MUST NOT** send an **ASPDN** message until it has completed the ASP Inactive Procedures with the intended SGP for every AS.

4.2.1.3. SGP or IPSP receiving ASPDN

An SGP or IPSP, upon rreceiving an **ASPDN** message from an ASP-ACTIVE ASP, **MUST** perform the ASP Inactive Procedures with regard to **CORID** (see Section 4.2.2.2) for every AS for which the ASP is ASP-ACTIVE and then complete the **ASPDN** procedures.

4.2.1.4. ASP receiving ASPDN ACK

An SGP or IPSP, upon receiving an unsolicited **ASPDN ACK** message from an active SGP, **MUST** perform the ASP Inactive Procedures with regard to **CORID** (see Section 4.2.2.3) for every AS for which the ASP is ASP-ACTIVE and then complete the **ASPDN ACK** procedures.

4.2.2. ASP Inactive Procedures

4.2.2.1. ASP sending ASPIA

When an ASP wishes to deactivate an Application Server with an SGP, the ASP **SHALL** perform the following actions for traffic pertaining to the AS:

(i) The ASP withholds sending AS traffic to the SGP or IPSP.

(ii) The ASP stops processing AS traffic recevied from the SGP or IPSP. Any messages received for the Application Server after the last processed message **MAY** be discarded.

(iii) The ASP starts a T(divert) timer.

(iv) The ASP **SHALL** perform the applicable UA ASP Inactive Procedures[10].

4.2.2.2. SGP receiving ASPIA or sending ASPIA ACK

An SGP receiving an **ASPIA** message for an AS, or wishing to send an unsolicited **ASPIA ACK** to deactivate an AS, **SHALL** perform the following actions for the traffic pertaining to each AS for which deactivation is performed:

(i) The SGP withholds sending AS traffic to the ASP.

(ii) The SGP stops processing AS traffic received from the ASP. Any messages received for the AS at the SGP after receiving the **ASPIA** message **MUST** be discarded.

(iii) The SGP marks for diversion all local copies of AS messages sent to the ASP.

(iv) The SGP then **SHALL** perform the actions described in Section 4.1.7.1.

(v) The ASP **SHALL** perform the applicable UA ASP Inactive Procedures[10].

4.2.2.3. ASP receiving ASPIA ACK

Upon receiving an **ASPIA ACK** message the ASP **SHALL** perform the following actions for the traffic pertaining to the AS identified by the *Routing Context* in the received **ASPIA ACK** message or implied by the SCTP association on which the **ASPIA ACK** message was received:

(i) The T(divert) timer is cancelled (if running).

(ii) The ASP marks for diversion any local copies of AS messages sent to the SGP.

(iii) The ASP then **SHALL** perform the actions described in Section 4.1.7.1.

(iv) The ASP **SHALL** perform the applicable UA ASP Inactive Procedures[10].

4.2.2.4. T(divert) timer expiry

If the T(divert) timer expires before receiving an **ASPIA ACK** for the AS, the ASP **SHALL** perform the actions described in Section 4.2.2.3.

4.2.3. ASP Active Procedures

4.2.3.1. ASP sending ASPAC

When an ASP wishes to activate an Application Server for an SGP, the ASP **SHALL** perform the following actions for traffic pertaining to the AS:

(i) The ASP determines the correlation id of the last message sent to this SGP for the AS for each traffic flow.

(ii) If the ASP has not sent a message to the SGP for the traffic flow, the correlation id zero (0) is used.

(iii) If the ASP has sent messages to the SGP for the traffic flow, but cannot determine the correlation id of the last message sent due to local failure, the correlation id zero (0) is used.

(iv) The ASP includes the correlation id(s) determined above in the *Correlation Id* parameter in the **ASPAC** message used to active the AS. (See Section 3.1.1.)

(v) The ASP **SHALL** perform the applicable UA ASP Active Procedures[11].

4.2.3.2. SGP receiving ASPAC

When an SGP receives an **ASPAC** message for an Application Server, the SGP **SHALL** perform the following actions with regard to traffic for the AS:

(i) The SGP sets the correlation id of the next received message from the ASP for each traffic flow to the value, contained in the *Correlation Id* parameter in the **ASPAC ACK** message, plus one (1).

(ii) The SGP determines the correlation id of the last message sent to this SGP for each traffic flow.

(iii) If the SGP has not sent a message to the ASP for a traffic flow, the correlation id zero (0) is used.

(iv) If the SGP has sent messages to the ASP for a traffic flow, but cannot determine the correlation id of the last message sent due to local failure, the correlation id zero (0) is used.

(v) The SGP includes the correlation id(s) determined above in the *Correlation Id* parameter in the **ASPAC ACK** message used to acknowledge activation of the AS. (See Section 3.1.1.)

(vi) The ASP **SHALL** perform the applicable UA ASP Active Procedures[11], including the sending of **ASPIA ACK**.

(vii) The ASP then **SHALL** perform the actions described in Section 4.1.7.2.

4.2.3.3. ASP receiving ASPAC ACK

When an ASP receives an expected **ASPAC ACK** message for an Application Server, the ASP **SHALL** perform the following actions with regard to AS traffic:

(i) The ASP sets the correlation id of the next received message from the SGP for each traffic flow to the value, contained in the *Correlation Id* parameter in the **ASPAC ACK** message, plus one (1).

(ii) The ASP **SHALL** perform the applicable UA ASP Active Procedures[11].

(iii) The ASP then **SHALL** perform the actions described in Section 4.1.7.2.

If an ASP receives an unexpected **ASPAC ACK** (i.e, one for which no ASPAC was sent and the ASP is already in the ASP-ACTIVE state for the AS), then the ASP **SHALL** ignore the message for the purposes of **CORID**. The ASP **SHALL**, however, perform the applicable UA ASP Active Procedures[11].

4.3. Interworking Procedures

Because the **CORID** procedures provided here rely upon close synchronization of correlation identifiers between SPP, if one of the SPP does not support these **CORID** procedures, neither SPP is able to take advantage of the full benefits of the procedures. The SPP supporting **CORID MAY** fall back to the interworking procedures provided in this section, or to procedures based on the original (non-**CORID**) UA procedures.

A peer SPP that does not support the **CORID** procedures can either be identified by local configuration information, the ASP Extenstions [ASPEXT] procedure, or at ASP Activation time. The lack of support for **CORID** can be determined at ASP Activation time when the peer SPP does not place a **Correlation Id** parameter (as it **MUST** if both peers support **CORID**) in the **ASPAC (ACK)** message.

When interworking to an SPP that does not support **CORID**, the SPP supporting **CORID SHALL** perform all of the procedures as though the peer SPP supported **CORID** with the following exceptions:

(i) The SPP **MUST NOT** send messages marked for diversion and tagged to the peer SPP not supporting **CORID**. All such messages **MAY** be discarded.

(ii) When diverting traffic between a failed, deactivated or overriden peer SPP and an alternate peer SPP not supporting **CORID**, the actions described in Section 4.1.7.1.2 **MUST** always be used instead of the procedures in Section 4.1.7.1.1, except when there is no alternate SPP.

(iii) The SPP **MUST NOT** place a *Correlation Id* parameter in the **ASPAC** or **ASPACK ACK**. So, the actions described in Sections 4.2.3.1(i)-(iv), 4.2.3.2(i)-(v) and 4.2.3.3(i)-(ii) do not apply.

5. Examples

5.1. Example Configuration

5.2. Initialization

Figure 5 illustrates the initialization sequence that is used for all of the examples .

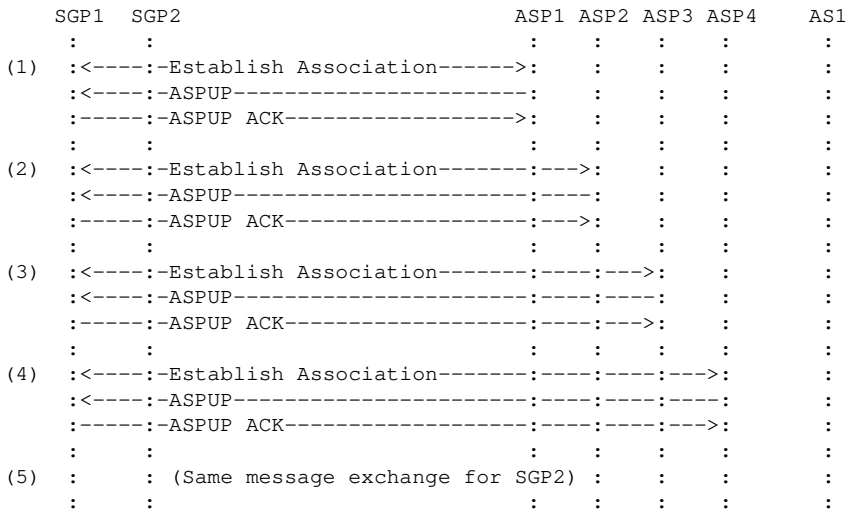

Figure 5. Example – Starting Traffic

The sequence of events in the exmaple illustrated in *Figure 5* is as follows:

(1) ASP1 establishes an SCTP association to SG1 and zents

(2)

(3)

(4)

5.3. Starting Traffic

Figure 6 illustrates

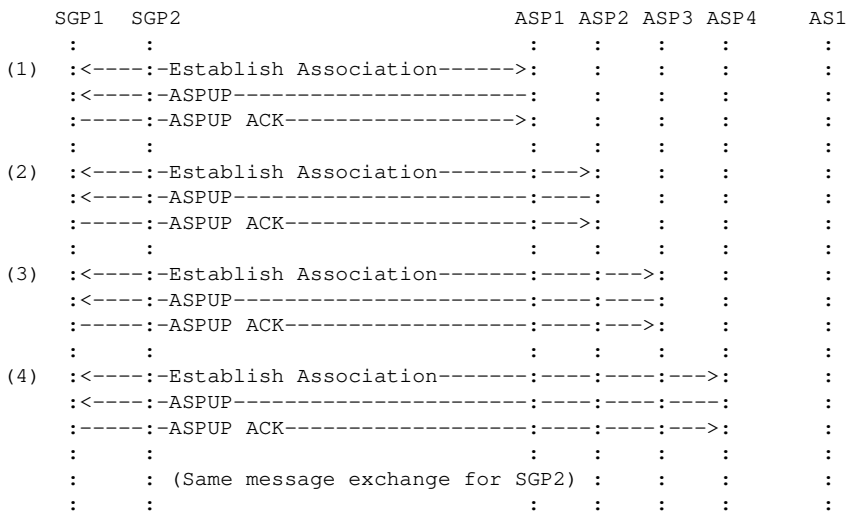

Figure 6. Example – Starting Traffic

The sequence of events in the exmaple illustrated in *Figure 6* is as follows:

(1)

(2)

(3)

(4)

(5)

(6)

(7)

5.3.1. Initial Startup

5.3.2. Joining a Broadcast

5.4. Fail-Over

5.4.1. Assocation Failure – Override

5.4.2. Deactivation – Loadshare

5.4.3. Management Blocking – Override

5.5. Recovery

5.5.1. Association Recovery – Loadshare

5.5.2. AS-Pending Recovery

5.6. Interworking

5.6.1. ASP does not Support CORID

6. Security

CORID does not introduce any new security risks or considerations that are not already inherent in the UA [M3UA, SUA, TUA] Please see the "Security" sections of M3UA, SUA and TUA [M3UA, SUA, TUA] for security considerations and recommendations that are applicable to each of these UAs.

7. IANA Considerations

CORID redefines the format of the *Correlation Id* parameter for M3UA, SUA and TUA. **CORID** also redifines the **ASPAC** and **ASPAC ACK** messages to include the *Correlation Id* parameter as a mandatory parameter of those messages.

8. Timers

Following are the **RECOMMENDED** timer values:

T(divert)	0.5-2 seconds
T(restore)	0.5-2 seconds

Acknowledgments

The authors would like to thank Ken Morneault, Greg Sidebottom, John Loughney, Sandeep Mahajan, Barry Nagelberg, and Nitin Vairagare for their valuable comments and suggestions.

Notes

[1] As described in the UA documents.

[2] For illustration purposes only, all ASPs in *Figure 2* are members of the one Application Server which is represented at all of the SGPs.

[3] See Section 4.3.4.3 of M3UA, SUA or TUA [M3UA, SUA, TUA].

[4] See, for example, Clause 5 "Changeover", Clause 6 "Changeback", Clause 7 "Forced Rerouting" and Clause 8 "Controlled Rereouting" of the MTP3 specifications [Q.704].

[5] This is true for all User Adaptation layers with the exception of M2UA [M2UA]. In M2UA, the Application Server and traffic flows are identified by an equivalent of the Routing Context: the Interface Identifier. An Application Server may also represent multiple Interface Identifiers.

[6] **IMPLEMENTATION NOTE:–** A simple way to meet the requirements for keeping local copies of messages is to keep a local copy of all messages sent to an SPP supporting **CORID** until a fixed buffer allocation is exceeded, or until the local copy lifetime expires. T(lifetime) and buffer capacity can then be adjusted to ensure that local copies of messages are not discarded too early resulting in message loss during fail-over.

[7] **IMPLEMENTATION NOTE:–** Determining which messages have already been processed for the AS may require some ASP-to-ASP or SGP-to-SGP synchronization that is outside the scope of the UA documents [M3UA, SUA, TUA] and also outside the scope of this document.

If the received traffic flow id matches that of the SPP on which the message was received, this might be a simple matter of comparing the correlation number of the message to the correlation number of the last message processed for the Application Server.

[8] **IMPLEMENTATION NOTE:–** The reason for discarding tagged messages at the receiver for which it cannot be determined with any certainty whether the message was processed for the AS or not is because, for SS7, message loss is preferrable to message duplication [Q.706].

[9] **IMPLEMENTATION NOTE:–** Although the unique identifier placed in the *Heartbeat Data* is implementation dependent, a useful identifier would be the tuple formed by the Routing Context, Correlation Id corresponding to the last message sent to the SPP from which traffic is to be diverted.

[10] For the "ASP Inactive Procedures", see Section 4.3.4.4 of M3UA, SUA, and TUA. [M3UA, SUA, TUA]

[11] For the "ASP Active Procedures", see Section 4.3.4.3 of M3UA, SUA, and TUA. [M3UA, SUA, TUA]

References

RFC 2960.
R. Stewart, Q. Xie, K. Morneault, C. Sharp, H. J. Schwarzbauer, T. Taylor, I. Rytina, H. Kalla, L. Zhang and V. Paxson, “Stream Control Transmission Protocol (SCTP),” RFC 2960, The Internet Society (February 2000).

M3UA. G. Sidebottom, J. Pastor-Balbes, I. Rytina, G. Mousseau, L. Ong, H. J. Schwarzbauer, K. Gradischnig, K. Morneault, M. Kalla, N. Glaude, B. Bidulock and N. Glaude, “SS7 MTP3-User Adaptation Layer (M3UA),” <draft-ietf-sigtran-m3ua-10.txt>, Internet Engineering Task Force - Signalling Transport Working Group (November, 2001). Work In Progress.

SUA. J. Loughney, G. Sidebottom, G. Mousseau, S. Lorusso, L. Coene, G. Verwimp, J. Keller, F. E. Gonzalez, W. Sully, S. Furniss and B. Bidulock, “SS7 SCCP-User Adaptation Layer (SUA),” <draft-ietf-sigtran-sua-09.txt>, Internet Engineering Task Force - Signalling Transport Working Group (June 15, 2001). Work In Progress.

TUA. B. Bidulock, “SS7 TCAP-User Adaptation Layer (TUA),” <draft-bidulock-sigtran-tua-00.txt>, Internet Engineering Task Force - Signalling Transport Working Group (January 2002). Work In Progress.

Q.704. ITU, “Message Transfer Part – Signalling Network Functions and Messages,” ITU-T Recommendation Q.704, ITU-T Telecommunication Standardization Sector of ITU, Geneva (March 1993). (Previously "CCITT Recommendation")

M2UA. K. Morneault, R. Dantu, G. Sidebottom, T. George, B. Bidulock and J. Heitz, “SS7 MTP2-User Adaptation Layer (M2UA),” <draft-ietf-sigtran-m2ua-11.txt>, Internet Engineering Task Force - Signalling Transport Working Group (November, 2001). Work In Progress.

Q.700. ITU, “Introduction to CCITT Signalling System No. 7,” ITU-T Recommendation Q.700, ITU-T Telecommunication Standardization Sector of ITU, Geneva (March 1993). (Previously "CCITT Recommendation")

LOADSEL.
B. Bidulock, “Load Selection Extension for Signalling User Adaptation Layers (LOADSEL),” <draft-bidulock-sigtran-loadsel-00.txt>, Internet Engineering Task Force - Signalling Transport Working Group (January 2002). Work In Progress.

LOADGRP.
B. Bidulock, “Load Grouping Extension for Signalling User Adaptation Layers (LOADGRP),” <draft-bidulock-sigtran-loadgrp-00.txt>, Internet Engineering Task Force - Signalling Transport Working Group (January 2002). Work In Progress.

Q.706. ITU, “Signalling System No. 7 – Message Transfer Part Signalling Performance,” ITU Recommendation Q.706, ITU-T Telecommunication Standardization Sector of ITU, Geneva (March 1993). (Previously "CCITT

Recommendation")

Q.705. ITU, “Signalling System No. 7 – Signalling Network Structure,” ITU-T Recommendation Q.705, ITU-T Telecommunication Standardization Sector of ITU, Geneva (March 1993). (Previously "CCITT Recommendation")

RFC 2119.
S. Bradner, “Key words for use in RFCs to Indicate Requirement Levels,” RFC 2119 - BCP 14, Internet Engineering Task Force (March 1997).

ASPEXT.
B. Bidulock, “Application Server Process (ASP) Extension Framework,” <draft-bidulock-sigtran-aspext-00.txt>, Internet Engineering Task Force - Signalling Transport Working Group (January 2002). Work In Progress.

Author’s Addresses

Brian Bidulock
OpenSS7 Corporation
4701 Preston Park Boulevard
Suite 424
Plano, TX 75093
USA

Phone: +1-972-839-4489
Email: bidulock@openss7.org
URL: http//www.openss7.org/

This draft expires July, 2002.

List of Tables

Table 1 Divertable Messages by UA 11

List of Illustrations

Figure 1 Buffer Categories at SCTP Association Failure 3
Figure 2 Example (A) Configuration of ASPs and SGPs 4
Figure 3 Restoration of a Traffic Flow 5
Figure 4 Example (B) Sample Multiple-SG Configuration 7
Figure 5 Example – Starting Traffic 17
Figure 6 Example – Starting Traffic 17

Table of Contents

1 Introduction 2
1.1 Scope 2
1.2 Terminology 2
1.3 Overview 2
1.3.1 Configuration 3
1.3.2 Conditions at Fail-Over 3
1.3.3 Sources of Message Loss and Duplication 4
1.3.4 Conditions at Recovery 4
1.3.5 Sources of Message Mis-Sequencing 5
1.4 Functional Areas 5
1.4.1 Identification of Traffic Flows 6
1.5 Sample Configurations 7
2 Conventions 7
3 Protocol Elements 7
3.1 Parameters 7
3.1.1 Correlation Id 8
3.2 Messages 9
3.2.1 ASP Active (ASPAC) 9
3.2.2 ASP Active Acknowledgement (ASPAC ACK) 9
4 Procedures 10
4.1 Traffic Handling 10
4.1.1 Classification 10
4.1.2 Correlation 11
4.1.3 Tagging 12
4.1.4 Buffering 12
4.1.5 Message Handling 12
4.1.6 Diversion 13
4.2 ASP Management Procedures 14
4.2.1 ASP Down Procedures 14
4.2.2 ASP Inactive Procedures 14
4.2.3 ASP Active Procedures 15
4.3 Interworking Procedures 16
5 Examples 16
5.1 Example Configuration 16
5.2 Initialization 16

5.3 Starting Traffic 17
5.3.1 Initial Startup 18
5.3.2 Joining a Broadcast 18
5.4 Fail-Over 18
5.4.1 Assocation Failure – Override 18
5.4.2 Deactivation – Loadshare 18
5.4.3 Management Blocking – Override 18
5.5 Recovery 18
5.5.1 Association Recovery – Loadshare 18
5.5.2 AS-Pending Recovery 18
5.6 Interworking 18
5.6.1 ASP does not Support CORID 18
6 Security 18
7 IANA Considerations 18
8 Timers 18

Copyright Statement

Copyright © The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedure for copyrights defined in the Internet Standards process must be followed, or as required to translate into languages other than English.

The limited permission granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and **THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MECHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.**