

Wireless Transaction Protocol

Version 27 Aug 2002

Open Mobile Alliance
WAP-224-WTP-20020827-a

Continues the Technical Activities
Originated in the WAP Forum



Use of this document is subject to all of the terms and conditions of the Use Agreement located at <http://www.openmobilealliance.org/UseAgreement.html>.

Unless this document is clearly designated as an approved specification, this document is a work in process, is not an approved Open Mobile Alliance™ specification, and is subject to revision or removal without notice.

You may use this document or any part of the document for internal or educational purposes only, provided you do not modify, edit or take out of context the information in this document in any manner. Information contained in this document may be used, at your sole risk, for any purposes. You may not use this document in any other manner without the prior written permission of the Open Mobile Alliance. The Open Mobile Alliance authorizes you to copy this document, provided that you retain all copyright and other proprietary notices contained in the original materials on any copies of the materials and that you comply strictly with these terms. This copyright permission does not constitute an endorsement of the products or services. The Open Mobile Alliance assumes no responsibility for errors or omissions in this document.

Each Open Mobile Alliance member has agreed to use reasonable endeavors to inform the Open Mobile Alliance in a timely manner of Essential IPR as it becomes aware that the Essential IPR is related to the prepared or published specification. However, the members do not have an obligation to conduct IPR searches. The declared Essential IPR is publicly available to members and non-members of the Open Mobile Alliance and may be found on the “OMA IPR Declarations” list at <http://www.openmobilealliance.org/ipr.html>. The Open Mobile Alliance has not conducted an independent IPR review of this document and the information contained herein, and makes no representations or warranties regarding third party IPR, including without limitation patents, copyrights or trade secret rights. This document may contain inventions for which you must obtain licenses from third parties before making, using or selling the inventions. Defined terms above are set forth in the schedule to the Open Mobile Alliance Application Form.

NO REPRESENTATIONS OR WARRANTIES (WHETHER EXPRESS OR IMPLIED) ARE MADE BY THE OPEN MOBILE ALLIANCE OR ANY OPEN MOBILE ALLIANCE MEMBER OR ITS AFFILIATES REGARDING ANY OF THE IPR'S REPRESENTED ON THE “OMA IPR DECLARATIONS” LIST, INCLUDING, BUT NOT LIMITED TO THE ACCURACY, COMPLETENESS, VALIDITY OR RELEVANCE OF THE INFORMATION OR WHETHER OR NOT SUCH RIGHTS ARE ESSENTIAL OR NON-ESSENTIAL.

THE OPEN MOBILE ALLIANCE IS NOT LIABLE FOR AND HEREBY DISCLAIMS ANY DIRECT, INDIRECT, PUNITIVE, SPECIAL, INCIDENTAL, CONSEQUENTIAL, OR EXEMPLARY DAMAGES ARISING OUT OF OR IN CONNECTION WITH THE USE OF DOCUMENTS AND THE INFORMATION CONTAINED IN THE DOCUMENTS.

© 2003 Open Mobile Alliance Ltd. All Rights Reserved.

Used with the permission of the Open Mobile Alliance Ltd. under the terms set forth above.

Contents

1. SCOPE	7
2. REFERENCES	8
2.1 NORMATIVE REFERENCES	8
2.2 INFORMATIVE REFERENCES	8
3. TERMINOLOGY AND CONVENTIONS	9
3.1 CONVENTIONS	9
3.2 DEFINITIONS	9
3.3 ABBREVIATIONS	10
4. PROTOCOL OVERVIEW	11
4.1 PROTOCOL FEATURES	11
4.2 TRANSACTION CLASSES	11
4.2.1 Class 0: Unreliable Invoke Message with No Result Message	11
4.2.2 Class 1: Reliable Invoke Message with No Result Message.....	11
4.2.3 Class 2: Reliable Invoke Message with One Reliable Result Message.....	12
4.3 RELATION TO OTHER PROTOCOLS	12
4.4 SECURITY CONSIDERATIONS	13
4.5 MANAGEMENT ENTITY	13
4.6 INTEROPERABILITY CONSIDERATIONS	13
4.7 OTHER WTP USERS	14
5. ELEMENTS FOR LAYER-TO-LAYER COMMUNICATION	15
5.1 NOTATIONS USED	15
5.1.1 Definition of Service Primitives and Parameters	15
5.1.2 Primitive Types.....	15
5.1.3 Service Parameter Tables.....	15
5.2 REQUIREMENTS ON THE UNDERLYING LAYER	16
5.3 SERVICES PROVIDED TO UPPER LAYER	16
5.3.1 TR-Invoke.....	16
5.3.2 TR-InvokeData	18
5.3.3 TR-Result.....	18
5.3.4 TR-ResultData	18
5.3.5 TR-Abort.....	19
6. CLASSES OF OPERATION	20
6.1 CLASS 0 TRANSACTION	20
6.1.1 Motivation.....	20
6.1.2 Protocol Data Units.....	20
6.1.3 Procedure	20
6.2 CLASS 1 TRANSACTION	20
6.2.1 Motivation.....	20
6.2.2 Service Primitive Sequences	20
6.2.3 Protocol Data Units.....	21
6.2.4 Procedure	21
6.3 CLASS 2 TRANSACTION	21
6.3.1 Motivation.....	21
6.3.2 Service Primitive Sequences	21
6.3.3 Protocol Data Units.....	22
6.3.4 Procedure	22
7. PROTOCOL FEATURES	23
7.1 MESSAGE TRANSFER	23
7.1.1 Description	23
7.1.2 Service Primitives	23
7.1.3 Transport Protocol Data Units	23

7.1.4	Timer Intervals and Counters.....	23
7.1.5	Procedure	24
7.2	RE-TRANSMISSION UNTIL ACKNOWLEDGEMENT.....	24
7.2.1	Motivation.....	24
7.2.2	Transport Protocol Data Units	25
7.2.3	Timer Intervals and Counters.....	25
7.2.4	Procedure	25
7.3	USER ACKNOWLEDGEMENT	25
7.3.1	Motivation.....	25
7.3.2	Protocol Data Units.....	27
7.3.3	Procedure	27
7.4	INFORMATION IN LAST ACKNOWLEDGEMENT	28
7.4.1	Motivation.....	28
7.4.2	Service Primitives	28
7.4.3	Protocol Data Units.....	28
7.4.4	Procedure	28
7.5	CONCATENATION AND SEPARATION	28
7.5.1	Motivation.....	28
7.5.2	Procedure	28
7.6	ASYNCHRONOUS TRANSACTIONS	29
7.6.1	Motivation.....	29
7.7	TRANSACTION ABORT.....	29
7.7.1	Motivation.....	29
7.7.2	Service Primitives	29
7.7.3	Transport Protocol Data Units	29
7.7.4	Procedure	29
7.8	TRANSACTION IDENTIFIER	30
7.8.1	Motivation.....	30
7.8.2	Procedure at the Responder.....	30
7.8.3	Procedure at the Initiator.....	31
7.9	TRANSACTION IDENTIFIER VERIFICATION	32
7.9.1	Motivation.....	32
7.9.2	Protocol Data Units.....	32
7.9.3	Procedure	32
7.10	TRANSPORT INFORMATION ITEMS (TPIs)	33
7.10.1	Motivation.....	33
7.10.2	Procedure	33
7.11	TRANSMISSION OF PARAMETERS	33
7.11.1	Motivation.....	33
7.11.2	Procedure	34
7.12	ERROR HANDLING	34
7.12.1	Motivation.....	34
7.12.2	Protocol Data Units.....	34
7.12.3	Procedure	34
7.13	VERSION HANDLING	34
7.13.1	Motivation.....	34
7.13.2	Protocol Data Units.....	34
7.13.3	Procedure	34
7.14	SEGMENTATION AND RE-ASSEMBLY (OPTIONAL)	34
7.14.1	Motivation.....	34
7.14.2	Procedure for Segmentation.....	35
7.14.3	Procedure for Packet Groups	35
7.14.4	Procedure for Selective Re-transmission	35
7.15	EXTENDED SEGMENTATION AND RE-ASSEMBLY (OPTIONAL)	36
7.15.1	Motivation.....	36
7.15.2	Procedure for Segmentation.....	36
7.15.3	Procedure for Sliding Window	37

7.15.4	Procedure for Reliability	38
8.	STRUCTURE AND ENCODING OF PROTOCOL DATA UNITS.....	40
8.1	GENERAL	40
8.2	COMMON HEADER FIELDS	40
8.2.1	Continue Flag (CON).....	40
8.2.2	Group Trailer (GTR) and Transmission Trailer (TTR) Flag.....	41
8.2.3	Packet Sequence Number.....	41
8.2.4	PDU Type	41
8.2.5	Reserved (RES)	41
8.2.6	Re-transmission Indicator (RID).....	41
8.2.7	Transaction Identifier (TID).....	41
8.3	FIXED HEADER STRUCTURE	41
8.3.1	Invoke PDU	41
8.3.2	Result PDU	42
8.3.3	Acknowledgement PDU	42
8.3.4	Abort PDU	43
8.3.5	Segmented Invoke PDU (Optional)	44
8.3.6	Segmented Result PDU (Optional)	44
8.3.7	Negative Acknowledgement PDU (PDU).....	44
8.4	TRANSPORT INFORMATION ITEMS	44
8.4.1	General	44
8.4.2	Error TPI.....	45
8.4.3	Info TPI.....	46
8.4.4	Option TPI	46
8.4.5	Packet Sequence Number TPI (Optional)	48
8.4.6	SDU Boundary TPI.....	48
8.4.7	Frame Boundary TPI.....	48
8.5	STRUCTURE OF CONCATENATED PDUS.....	49
9.	STATE TABLES	50
9.1	GENERAL	50
9.2	EVENT PROCESSING.....	50
9.3	ACTIONS	50
9.3.1	Timers	50
9.3.2	Counters	51
9.3.3	Messages.....	51
9.4	TIMERS, COUNTERS AND VARIABLE	51
9.4.1	Timers	51
9.4.2	Counters.....	52
9.4.3	Variables	52
9.5	WTP INITIATOR.....	53
9.6	WTP RESPONDER	55
10.	EXAMPLES OF PROTOCOL OPERATION.....	57
10.1	INTRODUCTION.....	57
10.2	CLASS 0 TRANSACTION.....	57
10.2.1	Basic Transaction.....	57
10.3	CLASS 1 TRANSACTION.....	57
10.3.1	Basic Transaction.....	57
10.4	CLASS 2 TRANSACTION.....	58
10.4.1	Basic Transaction.....	58
10.4.2	Transaction with "Hold on" Acknowledgement	58
10.5	TRANSACTION IDENTIFIER VERIFICATION	59
10.5.1	Verification Succeeds	59
10.5.2	Verification Fails	59
10.5.3	Transaction with Out-of-Order Invoke	60
10.6	SEGMENTATION AND RE-ASSEMBLY	60
10.6.1	Selective Re-transmission.....	61

10.6.2	Re-transmission of the GTR/TTR Packet	61
10.6.3	SAR and TID Verification	62
10.6.4	Flow Control Using Option TPI (Maximum Group) Conjointly with SAR.....	62
10.6.5	Basic Extended SAR.....	64
10.6.6	Example of Re-transmission Hold-off	65
10.6.7	Another Example of Re-transmission Hold-off	66
10.6.8	SAR, NON-SAR Interactions	67
APPENDIX A.	DEFAULT TIMER AND COUNTER VALUES (NORMATIVE).....	69
APPENDIX B.	IMPLEMENTATION NOTES (INFORMATIVE)	71
APPENDIX C.	STATIC CONFORMANCE REQUIREMENTS (NORMATIVE)	73
APPENDIX D.	CHANGE HISTORY (INFORMATIVE).....	75
D.1	APPROVED VERSION HISTORY	75

1. Scope

A transaction protocol is defined to provide the services necessary for interactive "browsing" (request/response) applications. During a browsing session, the client requests information from a server, which MAY be fixed or mobile, and the server responds with the information. The request/response duo is referred to as a "transaction" in this document. The objective of the protocol is to reliably deliver the transaction while balancing the amount of reliability required for the application with the cost of delivering the reliability.

WTP runs on top a datagram service and optionally a security service. WTP has been defined as a light weight transaction oriented protocol that is suitable for implementation in "thin" clients (mobile stations) and operates efficiently over wireless datagram networks. The benefits of using WTP include:

- Improved reliability over datagram services. WTP relieves the upper layer from re-transmissions and acknowledgements which are necessary if datagram services are used.
- Improved efficiency over connection oriented services. WTP has no explicit connection set up or teardown phases.
- WTP is message oriented and designed for services oriented towards transactions, such as "browsing".

2. References

2.1 Normative References

- [IOPProc] “OMA Interoperability Policy and Process”. Open Mobile Alliance™. OMA-IOP-Process-v1_0.
[URL:http://www.openmobilealliance.org/](http://www.openmobilealliance.org/)
- [RFC2119] “Key words for use in RFCs to Indicate Requirement Levels”. S. Bradner. March 1997.
[URL:http://www.ietf.org/rfc/rfc2119.txt](http://www.ietf.org/rfc/rfc2119.txt)
- [WDP] “Wireless Datagram Protocol”, Open Mobile Alliance™, WAP-259-WDP,
<http://www.openmobilealliance.org/>.

2.2 Informative References

- [FLEX] FLEX™ Protocol Specification and FLEX™ Encoding and Decoding Requirements, Version G1.9, Document Number 68P81139B01, March 16, 1998, Motorola.
- [FLEXSuite] FLEX™ Suite of Application Protocols, Version 1.0, Document Number 6881139B10, October 29, 1997, Motorola.
- [GSM0260] ETSI European Digital Cellular Telecommunication Systems (phase 2+): General Packet Radio Service (GPRS) – stage 1 (GSM 02.60)
- [GSM0290] ETSI European Digital Cellular Telecommunication Systems (phase 2): Unstructured Supplementary Service Data(USSD) - stage 1 (GSM 02.90)
- [GSM0340] ETSI European Digital Cellular Telecommunication Systems (phase 2+): Technical realisation of the Short Message Service (SMS) Point-to-Point (P) (GSM 03.40)
- [GSM0360] ETSI European Digital Cellular Telecommunication Systems (phase 2+): General Packet Radio Service (GPRS) – stage 2 (GSM 03.60)
- [GSM0390] ETSI European Digital Cellular Telecommunication Systems (phase 2): Unstructured Supplementary Service Data(USSD) - stage 2 (GSM 03.90)
- [GSM0490] ETSI European Digital Cellular Telecommunication Systems (phase 2): Unstructured Supplementary Service Data(USSD) - stage 3 (GSM 04.90)
- [IS130] EIA/TIA IS-130
- [IS135] EIA/TIA IS-135
- [IS136] EIA/TIA IS-136
- [IS176] EIA/TIA IS-176 – CDPD 1.1 specifications
- [IS637] TIA/EIA/IS-637: Short Message Services for Wideband Spread Spectrum Cellular Systems
- [ISO7498] ISO 7498 OSI Reference Model
- [ISO8509] ISO TR 8509 Service conventions
- [RFC768] "User Datagram Protocol", J. Postel, August 1980, <http://www.ietf.org/rfc/rfc768.txt>
- [RFC791] "IP: Internet Protocol", J. Postel, <http://www.ietf.org/rfc/rfc791.txt>
- [ReFLEX] ReFLEX25 Protocol Specification Document, Version 2.6, Document Number 68P81139B02-A, March 16, 1998, Motorola.
- [TR45.3.6] General UDP Transport Teleservice (GUTS) ñ Stage III, TR45.3.6/97.12.15
- [WAPARCH] “Wireless Application Protocol Architecture Specification”. WAP Forum™. WAP-100-WAPArch. URL: <http://www.wapforum.org/>
- [WSP] “Wireless Session Protocol”, Open Mobile Alliance™, WAP-230-WSP,
<http://www.openmobilealliance.org/>.

3. Terminology and Conventions

3.1 Conventions

The key words “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, “SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in [RFC2119].

All sections and appendixes, except “Scope”, are normative, unless they are explicitly indicated to be informative.

3.2 Definitions

Device Address

The unique network address assigned to a device and following the format defined by an international standard such as E.164 for MSISDN addresses, X.121 for X.25 addresses or RFC 791 for IPv4 addresses. An address uniquely identifies the sending and/or receiving device.

Initiator

The WTP provider initiating a transaction is referred to as the Initiator.

Mobile Device

Refers to a device, such as a phone, pager, or PDA, connected to the wireless network via a wireless link. While the term ‘mobile’ implies the device is frequently moving, it MAY also include fixed or stationary wireless devices (i.e. wireless modems on electric meters) connected to a wireless network.

Network Type

Network type refers to any network, which is classified by a common set of characteristics (i.e. air interface) and standards. Examples of network types include GSM, CDMA, IS-136, iDEN™, FLEX, ReFLEX, and Mobitex. Each network type may contain multiple underlying bearer services.

Protocol Control Information (PCI)

Information exchanged between WTP entities to coordinate their joint operation.

Protocol Data Unit (PDU)

A unit of data specified in the WTP protocol and consisting of WTP protocol control information and possibly user data.

Responder

The WTP provider responding to a transaction is referred to as the Responder.

Service Data Unit (SDU)

Unit of information from an upper level protocol that defines a service request to a lower layer protocol.

Service Primitive

An abstract, implementation independent interaction between a WTP user and the WTP provider.

Transaction

The transaction is the unit of interaction between the Initiator and the Responder. A transaction begins with an invoke message generated by the Initiator. The Responder becomes involved with a transaction by receiving the invoke. In WTP several transaction classes have been defined. The invoke message identifies the type of transaction requested which defines the action required to complete the transaction.

User Data

The data transferred between two WTP entities on behalf of the upper layer entities (e.g. session layer) for whom the WTP entities are providing services.

WTP Provider

An abstract machine which models the behaviour of the totality of the entities providing the WTP service, as viewed by the user.

WTP User

An abstract representation of the totality of those entities in a single system that make use of the WTP service. Examples of WTP users include the WAP session protocol WSP or an application that runs directly onto WTP.

3.3 Abbreviations

API	Application Programming Interface
CDMA	Code Division Multiple Access
CDPD	Cellular Digital Packet Data
ESAR	Extended Segmentation And Reassembly
ETSI	European Telecommunication Standardisation Institute
GPRS	General Packet Radio Service
GSM	Global System for Mobile Communication
GTR	Group Trailer, indicates the end of packet group
GUTS	General UDP Transport Service
IDEN	Integrated Digital Enhanced Network
IP	Internet Protocol
LSB	Least significant bits
MDC	More Data Flag Cleared
MPL	Maximum Packet Lifetime
MS	Mobile Station
MSB	Most significant bits
MSISDN	Mobile Subscriber ISDN (Telephone number or address of device)
PCI	Protocol Control Information
PDU	Protocol Data Unit
PSN	Packet Sequence Number
RTT	Round-Trip Time
SAR	Segmentation and Re-assembly
SAP	Service Access Point
SDU	Service Data Unit
SMS	Short Message Service
SPT	Server Processing Time
TIA/EIA	Telecommunications Industry Association/Electronic Industry Association
TPI	Transport Information Item
TTR	Transmission Trailer
UDP	User Datagram Protocol
USSD	Unstructured Supplementary Service Data
WAP	Wireless Application Protocol
WDP	Wireless Datagram Protocol
WSP	Wireless Session Protocol
WTP	Wireless Transaction Protocol

4. Protocol Overview

4.1 Protocol Features

The following list summarises the features of WTP.

- Three classes of transaction service:
 - Class 0: Unreliable invoke message with no result message
 - Class 1: Reliable invoke message with no result message
 - Class 2: Reliable invoke message with exactly one reliable result message
- Reliability is achieved through the use of unique transaction identifiers, acknowledgements, duplicate removal and re-transmissions.
- No explicit connection set up or tear down phases. Explicit connection open and/or close imposes excessive overhead on the communication link.
- Optionally user-to-user reliability: the WTP user confirms every received message.
- Optionally, the last acknowledgement of the transaction MAY contain out of band information related to the transaction. For example, performance measurements.
- Concatenation MAY be used, where applicable, to convey multiple Protocol Data Units in one Service Data Unit of the datagram transport.
- Message orientation. The basic unit of interchange is an entire message and not a stream of bytes.
- The protocol provides mechanisms to minimise the number of transactions being replayed as the result of duplicate packets.
- Abort of outstanding transaction, including flushing of unsent data both in client and server. The abort can be triggered by the user canceling a requested service.
- For reliable invoke messages, both success and failure is reported. If an invoke can not be handled by the Responder, an abort message will be returned to the Initiator instead of the result.
- The protocol allows for asynchronous transactions. The Responder sends back the result as the data becomes available.

4.2 Transaction Classes

The following subsections describe the transaction classes of WTP. The WTP provider initiating a transaction is referred to as the Initiator. The WTP provider responding to a transaction is referred to as the Responder. The transaction class is set by the Initiator and indicated in the invoke message sent to the Responder. Transaction classes can not be negotiated.

4.2.1 Class 0: Unreliable Invoke Message with No Result Message

Class 0 transactions provide an unreliable datagram service. It can be used by applications that require an "unreliable push" service. This class is intended to augment the transaction service with the capability for an application using WTP to occasionally send a datagram within the same context of an existing session using WTP. It is not intended as a primary means of sending datagrams. Applications requiring a datagram service as their primary means of data delivery SHOULD use WDP [WDP].

The basic behaviour for class 0 transactions is as follows: One invoke message is sent from the Initiator to the Responder. The Responder does not acknowledge the invoke message and the Initiator does not perform re-transmissions. At the Initiator, the transaction ends when the invoke message has been sent. At the Responder, the transaction ends when the invoke has been received. The transaction is stateless and can not be aborted.

4.2.2 Class 1: Reliable Invoke Message with No Result Message

Class 1 transactions provide a reliable datagram service. It can be used by applications that require a "reliable push" service.

The basic behaviour for class 1 transactions is as follows: One invoke message is sent from the Initiator to the Responder. The invoke message is acknowledged by the Responder. The Responder maintains state information for some time after the

acknowledgement has been sent to handle possible re-transmissions of the acknowledgement if it gets lost and/or the Initiator re-transmits the invoke message. At the Initiator, the transaction ends when the acknowledgement has been received. The transaction can be aborted at any time.

If the User acknowledgement function is enabled, the WTP user at the Responder confirms the invoke message before the acknowledgement is sent to the Initiator.

4.2.3 Class 2: Reliable Invoke Message with One Reliable Result Message

Class 2 transactions provide the basic invoke/response transaction service. One WSP session MAY consist of several transactions of this type.

The basic behaviour for class 2 transactions is as follows: One invoke message is sent from the Initiator to the Responder. The Responder replies with exactly one result message that implicitly acknowledges the invoke message. If the Responder takes longer to service the invoke than the Responder's acknowledgement timer interval, the Responder MAY reply with a "hold on" acknowledgement before sending the result message. This prevents the Initiator from unnecessarily re-transmitting the invoke message. The Responder sends the result message back to the Initiator. The result message is acknowledged by the Initiator. The Initiator maintains state information for some time after the acknowledgement has been sent. This is done in order to handle possible re-transmissions of the acknowledgement if it gets lost and/or the Responder re-transmits the result message. At the Responder the transaction ends when the acknowledgement has been received. The transaction can at any time be aborted.

If the User acknowledgement function is enabled, the WTP user at the Responder confirms the invoke message before the result is generated. The WTP user at the Initiator confirms the result message before the acknowledgement is sent to the Responder.

4.3 Relation to Other Protocols

This chapter describes how WTP relates to other WAP protocols. For a complete description of the WAP Architecture refer to [WAP]. The following table illustrates the where the services provided to the WTP user are located.

	WTP User (e.g. WSP)
WTP	<input type="checkbox"/> Transaction handling <input type="checkbox"/> Re-transmissions, duplicate removal, acknowledgements <input type="checkbox"/> Concatenation and separation
[WTLS]	<input type="checkbox"/> Optionally compression <input type="checkbox"/> Optionally encryption <input type="checkbox"/> Optionally authentication
Datagram Transport (e.g. WDP)	<input type="checkbox"/> Port number addressing <input type="checkbox"/> Segmentation and re-assembly (if provided) <input type="checkbox"/> Error detection (if provided)
Bearer Network (e.g. IP, GSM SMS/USSD, IS-136 GUTS)	<input type="checkbox"/> Routing <input type="checkbox"/> Device addressing (IP address, MSISDN) <input type="checkbox"/> Segmentation and re-assembly (if provided) <input type="checkbox"/> Error detection (if provided)

WTP is specified to run over a datagram transport service. The WTP protocol data unit is located in the data portion of the datagram. Since datagrams are unreliable, WTP is required to perform re-transmissions and send acknowledgement in order to provide a reliable service to the WTP user. WTP is also responsible for concatenation (if possible) of multiple protocol data units into one transport service data unit.

The datagram transport for WAP is defined in [WDP]. The datagram transport is required to route an incoming datagram to the correct WDP user. Normally the WDP user is identified by a unique port number. The responsibility of WDP is to provide a datagram service to the WDP user, regardless of the capability of the bearer network type. Fortunately, datagram

service is a common transport mechanism, and most bearer networks already provide such a service. For example, for IP-based utilise UDP for this service.

The bearer network is responsible for routing datagrams to the destination device. Addressing is different depending on the type of bearer network (IP addresses or phone numbers). In addition, some networks are using dynamic allocation of addresses, and a server has to be involved to find the current address for a specific device. Network addresses within the WAP stack MAY include the bearer type and the address (e.g. [IP; 123.456.789.123]). The multiplexing of data to and from multiple bearer networks with different address spaces to the same WAP stack has not been specified.

4.4 Security Considerations

WTP has no security mechanisms.

4.5 Management Entity

The WTP Management Entity is used as an interface between the WTP layer and the environment of the device. The WTP Management Entity provides information to the WTP layer about changes in the device environment, which MAY impact the correct operation of WTP.

The WTP protocol is designed around an assumption that the environment in which it is operating is capable of transmitting and receiving data. For example, this assumption includes the following basic capabilities that MUST be provided by the mobile device:

- the mobile is within a coverage area applicable to the bearer service being invoked;
- the mobile having sufficient power and the power being on;
- sufficient resources (processing and memory) within the mobile are available to WTP;
- the WTP protocol is correctly configured, and ;
- the user is willing to receive/transmit data.

The WTP Management Entity monitors the state of the above services/capabilities of the mobile's environment and would notify the WTP layer if one or more of the assumed services were not available. For example if the mobile roamed out of coverage for a bearer service, the Bearer Management Entity SHOULD report to the WTP Management Entity that transmission/reception over that bearer is no longer possible. In turn, the WTP Management Entity would indicate to the WTP layer to close all active connections over that bearer. Other examples such as low battery power would be handled in a similar way by the WTP Management Entity.

In addition to monitoring the state of the mobile environment the WTP Management Entity MAY be used as the interface to the user for setting various configuration parameters used by WTP, such as device address. It could also be used to implement functions available to the user such as a 'drop all data connections' feature. In general the WTP Management Entity will deal with all issues related to initialisation, configuration, dynamic re-configuration, and resources as they pertain to the WTP layer.

Since the WTP Management Entity MUST interact with various components of a mobile device which are manufacturer specific, the design and implementation of the WTP Management Entity is considered outside the scope of the WTP Specification and is an implementation issue.

4.6 Interoperability Considerations

The static conformance requirements define a minimum set of WTP features that need to be implemented to ensure that the implementation will be able to interoperate. The WTP user dictates which WTP features it needs. These features can be specified by referring to the WTP static conformance requirement tables in Appendix C using the notation from [IOPProc].

If the WTP provider is requested to execute a procedure it does not support, the transaction MUST be aborted with the an appropriate error code. For example, a Responder not supporting class 2 receiving a class 2 transaction aborts the transaction with the NOTIMPLEMENTEDCL2 abort code.

Segmentation and re-assembly (SAR) and selective re-transmission MAY be implemented in order to enhance the WTP service. If SAR is not implemented in WTP, another layer in the stack should provide this functionality. For example, in IS-136 the SSAR layer handles SAR, in an IP network IP [RFC791] handles SAR and for GSM SMS/USSD SAR is achieved by using SMS concatenation [GSM0340]. The motivation for implementing WTP SAR is the selective re-transmission procedure, which MAY, if large messages are sent, improve the over-the-air efficiency of the protocol.

Extended Segmentation and Re-assembly complicates the picture further due to the fact that it MAY use either a sliding window based transmission or the traditional stop and wait mechanism.

Whether WTP SAR is supported or not is indicated by the Initiator when the transaction is invoked. The following table shows how WTP Initiators and Responders SHOULD guarantee interoperability between WTP providers that have and those that have not implemented WTP SAR.

Table 1 Interoperability between WTP Providers with no SAR, SAR, ESAR

Responder	Initiator		
	<i>No SAR</i>	<i>SAR</i>	<i>Extended SAR</i>
<i>No SAR</i>	Full interoperability	Responder aborts transaction with the abort code NOTIMPLEMENTEDSAR. Initiator MUST re-send without using SAR	Initiator MUST include NumGroups TPI in Invoke. Responder aborts transaction with the abort code NOTIMPLEMENTEDSAR. Initiator MUST re-send without using SAR
<i>SAR</i>	Responder MUST NOT respond with a segmented message	Full interoperability	Initiator MUST include NumGroups Option TPI in Invoke. Initiator will learn that Responder does not support ESAR by absence of this TPI in ACK, NACK or Result PDU.
<i>Extended SAR</i>	Responder MUST NOT respond with a segmented message	No NumGroups Option TPI is included in the Invoke message	Initiator MUST include NumGroups Option TPI in Invoke. Responder can respond using NumGroups Option TPI. Full interoperability.

Note 1) If a Responder not supporting WTP SAR receives a non-segmented message from an Initiator that supports WTP SAR, there is no need to abort the transaction. The Initiator will never be aware of the fact that the Responder does not support WTP SAR.

4.7 Other WTP Users

The intended use of this protocol is to provide WSP [WSP] with a reliable transaction service over an unreliable datagram service. However, the protocol can be used by other applications with similar communication needs.

5. Elements for Layer-to-Layer Communication

5.1 Notations Used

5.1.1 Definition of Service Primitives and Parameters

Communications between layers and between entities within the layer are accomplished by means of service primitives. Service primitives represent, in an abstract way, the logical exchange of information and control between the transaction layer and adjacent layers. They do not specify or constrain implementations.

Service primitives consist of commands and their respective responses associated with the services requested of another layer. The general syntax of a primitive is:

X - Generic name . Type (Parameters)

where X designates the layer providing the service. For this specification X is:

"TR" for the Transaction Layer.

An example of a service primitive for the WTP layer would be TR-Invoke.Request.

Service primitives are not the same as an application programming interface (API) and are not meant to imply any specific method of implementing an API. Service primitives are an abstract means of illustrating the services provided by the protocol layer to the layer above. The mapping of these concepts to a real API and the semantics associated with a real API are an implementation issue and are beyond the scope of this specification.

5.1.2 Primitive Types

The primitives types defined in this specification are

Type	Abbreviation	Description
Request	req	Used when a higher layer is requesting a service from the next lower layer
Indication	ind	A layer providing a service uses this primitive type to notify the next higher layer of activities related to the peer (such as the invocation of the request primitive) or to the provider of the service (such as a protocol generated event)
Response	res	A layer uses the response primitive type to acknowledge receipt of the indication primitive type from the next lower layer
Confirm	cnf	The layer providing the requested service uses the confirm primitive type to report that the activity has been completed successfully

5.1.3 Service Parameter Tables

The service primitives are defined using tables indicating which parameters are possible and how they are used with the different primitive types. For example, a simple confirmed primitive might be defined using the following:

Parameter	Primitive	TR-primitive			
		<i>req</i>	<i>ind</i>	<i>res</i>	<i>cnf</i>
Parameter 1		M	M(=)	-	-
Parameter 2		-	-	O	C(=)

In the example table above, *Parameter 1* is always present in *TR-primitive.request* and corresponding *TR-primitive.indication*. *Parameter 2* MAY be specified in *TR-primitive.response* and in that case it MUST be present and have the equivalent value also in the corresponding *TR-primitive.confirm*; otherwise, it MUST NOT be present.

If some primitive type is not possible, the column for it will be omitted. The entries used in the primitive type columns are defined in the following table:

Table 2. Parameter Usage Legend

M	Presence of the parameter is mandatory - it MUST be present
C	Presence of the parameter is conditional depending on values of other parameters
O	Presence of the parameter is a user option - it MAY be omitted
P	Presence of the parameter is a service provider option - an implementation MAY not provide it The parameter is absent
*	Presence of the parameter is determined by the lower layer protocol
(=)	The value of the parameter is identical to the value of the corresponding parameter of the preceding service primitive

5.2 Requirements on the Underlying Layer

The WTP protocol is specified to run on top of a datagram service. The datagram service **MUST** handle the following functions:

- Port numbers to route the incoming datagram to the WTP layer;
- Length information for the SDU passed up to the WTP layer.

The datagram service **MAY** handle the following functions

- Error detection. For example, by using a checksum.

In addition, Segmentation And Re-assembly (SAR) is expected to be provided by the underlying layers. However, it is usually done at a layer below the datagram layer. For example, in an IP network, the IP protocol handles SAR.

5.3 Services Provided to Upper Layer

5.3.1 TR-Invoke

This primitive is used to initiate a new transaction.

Parameter	Primitive	TR-Invoke			
		<i>req</i>	<i>ind</i>	<i>res</i>	<i>cnf</i>
Source Address		M	M (=)		
Source Port		M	M (=)		
Destination Address		M	M (=)		
Destination Port		M	M (=)		
Ack-Type		M	M (=)		
User Data		O	C (=)		
Class Type		M	M (=)		
Exit Info				O	C (=)
More Data		M	M (=)		
Frame Boundary		M	M (=)		
Handle		M	M	M	M

5.3.1.1 Source Address

The source address is the unique address of the device making a request to the WTP layer. The source address MAY be an MSISDN number, IP address, X.25 address or other identifier.

5.3.1.2 Source Port

The source port number associated with the source address.

5.3.1.3 Destination Address

The destination address of the user data submitted to the WTP layer. The destination address MAY be an MSISDN number, IP address, X.25 address or other identifier.

5.3.1.4 Destination Port

The destination port number associated with the destination address for the requested or existing transaction.

5.3.1.5 Ack-Type

This parameter is used to turn the User acknowledgement function on or off.

5.3.1.6 User Data

The user data carried by the WTP protocol. The unit of data submitted to or received from the WTP layer is also referred to as the Service Data Unit. This is the complete unit (message) of data that the higher layer has submitted to the WTP layer for transmission. The WTP layer will transmit the Service Data Unit and deliver it to its destination without any manipulation of its content.

5.3.1.7 Class Type

Indicates the WTP transaction class.

5.3.1.8 Exit Info

Additional user data to be sent to the originator on transaction completion. This parameter can be present only if *More Data* is cleared and *Class Type* is 1.

5.3.1.9 More Data

More Data is a Boolean flag that specifies whether additional invocations of the primitive will be following for the same transaction. This flag is valid only if the optional extended segmentation and re-assembly function is used.

5.3.1.10 Frame Boundary

Frame Boundary is a Boolean flag that specifies whether this User Data is the beginning of a new user defined frame. WTP will insert the Frame Boundary TPI into the very first data packet transmitted. This flag is valid only if the optional extended segmentation and re-assembly function is used.

5.3.1.11 Handle

The transaction handle is an index returned to the higher layer so the higher layer can identify the transaction and associate the data received with an active transaction. The TR-Handle uniquely identifies a transaction. TR-Handle is an alias for the source address, source port, destination address, and destination port of the transaction.

The TR-Handle has local significance only.

5.3.2 TR-InvokeData

This primitive is used to send additional data belonging to the same transaction in case of the optional extended segmentation and re-assembly function. This service primitive can be given only after the initial TR-Invoke service primitive started the transaction. The WTP user MUST issue Invoke.res and InvokeData.res primitives in the same order as Invoke.ind and InvokeData.ind primitives were received. Similarly, WTP issues Invoke.cnf and InvokeData.cnf primitives in the same order as Invoke.req and InvokeData.req primitives were submitted by the user. Note that the SDU associated with a InvokeData.req may be transmitted in multiple groups; the InvokeData.cnf MUST be issued only after the acknowledgement for the last group is received.

Parameter	Primitive	TR-InvokeData			
		<i>req</i>	<i>ind</i>	<i>res</i>	<i>cnf</i>
User Data		O	C (=)		
Exit Info				O	C (=)
More Data		M	M(=)		
Frame Boundary		M	M(=)		
Handle		M	M	M	M

5.3.2.1 Handle

The transaction handle is the index returned to the higher layer for the previous TR-Invoke, i.e. the TR-Invoke for which this TR-InvokeData is providing further data, so the higher layer can identify the transaction and associate the data received with the active transaction already started by the previous TR-Invoke service primitive. In the TR-Invoke.req the handle is passed up to the Initiator WTP user from the WTP provider. In the TR-InvokeData.req the handle is passed down to the WTP provider by the WTP user. The Handle uniquely identifies a transaction. Handle is an alias for the source address, source port, destination address, destination port, and the transaction identifier of the transaction.

The Handle has local significance only.

5.3.3 TR-Result

This primitive is used to send back a result of a previously initiated transaction.

Parameter	Primitive	TR-Result			
		<i>req</i>	<i>ind</i>	<i>res</i>	<i>cnf</i>
User Data		O	C (=)		
Exit Info				O	C (=)
More Data		M	M(=)		
Frame Boundary		M	M(=)		
Handle		M	M	M	M

5.3.3.1 Exit Info

Additional user data to be sent to the originator on transaction completion. This parameter can be present only if *More Data* is cleared.

5.3.4 TR-ResultData

This primitive is used to send additional data belonging to the same transaction in case of the optional extended segmentation and re-assembly function. This service primitive can be given only after the initial TR-Result service primitive started to send back the result of a the previously initiated transaction. The WTP user MUST issue Result.res and ResultData.res primitives in the same order as Result.ind and ResultData.ind primitives were received. Similarly, WTP issues Result.cnf and ResultData.cnf primitives in the same order as Result.req and ResultData.req primitives were submitted by the user. Note that

the SDU associated with a ResultData.req may be transmitted in multiple groups; the ResultData.cnf MUST be issued only after the acknowledgement for the last group is received.

Parameter	Primitive	TR-ResultData			
		<i>req</i>	<i>ind</i>	<i>res</i>	<i>cnf</i>
User Data	O	C (=)			
Exit Info				O	C (=)
More Data	M	M(=)			
Frame Boundary	M	M(=)			
Handle	M	M	M	M	M

5.3.4.1 Handle

The transaction handle is the index returned to the higher layer for the previous TR-Result, i.e. the TR-Result for which this TR-ResultData is providing further data, so the higher layer can identify the transaction and associate the data received with the active transaction already started by the previous TR-Invoke service primitive. In the TR-Invoke.req the handle is passed up to the Initiator WTP user from the WTP provider. In the TR-ResultData.req the handle is passed down to the WTP provider by the WTP user. The Handle uniquely identifies a transaction. Handle is an alias for the source address, source port, destination address, destination port, and the transaction identifier of the transaction.

The Handle has local significance only.

5.3.5 TR-Abort

This primitive is used to abort an existing transaction

Parameter	Primitive	TR-Abort	
		<i>req</i>	<i>ind</i>
Abort Code	O	O	C (=)
Handle	M	M	M

5.3.5.1 Abort Code

The abort code indicates the reason for the transaction being aborted. This can include abort codes generated by the WTP protocol and user defined local abort codes.

6. Classes of Operation

6.1 Class 0 Transaction

6.1.1 Motivation

Class 0 is an unreliable datagram service. It can be used by WSP [WSP], for example, to make an unreliable “push” within a session using the same socket association.

This class is intended to augment the transaction service with the capability for an application using WTP to occasionally send a datagram within the same context of an existing session using WTP. It is not intended as a primary means of sending datagrams. Applications requiring a datagram service SHOULD use WDP as defined in [WDP].

6.1.2 Protocol Data Units

The following PDU is used:

1. Invoke PDU

6.1.3 Procedure

A Class 0 transaction is initiated by the WTP user by issuing the TR-Invoke request primitive with the Transaction Class parameter set to Class 0. The WTP provider sends the invoke message and becomes the Initiator of the transaction. The remote WTP provider receives the invoke message and becomes the Responder of the transaction. The Initiator does not wait for or expect a response. If the invoke message is received by the Responder it is accepted immediately. There is no duplicate removal or verification procedure performed. However, the initiator MUST increment the TID counter between each transaction, but the responder MUST NOT update its cached TID.

This transaction class MUST be supported by the WTP provider. The WTP provider MUST be able to act as both Initiator and Responder.

An example of this class can be found in chapter 10.2.

6.2 Class 1 Transaction

6.2.1 Motivation

The Class 1 transaction is a reliable invoke message without any result message. This type of transaction can be used by WSP [WSP] to realise a reliable “push” service.

6.2.2 Service Primitive Sequences

The following table describes legal service primitive sequences. A primitive listed in the column header MAY only be followed by primitives listed in the row headers that are marked with an “X”.

Table 3 Primitive Sequence Table for Transaction Class 1

	TR-Invoke				TR-Abort	
	<i>req</i>	<i>ind</i>	<i>res</i>	<i>cnf</i>	<i>req</i>	<i>ind</i>
TR-Invoke.req						
TR-Invoke.ind						
TR-Invoke.res		X				
TR-Invoke.cnf	X					
TR-Abort.req	X	X	X			
TR-Abort.ind	X	X	X			

6.2.3 Protocol Data Units

The following PDUs are used:

1. Invoke PDU
2. Ack PDU
3. Abort PDU

6.2.4 Procedure

A Class 1 transaction is initiated by the WTP user by issuing the TR-Invoke request primitive with the Transaction Class parameter set to Class 1. The WTP provider sends the invoke message and becomes the Initiator of the transaction. The remote WTP provider receives the invoke message and becomes the Responder of the transaction. The Responder checks the Transaction Identifier and determines whether a verification has to be initiated. If not, it delivers the message to the user and returns the last acknowledgement to the Initiator. The Responder **MUST** keep state information in order to re-transmit the last acknowledgement if it gets lost.

This transaction class **MUST** be supported by the WTP provider. The WTP provider **MUST** be able to act as both Initiator and Responder.

An example of this class can be found in chapter 10.3.

6.3 Class 2 Transaction

6.3.1 Motivation

The Class 2 transaction is the basic request/response transaction service. This is the most commonly used transaction service. For example, it is used by WSP [WSP] for method invocations.

6.3.2 Service Primitive Sequences

The following table describes generally the legal service primitive sequences by generic names. A primitive listed in the column header **MAY** only be followed by primitives listed in the row header and marked with an "X". MDC indicates that the More Data flag is cleared, otherwise it is set.

Table 4 Primitive Sequence Table for Transaction Class 2

	TR-Invoke	TR-InvokeData	TR-Result	TR-ResultData	TR-Abort	TR-Invoke MDC	TR-InvokeData MDC	TR-Result MDC	TR-ResultData MDC
TR-Invoke									
TR-InvokeData	X	X							
TR-Result						X	X		
TR-ResultData			X	X					
TR-Abort	X	X	X	X		X	X	X	
TR-Invoke MDC									
TR-InvokeData MDC	X	X							
TR-Result MDC						X	X		
TR-ResultData MDC			X	X					

6.3.3 Protocol Data Units

The following PDUs are used:

1. Invoke PDU
2. Result PDU
3. Ack PDU
4. Abort PDU

6.3.4 Procedure

A Class 2 transaction is initiated by the WTP user by issuing the TR-Invoke request primitive with the Transaction Class parameter set to Class 2. The WTP provider sends the invoke message and becomes the Initiator of the transaction. The remote WTP provider receives the invoke message and becomes the Responder of the transaction. The Responder checks the Transaction Identifier and determines whether a verification has to be initiated. If not, it delivers the message to the WTP user and wait for the result. The Responder MAY send a hold on acknowledgement after a specified time.

The WTP user sends the result message by issuing the TR-Result request primitive. When the Initiator has received the result message it returns the last acknowledgement to the Responder. The Initiator MUST keep state information in order to re-transmit the last acknowledgement if it gets lost.

If the Responder does not support this transaction class it returns an Abort PDU with the abort reason NOTIMPLEMENTEDCL2 as a response to the invoke message.

An example of this class can be found in chapter 10.4.

7. Protocol Features

7.1 Message Transfer

7.1.1 Description

WTP consists of two types of messages: data messages and control messages. Data messages carry user data. Control messages are used for acknowledgements, error reporting, etc. and do not carry user data. This section gives the reader an overall picture of how transactions are realised by WTP. The procedures to guarantee reliable message transfer are outlined. Special functions like concatenation and separation, re-transmission until acknowledgement, transaction abort, user acknowledgement, and others are described in further detail in separate sections.

It is important to note that not all messages and functions are used by all transaction classes. The following table illustrates which messages are used for the different transaction classes.

Table 5 Summary of WTP message transfer

Message/function	Class 2	Class 1	Class 0
Invoke message	X	X	X (Note 2)
Verification	X	X	
Hold on acknowledgement	X (Note 1)		
Result message	X		
Last acknowledgement	X	X	

Note 1) Only sent in the case when the user takes longer time to service the invoke message than the Responder's acknowledgement timer interval.

Note 2) The class 0 transaction is unreliable. No response is expected from the Responder and no verification is performed.

7.1.2 Service Primitives

The following service primitives are used during nominal WTP transactions. Their use is transaction class dependent:

1. TR-Invoke
2. TR-Result

7.1.3 Transport Protocol Data Units

The following PDUs are used during nominal WTP transactions. It is important to note that not all PDUs are used in every transaction class.

1. Invoke PDU
2. Result PDU
3. Ack PDU

7.1.4 Timer Intervals and Counters

The following timer intervals and counters are used during a nominal WTP transaction. Their use is transaction class dependent.

1. Re-transmission interval
2. Re-transmission counter
3. Acknowledgement interval
4. Wait timeout interval

The values and relations between timer intervals and counters MAY depend on the transaction class being used. A detailed description of timers and counters is provided in a separate section.

7.1.5 Procedure

A transaction takes place between two WTP providers. A WTP user initiates a transaction by issuing the TR-Invoke request primitive. The TCL parameter of the primitive indicates the transaction class: 0, 1 or 2. In WTP, the Initiator is the WTP provider initiating the transaction and the Responder is the WTP provider responding to the initiated transaction.

7.1.5.1 Invoke Message

The invoke message is always the first message of a transaction and it is sent using the Invoke PDU. The Initiator administers the Transaction Identifier (TID) by incrementing the TID by one for every initiated transaction. The TID is conveyed in every PDU belonging to the transaction. When the Invoke PDU has been sent the Initiator starts the re-transmission timer and waits for a response. When the Responder receives the Invoke PDU with a valid TID, it delivers the message to the user by generating the TR-Invoke indication primitive.

7.1.5.2 Verification

When the Responder has received and accepted the invoke message it SHOULD cache the TID. This is done in order to filter out duplicate and old invoke messages that have lower or identical TID values (see section on Transaction Identifier). If the Responder determines the TID in the Invoke PDU is invalid, the Responder can verify whether the invoke message is a new or delayed message. This is accomplished by sending an Ack PDU which initiates a three way handshake towards the Initiator (see section on TID Verification). In this case, the Responder MUST NOT deliver the data to the user until the three-way handshake is successfully completed. If the three-way handshake attempt fails, the transaction is aborted by the Initiator.

7.1.5.3 Hold on Acknowledgement

When the invoke message has been delivered to the WTP user, the acknowledgement timer is started. If the WTP user requires more time to service the invoke message than the acknowledgement timer interval, the Responder MAY or SHOULD or MUST send a 'hold on' acknowledgement. This is done to prevent the Initiator from re-transmitting the Invoke PDU. When the Initiator receives the Ack PDU it stops re-transmitting the Invoke PDU and generates the TR-Invoke confirm primitive.

7.1.5.4 Result Message

Upon assembling the data, the WTP user sends a result message by initiating the TR-Result request primitive. The result message is transmitted using the Result PDU. When the Result PDU has been sent the Responder starts the re-transmission timer and waits for a response. After the Result PDU is received by the Initiator it generates the TR-Invoke confirm primitive if one has not already been issued and the forwards up the TR-Result indication primitive.

7.1.5.5 Last Acknowledgement

The last Ack PDU is sent when the last message of the transaction has been received. The sender of the acknowledgement MUST maintain state information required to handle a re-transmission of the previous message. This can be done by using a wait timer, or by keeping a transaction history that indicates the results of past transactions.

7.2 Re-transmission until Acknowledgement

7.2.1 Motivation

The re-transmission until acknowledgement procedure is used to guarantee reliable transfer of data from one WTP provider to another in the event of packet loss. To minimise the number of packets sent over-the-air, WTP uses implicit acknowledgements wherever possible. An example of this is the use of the Result message to implicitly acknowledge the Invoke message.

7.2.2 Transport Protocol Data Units

The following PDUs are used:

1. Invoke PDU
2. Result PDU
3. Ack PDU

7.2.3 Timer Intervals and Counters

The following timer intervals and counters are used:

1. Re-transmission interval
2. Re-transmission counter

The values and relationships between timers and counters MAY depend on the transaction class being used. A detailed description of timers and counters is provided in a separate section.

7.2.4 Procedure

When a packet has been sent, the re-transmission timer is started and the re-transmission counter is set to zero. If a response has not been received when the re-transmission timer expires, the re-transmission counter is incremented by one, the packet re-transmitted, and the re-transmission timer re-started. The WTP provider continues to re-transmit until the number of re-transmissions has exceeded the maximum re-transmission value. If no acknowledgement has been received when the retransmission counter is fully incremented and the timer expires, the transaction is terminated and the local WTP user is informed.

In an extended SAR transaction, the re-transmission timer is used to re-transmit a packet group. If there are no more packets to send, the re-transmission timer is started and the re-transmission counter is set to zero. If an acknowledgement or a negative acknowledgement has not been received when the re-transmission timer expires, the re-transmission counter is incremented by one, the last unacknowledged GTR or TTR packet is re-transmitted and the re-transmission timer is re-started. The WTP provider continues to re-transmit until the number of re-transmissions exceeds the maximum re-transmission value. If no acknowledgement has been received when the re-transmission counter is fully incremented and the timer expires, the transaction is terminated and the local WTP user is informed.

The first time a PDU is transmitted the re-transmission indicator (RID) field in the header is clear. For all re-transmissions the RID field is set. Other than the RID field, the WTP provider MUST NOT change any fields in the PDU header.

The motivation for the re-transmission indicator is for the receiver to detect messages that have been duplicated by the network. A WTP provider that receives two identical messages with the RID set to zero, can safely ignore the second message because it must have been duplicated by the network. Any subsequent retransmissions that have the RID flag set to one can not be ignored by the receiver. Re-transmitted messages that get duplicated by the network must be treated as valid messages by the provider. The receiver in this situation can no longer distinguish between provider retransmissions and network duplicated packets. In this case, if the message is an Invoke PDU, there is a risk that the transaction will be re-played. To avoid such an error, the WTP provider should make a TID validation (see chapter 7.8).

7.3 User Acknowledgement

7.3.1 Motivation

The User Acknowledgement function allows for the WTP user to confirm every message received by the WTP provider. In case of an extended SAR transaction if this function is enabled, the WTP user will acknowledge the last packet group of the data flow in the appropriate direction.

When this function is enabled, the WTP provider does not respond to a received message until after the WTP user has confirmed the indication service primitive (by issuing the response primitive). If the WTP user does not confirm the indication primitive after a specified time, the transaction is aborted by the provider. Note that this is a much stronger form of a confirmed service than the traditional definition [ISO8509]. The traditional definition of a confirmed service is that there is

a confirmation from the service provider, however, there is not necessarily any relationship to a response from the peer service user. In WTP, when the User Acknowledgement function is used, the service provider requires a response from the service user for each indication. As a result, when the confirmation primitive is generated, there is a guarantee that there was a response from the peer service user.

This function is optional within WTP however WSP does utilise the User Acknowledgement feature and therefore any implementation of WTP that will have WSP as the higher layer, must implement it (see Appendix C). WSP requires a feature that at the end of a request-response transaction, the server gets a positive indication that the client has processed the response. This is illustrated below.

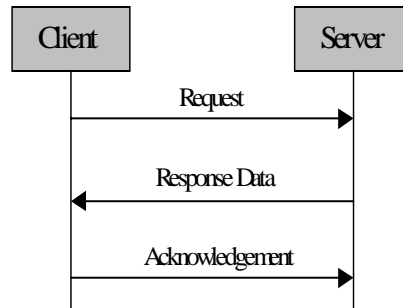


Figure 1 Generic WSP [WSP] transaction

In this model, the *Acknowledgement* is used to convey the fact that the response was received and processed by the client application. It is important to note that the *Client* and the *Server* in the figure refers to the client and server **Application**, and not only the protocol stack.

When the User Acknowledgement function is used the WSP -WTP primitive sequence for a Class 2 transaction becomes as illustrated below.

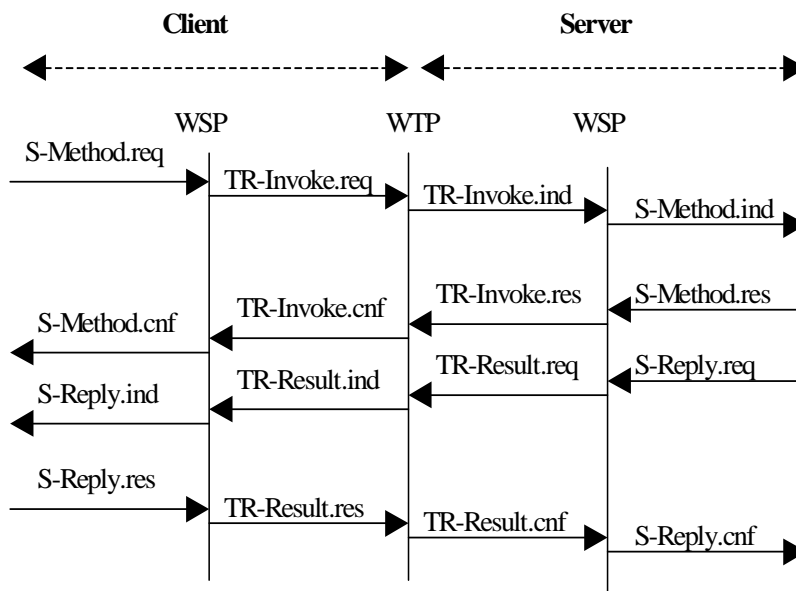


Figure 2 WSP-WTP primitive sequence for request-response

The primitive sequence started by the *S-Reply.res* and *TR-Result.res* primitives realises the *Complete* and *Confirm* primitives from Figure 1. If the application and/or the WSP for some reason does not issue these primitives, WTP aborts the transaction with the NORESPONSE reason. The abort is used by the WSP server as an indication that the result was not properly received or processed by the client.

The primitive sequence started by the *S-Method.res* and *TR-Invoke.res* primitives can be used by the client WSP to indicate to the application (and human user) that the invoke message has been received by the server WSP.

When this function is not used, WTP MAY acknowledge received messages independently of the WTP user. In Figure 2 this means that the response primitives MAY be ignored by the WTP provider. Put in other words: the WTP provider receives a message, returns an acknowledgement and indicates to the user that a message has been received. If there is an error, the transaction will be aborted by the WTP provider. If the WTP user is alive but can not process the message it MAY abort the transaction with an appropriate abort reason.

This function is optional. It applies to transaction class 1 and 2.

Note) Even though the WTP user has issued a response primitive there is no guarantee that it has interpreted the data and started processing. The WTP user MAY have only copied the data from one buffer to another, or issued the response primitive without any action taken at all. A WTP user can always abort a transaction if it discovers that the received data is corrupt or for some other reason not possible to process (see section on Transaction abort).

7.3.2 Protocol Data Units

The following PDUs are used:

1. Invoke PDU
2. Abort PDU

7.3.3 Procedure

The Initiator sets the U/P-flag in the Invoke PDU to indicate that User acknowledgement is required. A Responder not supporting this function aborts the transaction with the abort reason NOTIMPLEMENTEDUACK. The Initiator MAY then take the decision to re-initiate the transaction without the User acknowledgement function.

When the Responder receives the Invoke PDU with the U/P-flag set it generates the TR-Invoke indication and starts the acknowledgement timer. To give the WTP user time to read the parameters in the indication primitive and issue the TR-Invoke response primitive, the value of the timer MAY have a higher value than the provider's acknowledgement timer (see definitions of default timer values). The Responder MUST NOT return a response before the WTP user has issued the TR-Invoke response primitive. If the Initiator re-transmits Invoke PDUs due to lack of acknowledgement, the Responder MUST silently discard the PDU and restart the acknowledgement timer. When the WTP user issues the TR-Invoke response primitive, the Responder is enabled to send the Ack PDU. If the TR-Invoke response primitive has not been issued after a specified time, the provider aborts the transaction with the abort reason NORESPONSE. If the WTP user issues the TR-Result request primitive, the result is sent instead of the acknowledgement. The Initiator receiving the Ack PDU generates the confirm primitive which indicates that the remote WTP user has issued the corresponding response primitive.

For class 2 transactions, if the Initiator has indicated that the User Acknowledgement function shall be used, it is valid for the entire transaction. This means that when the Initiator has received the result and generated the TR-Result indication primitive it MUST wait for the TR-Result response primitive from the WTP user before the last acknowledgement can be sent. If the TR-Result response primitive has not been issued after a specified time, the provider aborts the transaction with the abort reason NORESPONSE. When the Responder receives the NORESPONSE abort it generates the TR-Abort indication primitive, indicating to the WTP user that the transaction failed.

7.4 Information in Last Acknowledgement

7.4.1 Motivation

The WTP user is allowed to attach information in the last, and only the last, acknowledgement of a transaction. This function is meant for transporting small amounts of information related to the transaction. The information can be, for example, performance measurements collected in order to evaluate the user's perceived quality of service.

For class 2 transactions, this function can be used by the Initiator to communicate some information back to the Responder. For a class 1 transaction, this function can be used by the Responder to communicate some information to the Initiator.

7.4.2 Service Primitives

The following service primitives and parameters are used:

1. TR-Result.res (Class 2)
2. TR-Invoke.res (Class 1)

7.4.3 Protocol Data Units

The following PDU is used:

1. Ack PDU

7.4.4 Procedure

For a class 2 transaction, information is attached to the last acknowledgement by issuing the TR-Result response primitive with the ExitInfo parameter.

For a class 1 transaction, information is attached to the last acknowledgement by issuing the TR-Invoke response primitive with the ExitInfo parameter.

The exit information is transferred as a Transport Information Item (TPI) in the variable part of the Ack PDU header.

For class 2 transactions, the ExitInfo parameter MUST NOT be included in the TR-Invoke response primitive and the Info TPI MUST NOT be included in the Ack PDU that acknowledges the Invoke PDU.

7.5 Concatenation and Separation

7.5.1 Motivation

Concatenation is the procedure to convey multiple WTP Protocol Data Units (PDUs) in one Datagram Service Data Unit (SDU) of the bearer network. When concatenation is done, a special mapping of the WTP PDUs to the SDUs is used. This is described in chapter 8.5.

Separation is the procedure to extract multiple PDUs from one SDU. When the PDUs have been separated they are dispatched to the transactions.

Concatenation and separation is used to provide over-the-air efficiency, since fewer transmissions over the air are required.

7.5.2 Procedure

Concatenation can only be done for messages with the same address information (source and destination port, source and destination device address).

Concatenation of PDU from different transactions can be done at any time. For example, the last acknowledgement of one transaction can be concatenated with the invoke message of the next transaction. Concatenation and separation is performed outside the WTP state machine.

The exact implementation of concatenation is not specified. Only the structure to be used when multiple packets are concatenated is specified. Exactly how the packets are buffered and concatenated is an implementation issue.

7.6 Asynchronous Transactions

7.6.1 Motivation

The implementation of the WTP provider SHOULD be able to initiate multiple transactions before it receives the response to the first transaction. Multiple transactions SHOULD be handled asynchronously. For example, the responses to transaction number 1, 2 and 3 MAY arrive to the Initiator as 3, 1 and 2. The Responder SHOULD send back the result as soon as it is ready, independently of other transactions.

The maximum number of outstanding transactions at any moment is limited by the maximum number of Transaction Identifiers. The Transaction Identifier is 16 bits, but the high order bit is used to indicate the direction of the message, so the maximum number of outstanding transactions is 2^{15} . The implementation environment will also set a limit to how many outstanding transactions it can handle simultaneously.

If the maximum number of outstanding transactions is exceeded the responder should ignore and discard the invoke message.

7.7 Transaction Abort

7.7.1 Motivation

An outstanding transaction can be aborted by the WTP user by issuing the TR-Abort request primitive. The user abort can be triggered by the application (e.g. input from human user) or it can be a negative result (e.g. the WTP user could not generate a result due to an error).

An outstanding transaction can also be aborted by the WTP provider due to a protocol error (e.g. reject the received data) or if a requested function is not implemented.

This function MUST be used with care. If the invoke message has already been sent, the response message MAY be on its way to the client and an abort will only increase network load.

7.7.2 Service Primitives

The following service primitive is used:

1. TR-Abort

7.7.3 Transport Protocol Data Units

The following PDU is used:

1. Abort PDU

7.7.4 Procedure

There are three special cases of the abort procedure:

- A) The sending WTP provider has not yet sent the message: the provider MUST discard the message from its memory.
- B) The sending WTP provider has sent the message to the peer, or is in the process of sending the message: the provider MUST send the Abort PDU to the remote peer to discard all data associated with the transaction.
- C) The receiving provider receives the Abort PDU: it generates the TR-Abort indication primitive and discards all transaction data.

When an Abort PDU is sent the reason for the abort is indicated in the abort reason field. There are two main types of aborts: User abort (USER) and Provider abort (PROVIDER). The user abort occurs when the WTP user has issued the TR-Abort request primitive. The provider abort occurs when there is an error in the WTP provider.

7.8 Transaction Identifier

7.8.1 Motivation

A transaction is uniquely identified by the socket pair (source address, source port, destination address and destination port) and the Transaction Identifier (TID). The Initiator increments the TID by one for every initiated transaction. This means that TIDs 1, 2 and 3 can go to server A, TIDs 4, 5 and 6 to server B and TIDs 7, 8 and 9 to server A.

The main use of the TID is to identify messages belonging to the same transaction. When a message is re-transmitted the TID is reused for the re-transmitted messages. A Responder MAY choose to remember the TID after an invoke message has been accepted and force TID verification in order to avoid replaying transactions. Also, the Initiator increments the TID by one for each transaction. This information can be used by the Responder to filter out new invoke messages from old and duplicated invoke messages: a new invoke always has a higher TID value.

Since transactions can be initiated simultaneously from both directions on the same socket association, the high order bit of the TID is used to indicate the direction of the transaction. The Initiator sets the high order bit to 0 in the Invoke PDU. Thereafter the high order bit is always inverted in the received TID before it is added to the response packet. By setting the high order bit of the TID field to 0 at the Initiator and 1 at the Responder, the Initiator can be guaranteed that the allocated TID will not collide with the remote entity.

The TID is 16-bits but the high order bit is used to indicate the direction. This means that the TID space is 2^{15} . The TID is an unsigned integer.

7.8.2 Procedure at the Responder

7.8.2.1 Variables

If the Responder caches old TID values for each different Initiator the old TID value is called LastTID. The TID in the received invoke message is called RcvTID.

7.8.2.2 Decisions when Receiving a New Invoke Message

When the Responder receives an invoke message it takes one of the following actions depending on whether the Responder is caching old TID values or not, the characteristics of the underlying transport and the outcome of the TID test (described in the following chapter):

Table 6 Decisions when receiving new invoke message

Event	Condition	Action
TID test Fail	Underlying transport service can guarantee there are no duplicates (Note 1)	Start transaction
	Underlying transport service can NOT guarantee there are no duplicates	Invoke TID verification
TID test Ok		LastTID = RcvTID Start transaction
TIDnew flag set		LastTID = 0 Invoke TID verification
No cache	Responder caches the TID for each Initiator for the Maximum Packet Lifetime (MPL) of the network and it has not been re-booted during this time period and lost the information. If the invoke was not a new one, the Responder would have had the latest TID in its cache.	Create new record for this Initiator LastTID = RcvTID Start transaction
	Responder does not cache TIDs (Note 2)	Invoke TID verification

Note 1) This is the case, for example, if a security layer is located under WTP and that can remove duplicates.

Note 2) This is not very efficient and SHOULD be avoided.

7.8.2.3 The TID Test

One method of validating the TID is to use a window mechanism. The Responder MAY cache the last valid TID (LastTID) from each different Initiator. When the Responder receives a new invoke message it compares the TID in the invoke message (RcvTID) with the cached one. Let W be the size of the window. If $W=2^{*}14$, it means that the boundary between two TID values occurs when they differ by $2^{*}14$, that is, half the TID space.

Table 7 TID test; RcvTID \geq LastTID

RcvTID \geq LastTID	
RcvTID - LastTID	TID test
0	Fail
$\leq W$	Ok
$> W$	Fail (see 7.8.2.4)

Table 8 TID test; RcvTID $<$ LastTID

RcvTID $<$ LastTID	
RcvTID - LastTID	TID test
$< W$	Fail (see 7.8.2.4)
$\geq W$	Ok

The above tables show different results from the TID test. If the test succeeds it is guaranteed that the received invoke message is new and not an old delayed one. This is under the assumption that all messages have a Maximum Packet Lifetime (MPL), and that after MPL seconds it is guaranteed that there are no duplicate messages present in the network (see Note). Furthermore, it is assumed that the TID is not incremented faster than $2^{*}14$ steps in $2^{*}MPL$.

Note) For some networks types, the average Maximum Packet Lifetime MAY have a very high variance. For example, in a store-and-forward network like GSM SMS, a short message MAY reside in the SMS-C for a very long time, before it gets delivered to the destination. This fact MAY in some cases violate the correctness of the TID validation.

7.8.2.4 Reception of Out-of-order Invoke Messages

Messages can arrive out-of-order. This means that even if the Initiator increments the TID by one for each transaction, a transaction with a lower TID value can arrive after a TID with a higher value. This MAY cause the TID test to fail and a TID verification to be started. This will not break the protocol, however, it will lead to degraded performance. One way to overcome this is to keep an array of TID values for past transactions. If the received TID is not in the array it can be accepted without any TID verification. This solution improves performance, but requires the Responder to maintain more information.

7.8.3 Procedure at the Initiator

7.8.3.1 Administration of TID

The Initiator is responsible for incrementing the TID by one for each transaction. This MUST NOT be done faster than $2^{*}14$ steps in $2^{*}MPL$.

7.8.3.2 Violating the Monotonic Property of the TID

There are cases when the Initiator MAY generate non-monotonic TID values, that is, the next TID MAY be smaller than the previous:

1. The Initiator has crashed and re-booted and randomly picked a smaller TID value than the previous.
2. The TID values have wrapped around the finite space. This can happen if, for example, the Initiator sends a transaction to Responder A, then sends $2^{*}14$ transactions to Responder B and finally returns to Responder A. The cached TID value at Responder A for this Initiator will now be smaller than current TID.

Neither of these two cases will break the protocol. However, TID verifications will be invoked and that will lead to lower efficiency.

In (1), if the Responder discards cached TID values after MPL seconds and the time to re-boot takes longer than that, the Responder will accept the new TID value without a TID verification (see 7.8.2.2). We have assumed that it will take longer time than $2 * MPL$ to increment the TID $2^{**}14$ steps. However, if the Responder caches the TID value longer than for MPL seconds it will initiate a TID verification in this case.

The wraparound in (2) will be detected only if the Initiator caches the last sent TID to each Responder.

In both (1) and (2) excessive use of the TID verification mechanism SHOULD be avoided by setting the TIDnew flag in the Invoke PDU, i.e. when the initiator has received multiple subsequent verification requests from the responder the TIDnew flag should be set in the next transaction. This will invalidate the Responder's cached TID for the Initiator (see 7.8.2.2). When the Initiator uses the TIDnew flag it MUST NOT initiate any subsequent transaction until the TID verification has been completed. The reason for this is that the TIDnew MAY be delayed in the network. If, during that time period, transactions with higher TID are initiated, duplicates from these will get erroneously accepted when Responder has updated its cache with the lower TID in the TIDnew packet.

7.9 Transaction Identifier Verification

7.9.1 Motivation

The transaction identifier verification procedure is a three-way handshake. A three-way handshake between an Initiator (I) and a Responder (R) has the following steps:

- (1) I → R This is the TID (Invoke PDU)
- (2) I ← R Do you have an outstanding transaction with this TID? (Ack PDU)
- (3) I → R Yes/No! (Ack PDU / Abort PDU)

The TID verification procedure is necessary to guarantee that the same invoke message is not accepted and delivered to the WTP user more than once, due to old duplicate packets.

The invoke message MUST NOT be delivered to the user before the TID verification procedure is completed successfully.

7.9.2 Protocol Data Units

The following PDUs are used:

1. Invoke PDU
2. Ack PDU
3. Abort PDU

7.9.3 Procedure

In the event that the Responder has received an Invoke PDU from an Initiator and has decided, using the rules for the Transaction Identifier procedure, to verify the TID, the following process is used.

The Responder sends an Ack PDU with Tve flag set indicating that it has received an invoke message with this TID.

When the Initiator receives the Ack PDU from the Responder it checks whether it has a corresponding outstanding transaction with this TID. In this case, the Initiator sends back an Ack PDU with TIDok flag set indicating that the TID is valid. This completes the three way handshake. If the Initiator does not have a corresponding outstanding transaction, it MUST abort the transaction by sending an Abort PDU with the Abort reason INVALIDTID.

Depending on the outcome of the TID verification WTP SHOULD take different actions. These are listed in the below table.

Table 9 Actions depending on result of TID verification

Result of TID verification	Condition	Action
Valid TID	TIDnew == True	Start transaction LastTID = RcvTID

Result of TID verification	Condition	Action
	TIDnew == False	Start transaction LastTID = LastTID
Invalid TID		Abort transaction

The TIDnew flag is set in the invoke message and is used by the Initiator to invalidate the Responder's cache.

An example of this procedure can be found in chapter 10.5.

7.10 Transport Information Items (TPIs)

7.10.1 Motivation

The variable portion of the header in a WTP PDU MAY consist of Transport Information Items (TPIs). If not, the variable part of the header MUST be empty. The use of TPIs allows for future extensions of the protocol.

7.10.2 Procedure

All TPIs follow the general structure: TPI identity, TPI length and TPI data; the length can be zero. The following table lists the currently defined TPI and in which section they are explained:

Table 10 WTP Transport Information Items (TPIs)

Transport Information Item	Described in section
Error	"Transport Information Items (TPIs)" section 7.10
Frame Boundary	"Procedure for Segmentation" section 7.15.2
Info	"Information in Last Acknowledgement" section 7.4
Option	"Transmission of Parameters" section 7.11
Packet Sequence Number	"Segmentation and Re-assembly" section 7.14
SDU Boundary	"Procedure for Segmentation" section 7.15.2

A WTP provider without error ignore a TPI it does not implement, assuming the general TPI structure is used by all TPIs.

The error TPI can be used to inform the sender that an unsupported or erroneous TPI was received. When a WTP provider receives a TPI that is not supported, the WTP provider returns the Error TPI with the ErrorCode indicating "Unknown TPI" along with the identity of the unsupported TPI. When a WTP provider receives a supported TPI, but fails to understand the content of the TPI, the WTP provider returns the Error TPI with the ErrorCode indicating "Known TPI, unknown content", and the identity of the TPI and the first octet of the content included as argument.

Note that if an unsupported or erroneous TPI is received in the last message of a transaction, the receiver can not notify the sender of the event.

7.11 Transmission of Parameters

7.11.1 Motivation

Protocol parameters can be transmitted between two WTP providers by using the Option TPI in the variable part of the PDU header.

No mandatory parameters have been defined. Optional parameters used by the segmentation and re-assembly function are listed in 8.4.4.

7.11.2 Procedure

A WTP provider MAY support only a subset of all parameters. The parameters are transported in the variable part of the PDU header by using the Option TPI. The first octet of the Option TPI identifies the parameter and the following octets contains the value of the parameter. A WTP provider not supporting a parameter ignores it and returns the Error TPI.

7.12 Error Handling

7.12.1 Motivation

When an unrecoverable error is detected during the transaction, the transaction MUST be aborted. Currently no recovery mechanisms have been defined.

7.12.2 Protocol Data Units

The following PDU is used:

1. Abort PDU

7.12.3 Procedure

When an error occurs in the WTP provider during a transaction, the transaction MUST be aborted with an appropriate Abort reason and the local WTP user informed. The abort procedure is described in a separate section.

7.13 Version Handling

7.13.1 Motivation

A WTP provider receiving an invoke message with a higher version number than what is supported MUST abort the transaction.

7.13.2 Protocol Data Units

The following PDUs and parameters are used:

1. Invoke PDU
2. Abort PDU

7.13.3 Procedure

The Initiator indicates its version in the version field of the Invoke PDU.

If the Responder does not support the version it MUST return an Abort PDU with the Abort Reason set to WTPVERSIONONE. This indicates that the WTP provider supports version one of the WTP protocol.

7.14 Segmentation and Re-assembly (Optional)

7.14.1 Motivation

If the length of a message exceeds the MTU for the current bearer, the message can be segmented by WTP and sent in several packets. When a message is sent as a large number of small packets, the packets MAY be sent and acknowledged in groups. The sender can exercise flow control by changing the size of the packet groups depending on the characteristics of the network.

Selective re-transmission allows for a receiver to request one or multiple lost packets. The alternative is for the sender to re-transmit the entire message, which MAY include packets that have been successfully received. This function minimises the number of packets sent by WTP.

This function is optional. If SAR is not implemented in WTP, this functionality has to be provided by another layer in the stack. For example, in IS-136 the SSAR layer handles SAR, in an IP network IP [RFC791] handles SAR and for GSM SMS/USSD SAR is achieved by using SMS concatenation [GSM0340]. The motivation for implementing WTP SAR is the selective re-transmission procedure, which MAY, if large messages are sent, improve the over-the-air efficiency of the protocol.

An example of this procedure can be found in chapter 10.6.

7.14.2 Procedure for Segmentation

For the sake of brevity only the procedure to segment an invoke message is described here (segmentation of a result message is identical except for the names of the PDUs.)

An invoke message which exceeds the MTU for the network is segmented into an ordered sequence of one Invoke PDU followed by one or more Segmented Invoke PDUs. The initial Invoke PDU has the implicit packet sequence number of zero, the following Segmented Invoke PDU has the packet sequence number one and all the following Segmented Invoke PDUs have packet sequence number that is one greater than the previous (n, n+1, n+2, etc). The Invoke PDU has an "implicit" packet sequence number since this number is not included as a field in the header. The client indicates in the Invoke PDU if the invoke message is segmented by clearing the TTR flag. If the invoke message is segmented, the server counts the Invoke PDU as packet number zero and waits for the following Segmented Invokes PDUs. The packet sequence number MUST NOT wrap. The packet sequence number field is 8 bits; and thus the maximum number of packets is 256.

7.14.3 Procedure for Packet Groups

The packets (Segmented Invoke PDUs and/or Segmented Result PDUs) are sent and acknowledged in groups. The sender MUST NOT send any new packets belonging to the same transaction until the previous packet group has been acknowledged. That is, packet groups are sent according to a stop-and-wait protocol. The sender determines the number of packets for each packet group. The size of a packet group SHOULD be decided with regards to the characteristics of the network and the device. No procedure for determining packet group size has been defined.

The packets in a packet group are sent in one batch. The last packet of the group has the GTR flag set. The last packet of the last packet group of the entire message has the TTR flag set. Since the first group is sent without knowing the status of the receiver the number of packets SHOULD not be too large. When the receiver receives a packet that is not a GTR or TTR packet it MUST store the packet and wait for a new one.

When the receiver receives a packet with the GTR flag set it MUST check whether it has received all packets belonging to that packet group. If the complete packet group has been received the receiver returns an Ack PDU with the PSN TPI containing the Packet Sequence Number of the GTR packet. If one or more packets are missing the receiver returns a Nack PDU including the sequence number(s) of missing packet(s). The missing packets are re-transmitted with the original Packet Sequence Numbers but with the Re-transmission Indicator flag set. When the receiver has received the complete packet group, including those that were re-transmitted, it acknowledges the GTR packet.

When the receiver has received a complete packet group and the last packet has the TTR flag set, it SHOULD be able to re-assemble the complete message.

If the sender has not received an acknowledgement when the re-transmission timer expires, only the GTR/TTR packet is re-transmitted, not the entire packet group.

7.14.4 Procedure for Selective Re-transmission

When a GTR or TTR packet has been received and one or more packets of the packet group are missing, the WTP provider returns the Nack PDU with the sequence number of the missing packet(s). For example, if the receiver has received packet number 2, 3, 5 and 7, and packet number 7 has the GTR flag set, it returns a Nack PDU with packet numbers 4 and 6, indicating missing packets. The packet sequence number of the missing packets are contained in the header part of the Nack PDU.

If the Nack PDU is received with the number of missing packets field set to zero, this means that the entire packet group shall be re-transmitted.

The missing packets are re-transmitted with the original Packet Sequence Numbers. When the sender has re-transmitted the requested packets, it reverts to wait for the original acknowledgement (for the GTR or TTR packet).

When the receiver has received all packets it acknowledges the GTR or TTR packet according to the normal procedure, using the Ack PDU.

A WTP provider not supporting this function **MUST** re-transmit the entire message when one or multiple packets are requested for re-transmission.

When the GTR or TTR packet has been received and one or more packets of the group are missing, the WTP provider **SHOULD** wait for some period of time, such as $\frac{1}{2}$ the median round-trip, before returning the Nack PDU with the sequence numbers of the missing packet(s). If the status of the group changes during the time, i.e. one of the missing packets is received, the waiting time **SHOULD** be reset.

7.15 Extended Segmentation and Re-assembly (Optional)

7.15.1 Motivation

The segmentation and re-assembly function described in section 7.14 is defined for a data transfer where the size exceeds the actual MTU. However, the overall size of data that can be transferred is limited to 256 packets by the fact that the PSN is 8 bits and therefore it does not address large data transfer issues. Thus an extended segmentation and re-assembly procedure has been defined which allows the transmission of a large amount of data, i.e. where the volume exceeds 256 packets. Two key objectives of the procedure are to ensure an efficient data transfer and not to limit the amount of data that can be transferred.

7.15.2 Procedure for Segmentation

For the sake of brevity only the procedure for applying extended segmentation to an invoke message sequence is described here. The Extended SAR mechanism applies to both a Class 1 and Class 2 WTP transaction. Extended SAR function is optional.

At the beginning of each extended SAR transaction a negotiation of the feature takes place. The NumGroups Option TPI is used to advertise the extended segmentation and re-assembly feature availability to the peer. All WTP implementations, which support extended SAR function, **MUST** include this TPI in the initial Invoke PDU. The absence of this TPI in the Ack, Nack or Result PDU will indicate to the receiving WTP that the sender of the PDU does not support the extended SAR. The Initiator **MAY** send the Invoke PDU in a packet group, but the transmission rules for the first group correspond to the SAR without extended mode (section 7.14).

The Initiator **MUST NOT** send more than one group until it has confirmed that the responder supports ESAR. The value of NumGroups indicates the maximal number of outstanding groups in the sliding window. If its value of length is zero, it defines the default value 1 (which means no sliding window has to be taken into use).

From this point on it is assumed that the negotiation was successful and only the normal operating procedure is described. If the negotiation fails, the SAR procedure without extended mode is applied, or the Initiator **MAY** abort the transaction with the NOTIMPLEMENTEDESAR abort code.

The SDU (user data) received in the TR-Invoke.req or in the TR-InvokeData.req service primitives is segmented into a series of one Invoke PDU and additional Segmented Invoke PDUs. So from the TR-Invoke.req SDU one Invoke PDU and a number of Segmented Invoke PDUs are generated and from each subsequent TR-InvokeData.req primitives subsequent only Segmented Invoke PDUs are created.

The initial Invoke PDU has a packet sequence number of zero, the following Segmented Invoke PDU has the packet sequence number of one and all following Segmented Invoke PDUs have one greater than the previous (n, n+1, n+2). When the least significant byte of PSN wraps around, the PSN TPI has to be taken into use. So after the Segmented Invoke PDU that has a PSN of 255 and does not have a PSN TPI a Segmented Invoke PDU with a PSN (least significant byte) of zero and a PSN TPI of length 1 and value 1 will follow.

If the PSN value wraps around it MUST be handled as if monotonically increased. The exception of wraparound will not be mentioned below.

The SDUs received in the TR-Invoke.req and TR-InvokeData.req service primitives are segmented into a series of packets. If the service primitive has the *MoreData* parameter cleared the last packet resulting from that SDU will have the TTR flag set. This denotes the last PDU and the end of the extended invoke. After that the result can be retrieved. If the Frame Boundary flag in an invoked service primitive is set, a Frame Boundary TPI is attached to the first packet resulting from that SDU and it marks the beginning of a new partial invoke message. The very last packet having the TTR flag set MUST NOT be appended with the SDU Boundary TPI. The TTR flag implies a frame and SDU boundary as well.

If the service primitive MoreData parameter is set, the WTP provider will attach the SDU Boundary TPI to the last GTR packet of the last group of the SDU indicating the peer the end of the SDU.

Ordering of service primitives MUST be enforced by the WTP provider. That is, confirm service primitives MUST be passed to the WTP user in the order that the request were made. Inside WTP layer the sequence of the data flow is ensured by the numbering of the packets. WTP is a serial channel, i.e. data MUST be delivered to the peer WTP user in the same order as submitted to WTP by the local user.

The order in which service primitives within the same transaction are submitted to WTP must be preserved. For example, say that SP1, SP2 and SP3 are service primitives being used to transfer fragments for data of the same transaction and are submitted to WTP in that order by the WTP user. The user data in SP1 must be sent entirely to the WTP peer before the user data for SP2 is sent, and similarly for SP3.

7.15.3 Procedure for Sliding Window

Packets are organized into groups. The last packet of the group has the GTR flag set. The last packet of the last group in the transmission has the TTR flag set and GTR flag cleared. So far this is the same as of the original SAR. If a packet carries the SDU Boundary TPI it MUST have the GTR flag set.

A single group cannot include packets referenced by different high order PSN values. For example, a group cannot include PSN = 255 and PSN = 256.

The packet sequence number PSN MUST always be increased one by one, so there cannot be any holes in the sequence number space. For example, a provider cannot generate PSN 1,2 and 4 without generating PSN 3.

The initiator MAY send a number of neighbouring groups in one batch sending the packets in the order of their extended PSNs. The responder MUST acknowledge each group that has been fully received. While a group is not acknowledged it is counted as outstanding. So at a time there can be more than one groups outstanding. The *window* comprises the amount of groups the peer is able to receive.

The size of the sender window in packets is equal to

$$\text{NumGroups} \times \lfloor \text{Maximum Group/BearerPacketSize} \rfloor$$

where $\lfloor \]$ means integer part of the division. The number of outstanding groups cannot exceed the value that has been received in the last *NumGroups TPI*. The size of any group created cannot exceed the value received in the *Maximum Group TPI* divided by the maximal packet size value of the underlying bearer network, so the formula above gives the actual size of the window measured in packets. These Maximum Group and NumGroups TPI values can be used for flow control and can be sent multiple times during a transaction. The initiator can start to transmit the next group only if the size of the next group, when added to the size of the groups are currently outstanding, does not exceed the negotiated window size. So the operation is sliding window based and the measure of sliding is one group. The Maximum Group Option TPI can be set in any PDU sent by the receiver to control the maximum number of packets in a group and so indirectly the window size. If no Maximum Group Option TPI was given during the extended SAR transfer, it gets a default value. This value defines the minimum size of data in bytes the receiver is able to receive by default.

The receiver can practise flow control by sending a NumGroups Option TPI (in Ack or Nack PDUs) having a value of zero. This will prevent the sender from starting the transmission of a new group, once all the packets in the outstanding groups have been successfully acknowledged. If there are no outstanding groups (ie. The transaction is idle for the time being), the NumGroups Option TPI can be transmitted by resending the latest Ack PDU. Reopen of the channel MUST be done by sending larger than 0 value in a NumGroups TPI in any appropriate PDU (Nack or Ack). The sender MAY clear its

retransmission counter each time it receives the NumGroups TPI set to zero and MAY keep sending probe packets in order to keep the transaction alive. If the receiver closes the window (by the NumGroups TPI zero), the sender has to start its retransmission timer and after the expiration it has to send a dataless one-packet group to the receiver to stimulate it to send its status, i.e. to check whether the window should be reopened. The probe packet sent by the sender must be the packet for which the Ack to reopen the window is expected. For example, if the receiver closed the window using the Ack for PSN 12, the sender has to retransmit packet PSN 12 as a probe packet. This packet does not contain data. The receiver reopens the window by sending an Ack for PSN 12 with NumGroups set to a value greater than zero.

The receiver may also send the closing TPI in the Nack PDU and it means the same: it will stop only the sliding of the window.

If the initiator supports a sliding window when acting as a receiver, i.e. more than 1 group can be sent before an Ack received, the initiator includes the NumGroups TPI value more than 1 in the Invoke PDU. This is used by the responder when sending Segmented Result PDUs to the initiator; it defines the number of groups of Segmented Results that can be sent to the initiator before an Ack received. The responder sending the NumGroups TPI value more than 1 in the Ack, Nack or Result PDU which indicates the number of groups of Segmented Invokes, which the initiator may send to the responder without waiting for an acknowledgement. The NumGroups TPI with value of length zero included in the Ack PDU signals, that the initiator MUST NOT send more than 1 group before waiting for an acknowledgement.

The initiator MAY send a number of Segmented Invoke PDUs after the Invoke has been sent and before the Ack is received from the responder. Since the Ack has not been received yet, the initiator does not know the status of the responder and therefore the group size SHOULD not be too large. The initiator MUST NOT send multiple groups at this point.

7.15.4 Procedure for Reliability

The sender cannot have a window that has more than 255 packets pending

When the receiver receives a packet with the GTR or TTR flag set or any of the groups becomes complete it MUST check whether all of the groups of the receiver window are complete or not. If every group is complete it returns an Ack PDU with the PSN of the packet with the highest extended PSN (the acknowledgement is cumulative) with the same values as the received packet has. If any of the groups is incomplete it MUST send a Nack PDU that:

- MUST include all the least significant PSN bytes of all the missing packets (holes in the packet sequence) within the window.
- MUST attach a PSN TPI having the value of the PSN of the packet with the highest extended PSN.

Note that the receiver MUST NOT send Ack PDU if there is any incomplete group.

The sender upon receiving this Nack MUST go through the list of missing packets and MUST interpret it in the following way:

If the least significant byte of the PSN of the missing packet is lower than, or the same as, the least significant byte of the value carried in the PSN TPI attached to the Nack PDU then the high order PSN bytes for the missing packet are the same as carried by the PSN TPI attached to the Nack PDU. Otherwise it is assumed to be in the previous PSN high order domain so the high order PSN bytes of the missing packet will be in this case one lower than the value carried in the PSN TPI attached to the Nack PDU.

$$\text{PSN}(\text{LSB})_{\text{missing}} \leq \text{PSN}(\text{LSB})_{\text{Nack}} \Rightarrow \text{PSN}(\text{HOB})_{\text{missing}} := \text{PSN}(\text{HOB})_{\text{Nack}}$$

$$\text{PSN}(\text{LSB})_{\text{missing}} > \text{PSN}(\text{LSB})_{\text{Nack}} \Rightarrow \text{PSN}(\text{HOB})_{\text{missing}} := \text{PSN}(\text{HOB})_{\text{Nack}} - 1$$

where LSB means Least Significant Byte and HOB designates the High Order Bytes.

If upon receiving a Nack it turns out that for a particular outstanding group every packet is received (no missing packet for the group is listed in the Nack and extended PSN of the Nack is not lower than the extended PSN of the trailing packet of the group), the receiver MUST treat the group as acknowledged. It can be removed at any time if the continuity of the window can be preserved. See example in section 10.6.7. Thus a Nack may acknowledge implicitly a group.

The sender upon sending the last packet of the sending window MUST start the retransmission timer. If the retransmission timer expires it has to increase the retransmission counter and MUST send the GTR or TTR packet which has the highest

extended PSN amongst the unacknowledged packets in order to force the receiver to send an Ack or Nack PDU. If the counter reaches the maximum the transaction **MUST** be aborted.

The sender - supporting sliding window mechanism - **MUST** be prepared to the situation where two Nacks are received within a very short time period listing some common packets. This can happen if there is a lost packet and the receiver sends Nacks as an answer to two subsequent GTR packets. The sender **SHOULD** run a retransmission hold-off timer to avoid excessive retransmissions when multiple Nacks for the same group are received in a short interval. One possible procedure for managing this timer is described in Appendix B. Moreover the sender **MUST** use exponential back-off (see section 9.4.1) when retransmitting.

8. Structure and Encoding of Protocol Data Units

8.1 General

A Protocol Data Unit, PDU, contains an integer number of octets and consists of:

- a) the header, comprising:
 1. the fixed part
 2. the variable part
- b) the data, if present

The fixed part of the headers contains frequently used parameters and the PDU code. The length and the structure of the fixed part are defined by the PDU code. The following PDU types are currently defined:

Table 11 WTP PDU Types

PDU Type	PDU Code
* NOT ALLOWED *	0x00 (Note 1)
Invoke	0x01
Result	0x02
Ack	0x03
Abort	0x04
Segmented Invoke	0x05 (Note 2)
Segmented Result	0x06 (Note 2)
Negative Ack	0x07 (Note 2)

The variable part is used to define less frequently used parameters. Variable parameters are carried in Transport Information Items, TPI.

The very first bit of the fixed header indicates whether the PDU has a variable header or not. The length of the fixed header is given by the PDU type. The variable header consists of TPIs. Every TPI has a length field for its own length. The very first bit of each TPI indicates whether it is the last TPI or not.

Network Octet order for multi-octet integer values is “big-endian”. In other words, the most significant octet is transmitted on the network first followed subsequently by the less significant octets.

The left most bit (bit number 0) of an octet or a bit field is the most significant. Bit fields described first are placed in the most significant bits of the octet. The transmission order in the network is determined by the underlying transport mechanism

Note 1) If the first octet of a datagram is 0x00, it will be interpreted as if the datagram contains multiple concatenated PDUs. See section on Encoding of Concatenated PDUs.

Note 2) This PDU is only applicable if the optional Segmentation and Re-assembly function is implemented.

8.2 Common Header Fields

8.2.1 Continue Flag (CON)

As the first bit of the fixed portion of the header, the Continue Flag indicates the presence of any TPIs in the variable part. If the flag is set, there are one or more TPIs in the variable portion of the header. If the flag is clear, the variable part of the header is empty.

This flag is also used as the first bit of a TPI, and indicates whether the TPI is the last of the variable header. If the flag is set, another TPI follows this TPI. If the flag is clear, the octet after this TPI is the first octet of the user data.

8.2.2 Group Trailer (GTR) and Transmission Trailer (TTR) Flag

When segmentation and re-assembly is implemented the TTR flag is used to indicate the last packet of the segmented message, the GTR flag is used to indicate the last packet of a packet group.

Table 12 GTR/TTR flag combinations

GTR	TTR	Description
0	0	Not last packet
0	1	Last packet of message
1	0	Last packet of packet group
1	1	Segmentation and Re-assembly NOT supported

The default setting SHOULD be GTR=1 and TTR=1, that is, WTP segmentation and re-assembly not supported. In the case where a message uses segmentation, if the TTR flag is set in the last segment, then the GTR flag must be ignored.

8.2.3 Packet Sequence Number

This is used by the PDUs belonging to the segmentation and re-assembly function. This number indicates the position of the packet in the segmented message.

8.2.4 PDU Type

The PDU Type field indicates what type of WTP PDU the PDU is (Invoke, Ack, etc). This provides information to the receiving WTP provider as to how the PDU data SHOULD be interpreted and what action is required.

8.2.5 Reserved (RES)

All reserved bits are to be set to the value 0x00 unless otherwise specified.

8.2.6 Re-transmission Indicator (RID)

Enables the receiver to differentiate between packets duplicated by the network and packets re-transmitted by the sender. In the original message the RID is clear. When the message gets re-transmitted the RID is set.

8.2.7 Transaction Identifier (TID)

The TID is used to associate a packet with a particular transaction.

8.3 Fixed Header Structure

8.3.1 Invoke PDU

Table 13 Structure of Invoke PDU

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	PDU Type = Invoke				GTR	TTR	RID
2	TID							
3								
4	Version	TIDnew	U/P	RES	RES	TCL		

8.3.1.1 Transaction Class, TCL

The Initiator indicates the desired transaction class in the invoke message.

Table 14 Encoding of Class field

Class	TCL
0	0x00
1	0x01
2	0x02

The transaction classes are explained in separate chapter.

8.3.1.2 TIDnew Flag

This is set when the Initiator has "wrapped" the TID value; that is, the next TID will be lower than the previous. When the Responder receives the Invoke PDU and the TIDnew flag is set, it invalidates its cached TID value for this Initiator.

8.3.1.3 Version

The current version is 0x00.

8.3.1.4 U/P Flag

When this flag is set it indicates that the Initiator requires a User acknowledgement from the server WTP user. This means that the WTP user confirms every received message.

When this flag is clear the WTP provider MAY respond to a message without a confirmation from the WTP user.

8.3.2 Result PDU

Table 15 Structure of Result PDU

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	PDU Type = Result				GTR	TTR	RID
2	TID							
3								

8.3.3 Acknowledgement PDU

Table 16 Structure of Ack PDU

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	PDU Type = Acknowledgement				Tve/Tok	RES	RID
2	TID							
3								

8.3.3.1 Tve/Tok Flag

In the direction from the responder to the initiator the Tve (TID Verify) means: -"Do you have an outstanding transaction with this TID?". In the opposite direction the Tok (TID OK) flag means: -"I have an outstanding transaction with this TID!"

8.3.4 Abort PDU

Table 17 Structure of Abort PDU

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	PDU Type = Abort				Abort type		
2	TID							
3								
4	Abort reason							

8.3.4.1 Abort type and Abort reason

Currently the following abort types are specified:

Table 18 WTP Abort Types

Abort type	Code	Description
Provider (PROVIDER)	0x00	The abort was generated by the WTP provider itself. The abort reason is specified below.
User (USER)	0x01	The abort was generated by the WTP user. The abort reason is provided to the WTP provider by the WTP user.

Abort reasons from the WTP provider

The following abort reasons are specified:

Table 19 WTP Provider Abort Codes

Abort reason (PROVIDER)	Code	Description
Unknown (UNKNOWN)	0x00	A generic error code indicating an unexpected error .
Protocol Error (PROTOERR)	0x01	The received PDU could not be interpreted. The structure MAY be wrong.
Invalid TID (INVALIDTID)	0x02	Only used by the Initiator as a negative result to the TID verification.
Not Implemented Class 2 (NOTIMPLEMENTEDCL2)	0x03	The transaction could not be completed since the Responder does not support Class 2 transactions.
Not Implemented SAR (NOTIMPLEMENTEDSAR)	0x04	The transaction could not be completed since the Responder does not support SAR.
Not Implemented User Acknowledgement (NOTIMPLEMENTEDUACK)	0x05	The transaction could not be completed since the Responder does not support User acknowledgements.
WTP Version One (WTPVERSIONONE)	0x06	Current version is 1. The initiator requested a different version that is not supported.
Capacity Temporarily Exceeded (CAPTEMPEXCEEDED)	0x07	Due to an overload situation the transaction can not be completed.
No Response (NORESPONSE)	0x08	A User acknowledgement was requested but the WTP user did not respond
Message too large (MESSAGETOOLARGE)	0x09	Due to a message size bigger than the capabilities of the receiver the transaction cannot be completed.
Not Implemented Extended SAR (NOTIMPLEMENTEDESAR)	0x0A	The transaction could not be completed since the Responder does not support extended SAR.

Abort reasons from the WTP user

The abort reasons from the WTP user are given to the local WTP provider in the T-TRAbort request primitive. The abort reason is specific to the WTP user. For example, if the WTP user is WSP, abort codes defined in [WSP] can be used.

8.3.5 Segmented Invoke PDU (Optional)

Table 20 Structure of Segmented Invoke PDU

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	PDU Type = Segmented Invoke				GTR	TTR	RID
2	TID							
3								
4	Packet Sequence Number							

8.3.6 Segmented Result PDU (Optional)

Table 21 Structure of Segmented Result PDU

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	PDU Type = Segmented Result				GTR	TTR	RID
2	TID							
3								
4	Packet Sequence Number							

8.3.7 Negative Acknowledgement PDU (PDU)

Table 22 Structure of Negative Acknowledgement PDU

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	PDU Type = Negative Ack				Reserved		RID
2	TID							
3								
4	Number of Missing Packets = N							
5	Packet Sequence Number(s) of Missing Packets							
...								
4+N								

8.3.7.1 Number of Missing Packets

Indicates the requested number of missing packets. If 0x00, this means that the entire packet group shall be re-transmitted.

8.3.7.2 Packet Sequence Number(s) of Missing Packets

List of packet sequence number for the request packets.

8.4 Transport Information Items

8.4.1 General

The variable part of the PDU can consist of one or several Transport Information Items, TPIs. The length field of a TPI can be 2 or 8 bits.

The long TPI (8 bits length) has the following structure:

Table 23 Long TPI structure

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	TPI Identity				1	RES	RES
2	TPI Length = N							
3	TPI Data							
...								
2+N								

The short TPI (2 bits length) is structured as

Table 24 Short TPI Structure

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	TPI Identity				0	TPI Length = M	
2	TPI Data							
...								
1+M								

In the above tables, $N=0..255$ and $M=0..3$. The data field of the TPI MUST contain an integer number of octets. In theory the maximum length of a TPI is 255 octets, however, it is also limited by the MTU size of the bearer network and the number of, and length of, other TPIs in the same PDU header.

The following TPIs are currently defined:

Table 25 Encoding of TPIs

TPI	TPI Identity	Comment
Error	0x00	
Info	0x01	
Option	0x02	
Packet Sequence Number (PSN)	0x03	Note 1
SDU Boundary	0x04	Note 2
Frame Boundary	0x05	Note 2

Note 1) This TPI is only applicable if the optional segmentation and re-assembly function is implemented.

Note 2) This TPI is only applicable if the optional extended segmentation and re-assembly function is implemented.

8.4.2 Error TPI

The Error TPI is returned to the sender of an erroneous or unsupported TPI. Currently the following error codes have been defined:

Table 26 Encoding of Error TPI

Error	Code	Argument
Unknown TPI	0x01	TPI Identity of unknown TPI
Known TPI, unknown content	0x02	TPI Identity and first octet of content

Depending on the ErrorCode the Error TPI can have a different structure.

Table 27 Structure of Error TPI (UNKNOWN)

Bit/Octet	0	1	2	3	4	5	6	7
1	CON	TPI Identity = 0x00				0	TPI Length = 0x01	
2	ErrorCode = 0x01				Bad TPI Identity			

Table 28 Structure of Error TPI (KNOWN)

Bit/Octet	0	1	2	3	4	5	6	7	
1	CON	TPI Identity = 0x00				0	TPI Length = 0x02		
	Error Code = 0x02				Bad TPI Identity				
3	First octet of TPI								

Note that this TPI is mandated to support by a WTP provider. Consequently, the WTP provider MUST also be able to recognise the general structure of a TPI.

8.4.3 Info TPI

This TPI is used to piggyback a small amount of data in the variable part of the PDU header. For example, the data can be performance measurements or statistical data.

The structure of the Info TPI is illustrated below.

Table 29 Structure of Info TPI

Bit/Octet	0	1	2	3	4	5	6	7	
1	CON	TPI Identity				0	TPI Length = N		
2	Information								
...									
1+N									

The above table shows the Info TPI as short TPI. If more information MUST be sent, the long TPI can be used.

8.4.4 Option TPI

The Option TPI is used to transfer parameters between two WTP entities. The parameter carried in the Option TPI is valid for the lifetime of the transaction. The following options are currently defined:

Table 30 Encoding of Option TPI

Option	Identity	Description	Comment
Maximum Receive Unit	0x01	This parameter is used by the Initiator to advertise the maximum unit of data in bytes that can be received in the result	
Total Message Size	0x02	This parameter can be sent in the first packet of a segmented message to inform the receiver about the total message size in bytes	Note 1
Delay Transmission Timer	0x03	This parameter can be sent in the Ack PDU when a packet group is acknowledged. The receiver MUST NOT send the next packet group until the specified time has elapsed. The time is in 1/10 seconds.	Note 1
Maximum Group	0x04	This parameter can be used by either transaction party to advertise the maximum group size, which can be received. The parameter indicates the maximum size in bytes of data in a single group. The default is 1405.	Note 1
Current TID	0x05	This parameter may be sent with an Ack PDU when a 3-way-handshake is requested by the server, i.e. the Verify flag is set. The use of the parameter is optional, and the interpretation by the client implementation dependent. When used, the value shall be the value cached by the server (LastTID).	
No Cached TID	0x06	This parameter may be sent with an Ack PDU when a 3-way-handshake is requested by the server, i.e. the Verify flag is set. The use of the parameter is optional, and the interpretation by the client implementation dependent. When used, the parameter indicated that there is no cached LastTID	
NumGroups	0x07	This parameter can be used by either transaction party to advertise the maximum number of outstanding groups which can be received. This value is expressed in number of groups. This TPI is used to signal the extended segmentation and re-assembly feature availability to the peer. A WTP implementation, which supports extended SAR function, MUST include this TPI in the initial Invoke PDU. The absence of this TPI in the Ack, Nack or Result PDU will indicate to the receiving WTP that the sender of the PDU does not support the extended SAR. The value of length zero indicates the usage of the default value, 1.	Note 2

Note 1) This parameter is only applicable if the optional segmentation and re-assembly function is implemented.

Note 2) This parameter is only applicable if the optional extended segmentation and re-assembly function is implemented.

The structure of the Option TPI is illustrated below.

Table 31 Structure of Option TPI

Bit/Octet	0	1	2	3	4	5	6	7	
1	CON	TPI Identity				0	TPI Length = N		
2	Option Identity								
3	Option Value								
...									
1+N									

The above table shows the Option TPI as a short TPI. If more information must be sent, the long TPI can be used.

8.4.5 Packet Sequence Number TPI (Optional)

The Ack PDU does not have a Packet Sequence Number (PSN) field. When Segmentation and Re-assembly is used this TPI is attached to the variable part of the Ack PDU header. The PSN included in the Ack PDU is the PSN of the acknowledged packet (GTR or TTR packet).

In case of extended SAR, the range of PSN is 24 bits, therefore the TPI length=M, where M may vary between 1 and 3. This TPI provides the high-order bytes of the sequence number, which are not included in the PDU itself. This TPI provides the high-order bytes of the sequence number in the case of Segmented Invoke/Result PDUs and the complete PSN when sent in the Nack or Ack PDU. PSN will be reset to zero once it reaches 2^{24} (wraps around).

Table 32 Structure of Packet Sequence Number TPI

Bit/Octet	0	1	2	3	4	5	6	7	
1	CON	TPI Identity = Packet Sequence Number				0	TPI Length = M		
2	First Byte (MSB) of PSN								
...	Optional Second Byte of PSN								
1 + M	Optional Third Byte of PSN								

8.4.6 SDU Boundary TPI

This TPI is used only in the optional extended segmentation and re-assembly function to put a framing boundary into the SDU data transmitted to the peer. This TPI will be attached to the GTR packet fix header of the last group by the protocol automatically, if the service primitive parameter MoreData is set.

Table 33 Structure of SDU Boundary TPI

Bit/Octet	7	6	5	4	3	2	1	0	
1	CON	TPI Identity = SDU Boundary				0	TPI Length = 0		

8.4.7 Frame Boundary TPI


This TPI is used only in the optional extended segmentation and re-assembly function to put a user controlled framing boundary at the beginning of SDU data transmitted to the peer. This TPI can be attached only to the very first packet's fixed header.

Table 34 Structure of Frame Boundary TPI

Bit/Octet	7	6	5	4	3	2	1	0	
1	CON	TPI Identity = Frame Boundary				0	TPI Length = 0		

8.5 Structure of Concatenated PDUs

One or more WTP Protocol Data Units (PDUs) MAY be contained in one datagram Service Data Unit (SDU). This is illustrated below.



The following table represents a datagram SDU with one WTP PDU. The PDU including header and data is N octets.

Table 35 WTP PDU without Concatenation

Bit/Octet	0	1	2	3	4	5	6	7
1	WTP PDU							
...								
N								

The following table represents two WTP PDUs concatenated in the same SDU of the bearer network. The first PDU is N octets and the second is M octets.

Table 36 Concatenated WTP PDUs

Bit/Octet	0	1	2	3	4	5	6	7
1	Concatenation Indicator = 0x00							
2	0	WTP PDU Length = N						
3	WTP PDU							
...								
N+2								
N+3	0	WTP PDU Length = M						
N+4	WTP PDU							
...								
N+M+3								

The concatenation indicator is used to indicate that the SDU contains multiple WTP PDUs. The number of PDUs is limited only by the maximum size of the SDU.

The PDU Length field can be 7 or 15 bits. If the first bit in the PDU Length field is set, the length field is 15 bits, if not, it is 7 bits. This means that the PDU Length field takes up 8 or 16 bits depending on whether the first bit is set or not. In the above table the first bit is 0 and thus the length field is 7 bits.

9. State Tables

9.1 General

This chapter defines state tables for the core WTP protocol without the optional segmentation and re-assembly function.

9.2 Event Processing

The WTP provider initiating a transaction is called the Initiator. The WTP provider responding to an initiated transaction is called the Responder. An implementation of the WTP protocol is not required to have both Initiator and Responder functionality. For example, if the WTP user is the WSP client protocol the WTP provider MAY only support initiation of transactions, that is, no Responder functionality. See Static Conformance Requirements appendix for details on what MUST be implemented in order to conform to the standard.

The interface to the next higher layer is defined by the WTP service primitives. The next lower layer is typically a datagram service and the only service primitives are the UnitData indication and requests. The request and response service primitives from the next higher layer together with indication primitive from the next lower layer are termed *events*. If multiple PDUs are concatenated in the SDU from the next lower layer, they MUST be separated and dispatched to the transactions. In addition to the external events, there will also be internal events such as timer expirations and errors.

An event is validated before it is processed. The following tests are performed, and if no action is taken, the event is processed according to the state tables.

Table 37 Test of incoming events

Test	Action
UnitData.ind on the Responder: Invoke PDU	Create a new transaction
UnitData.ind on the Initiator: Ack PDU with the TIDve flag set, no matching outstanding transaction	Send Abort PDU (INVALIDTID)
UnitData.ind: Ack PDU, Nack PDU, Result PDU or Abort PDU, no matching outstanding transaction	Ignore
Illegal PDU type or erroneous header structure	Refer to entry for `RcvErrorPdu` in state tables
Buffer overflow or out-of-memory errors	Send Abort PDU (CAPTEMPEXCEED)
UnitData.ind on the Responder: Invoke PDU requesting Class 2 transaction and Class 2 is not supported	Send Abort PDU (NOTIMPLEMENTEDCL2)
UnitData.ind on the Responder: Invoke PDU using SAR and SAR is not supported	Send Abort PDU (NOTIMPLEMENTEDSAR)
UnitData.ind on the Responder: Invoke PDU requesting User acknowledgement and User acknowledgement is not supported	Send Abort PDU (NOTIMPLEMENTEDUACK)
UnitData.ind on the Responder: Invoke PDU with Version != 0x00	Send Abort PDU (WTPVERSIONONE)
UnitData.ind on the Responder: Invoke PDU when no more transaction requests can be accepted	Ignore.

9.3 Actions

9.3.1 Timers

The following timer actions can be used in the state tables:

Start timer, <value>

Starts the timer with the specified interval value. If the timer is already running, it is re-started with the new value.

Stop timer

Stop the timer without generating an event.

9.3.2 Counters

The following counter actions can be used in the state tables:

Reset counter

Set the counter to zero.

Increment counter

Increment the counter with one.

9.3.3 Messages

The following message actions can be used in the state tables:

Queue (Time T)

Queuing a PDU causes it to be queued for eventual delivery. The message **MUST NOT** be queued for longer time than T time units and is queued only until the already started timer T expires. The timer is not restarted

Send

Sending a PDU causes it and any queued PDUs to be sent immediately.

The queuing mechanism is used to concatenate messages from different transactions. This can be seen as a concatenation layer that operates below the transaction state machine. The realisation of the concatenation layer is implementation dependent and not specified.

9.4 Timers, Counters and Variable

9.4.1 Timers

The following timers are used by WTP:

Table 38 WTP Timers

Timer	Description
Transaction timer	Each transaction has a timer associated with it. The timer is used for both the retry interval, acknowledgement interval and wait timeout interval.

A timer can be started with different timer values depending on the type of transaction and the current state of the transaction. Timer values are grouped according to their purpose. This is shown in the following table.

Table 39 WTP Timer Intervals

Timer interval (name)	Description
Acknowledgement interval (A)	This sets a bound for the amount of time to wait before sending an acknowledgement.
Retry interval (R)	This sets a bound for the amount of time to wait before re-transmitting a PDU.
Wait timeout interval (W)	This sets a bound for the amount of time to wait before state information about a transaction is released. Only Class 2 Initiator and Class 1 Responder

The Retry interval **MAY** be implemented as an array with the re-transmission counter as an index, R[RCR]. An exponential back off algorithm can be implemented by populating R[] with values that are increasing successive powers of 2.

The value of a timer interval depends on the following parameters:

- The characteristics of the bearer network
- The transaction class

- The state of the transaction (which message is being re-tried or acknowledged)

Timer interval values for different bearer networks can be found in Appendix A.

9.4.2 Counters

The following counters are used by the WTP:

Table 40 WTP Counters

Counter (name)	Description
Re-transmission Counter (RCR)	This set a bound for the maximum number of re-transmissions of any PDU. The max value is defined as RCR_MAX.
Acknowledgment Expiration Counter (AEC)	This sets a bound for the maximum number of times the transaction timer, initialised with the acknowledgement interval, is allowed to expire and be re-started before the transaction is aborted. The max value is defined as AEC_MAX.

9.4.3 Variables

The following variables are used by WTP at the Initiator and Responder.

Table 41 WTP Variables

WTP variables			
Variables	Type	Description	
GenTID	Uint16	The TID to use for the next transaction. Incremented by one for every initiated transaction.	Global Only Initiator
SendTID	Uint16	The TID value to send in all PDUs in this transaction	One per transaction
RcvTID	Uint16	The TID values expected to receive in every PDU in this transaction. $RcvTID = SendTID \text{ XOR } 0x8000$	One per transaction
LastTID	Uint16	The last received TID from a certain remote host	One per remote host Only Responder
HoldOn	BOOL	True if HoldOn acknowledgement has been received	One per class 2 transaction
Uack	BOOL	True if User Acknowledgement has been requested for this transaction	One per transaction

The Uint16 type is an unsigned 16-bit integer. The BOOL type is a Boolean value that only can take the value of True or False.

9.5 WTP Initiator

WTP Initiator NULL			
Event	Condition	Action	Next State
TR-Invoke.req	Class == 2 1	SendTID = GenTID Send Invoke PDU Reset RCR Start timer, R [RCR] Uack = False	RESULT WAIT
	Class == 2 1 UserAck	SendTID = GenTID Send Invoke PDU Reset RCR Start timer, R [RCR] Uack = True	
	Class == 0	SendTID = GenTID Send Invoke PDU	NULL

WTP Initiator RESULT WAIT			
Event	Condition	Action	Next State
TR-Abort.req		Abort transaction Send Abort PDU (USER)	NULL
RcvAck	Class == 2 HoldOn == False	Stop timer Generate T-TRInvoke.cnf HoldOn = True	RESULT WAIT
	Class == 2 HoldOn == True	Ignore	RESULT WAIT
	Class == 1	Stop timer Generate T-TRInvoke.cnf	NULL
	TIDve Class == 2 1 RCR < MAX_RCR	Send Ack(TIDok) Increment RCR Start timer, R [RCR]	RESULT WAIT
	TIDve Class == 2 1	Ignore	RESULT WAIT
RcvAbort		Abort transaction Generate TR-Abort.ind	NULL
RcvErrorPDU		Abort Transaction Send Abort PDU (PROTOERR) Generate TR-Abort.ind	NULL
TimerTO_R	RCR < MAX_RCR	Increment RCR Start timer R [RCR] Send Invoke PDU	RESULT WAIT
	RCR < MAX_RCR, Ack(TIDok) already sent	Increment RCR Start timer R [RCR] Send Ack(TIDok)	RESULT WAIT
	RCR == MAX_RCR	Abort transaction Generate TR-Abort.ind	NULL
RcvResult	Class == 2 HoldOn == True	Stop timer Generate TR-Result.ind Start timer, A	RESULT RESP WAIT
	Class == 2 HoldOn == False	Stop timer Generate TR-Invoke.cnf Generate TR-Result.ind Start timer, A	

WTP Initiator RESULT RESP WAIT (class 2 only)			
Event	Condition	Action	Next State
TR-Result.res		Queue(A) Ack PDU Start timer, W	WAIT TIMEOUT
	ExitInfo	Queue(A) Ack PDU with Info TPI Start timer, W	
RcvAbort		Abort transaction Generate T-TRAbort.ind	NULL
TR-Abort.req		Abort transaction Send Abort PDU (USER)	
RcvErrorPDU		Abort Transaction Send Abort PDU (PROTOERR) Generate TR-Abort.ind	NULL
RcvResult		Ignore	RESULT RESP WAIT
TimerTO_A	AEC < AEC_MAX	Increment AEC Start timer, A	RESULT RESP WAIT
	AEC == AEC_MAX	Abort transaction Send Abort PDU (NORESPONSE)	NULL
	Uack == False	Queue(A) Ack PDU Start timer, W	WAIT TIMEOUT

WTP Initiator WAIT TIMEOUT (class 2 only)			
Event	Condition	Action	Next State
RcvResult	RID=0	Ignore	WAIT TIMEOUT
RcvResult	RID=1	Send Ack PDU	WAIT TIMEOUT
RcvResult	RID=1, ExitInfo	Send Ack PDU with info TPI	WAIT TIMEOUT
RcvAbort		Abort transaction Generate T-TRAbort.ind	NULL
RcvErrorPDU		Abort Transaction Send Abort PDU (PROTOERR) Generate TR-Abort.ind	NULL
TimerTO_W		Clear Transaction	NULL
TR-Abort.req		Abort transaction Send Abort PDU (USER)	NULL

9.6 WTP Responder

WTP Responder LISTEN			
Event	Condition	Action	Next State
RcvInvoke	Class == 2 1 Valid TID U/P flag	Generate TR-Invoke.ind Start timer, A Uack = True	INVOKE RESP WAIT
	Class == 2 1 Valid TID	Generate TR-Invoke.ind Start timer, A Uack = False	
	Class == 0	Generate TR-Invoke.ind	LISTEN
	Class == 2 1 Invalid TID	Send Ack(TIDve)	TIDOK WAIT
RcvErrorPDU		Send Abort PDU (PROTOERR)	LISTEN

WTP Responder TIDOK WAIT			
Event	Condition	Action	Next State
RcvAck	Class == 2 1 TIDok	Generate TR-Invoke.ind Start timer, A	INVOKE RESP WAIT
RcvErrorPDU		Send Abort PDU (PROTOERR) Abort Transaction	LISTEN
RcvAbort		Abort transaction	LISTEN
RcvInvoke	RID=0	Ignore	TIDOK WAIT
	RID=1	Send Ack(TIDve)	TIDOK WAIT

WTP Responder INVOKE RESP WAIT			
Event	Condition	Action	Next State
TR-Invoke.res	Class == 1 ExitInfo	Queue(A) Ack PDU with InfoTPI Start timer, W	WAIT TIMEOUT
	Class == 1	Queue(A) Ack PDU Start timer, W	
	Class == 2	Start timer, A	RESULT WAIT
TR-Result.req		Reset RCR Start timer R[RCR] Send Result PDU	RESULT RESP WAIT
TR-Abort.req		Abort transaction Send Abort PDU (USER)	LISTEN
RcvAbort		Generate TR-Abort.ind Abort transaction	LISTEN
RcvInvoke		Ignore	INVOKE RESP WAIT
RcvErrorPDU		Abort Transaction Send Abort PDU (PROTOERR) Generate TR-Abort.ind	LISTEN
TimerTO_A	AEC < AEC_MAX	Increment AEC Start timer, A	INVOKE RESP WAIT
	AEC == AEC_MAX	Abort transaction Send Abort PDU (NORESPONSE) Generate TR-Abort.ind	LISTEN
	Class == 1 Uack == False	Queue(A) Ack PDU Start timer, W	WAIT TIMEOUT
	Class == 2 Uack == False	Send Ack PDU	RESULT WAIT

WTP Responder RESULT WAIT (class 2 only)			
Event	Condition	Action	Next State
TR-Result.req		Reset RCR Start timer, R[RRCR] Send Result PDU	RESULT RESP WAIT
RcvInvoke	RID=0	Ignore	RESULT WAIT
	RID=1	Ignore	RESULT WAIT
	RID=1, Ack PDU already sent	Resend Ack PDU	RESULT WAIT
RcvErrorPDU		Abort Transaction Send Abort PDU (PROTOERR) Generate TR-Abort.ind	LISTEN
TR-Abort.req		Abort transaction Send Abort PDU (USER)	LISTEN
RcvAbort		Generate T-TRAbort.ind Abort transaction	LISTEN
TimerTO_A		Send Ack PDU	RESULT WAIT

WTP Responder RESULT RESP WAIT (class 2 only)			
Event	Condition	Action	Next State
TR-Abort.req		Abort transaction Send Abort PDU (USER)	LISTEN
RcvAbort		Generate T-TRAbort.ind Abort transaction	LISTEN
RcvAck	TIDok	Ignore	RESULT RESP WAIT
RcvAck		Generate TR-Result.cnf	LISTEN
RcvErrorPDU		Abort Transaction Send Abort PDU (PROTOERR) Generate TR-Abort.ind	LISTEN
TimerTO_R	RCR < MAX_RCR	Increment RCR Send Result PDU Start timer, R [RRCR]	RESULT RESP WAIT
	RCR == MAX_RCR	Generate T-TRAbort.ind Abort transaction	LISTEN

WTP Responder WAIT TIMEOUT (class 1 only)			
Event	Condition	Action	Next State
RcvInvoke	RID=0	Ignore	WAIT TIMEOUT
RcvInvoke	RID=1	Send Ack PDU	WAIT TIMEOUT
RcvInvoke	RID=1, ExitInfo	Send Ack PDU with Info TPI	WAIT TIMEOUT
RevAck	TIDok, RID=1	Send Ack PDU	WAIT TIMEOUT
RevAck	TIDok, RID=1, ExitInfo	Send Ack PDU with Info TPI	WAIT TIMEOUT
RcvErrorPDU		Abort Transaction Send Abort PDU (PROTOERR) Generate TR-Abort.ind	LISTEN
RcvAbort		Abort transaction Generate T-TRAbort.ind	LISTEN
TimerTO_W		Clear Transaction	LISTEN
TR-Abort.req		Abort transaction Send Abort PDU (USER)	LISTEN

10.Examples of Protocol Operation

10.1 Introduction

The examples in this chapters attempt to illustrate and clarify how the protocol operates. For the sake of brevity, only header fields relevant for the specific example are included in the diagrams. Each flag in the Flag field of the PDU header is indicated by one character. The below table shows the different characters that can appear in the examples.

Table 42 Abbreviations Used in the Examples

Abbreviation	Meaning
N	TIDnew flag is set
V	TIDve flag is set
O	TIDok flag is set
U	U/P flag is set
G	GTR flag is set
T	TTR flag is set
TG	Both TTR and GTR flags are set to indicate that SAR is not supported
RID = X	Re-transmission Indicator is X
TID = N	Transaction Identifier is N
c0	The TCL field indicates class 0 transaction
c1	The TCL field indicates class 1 transaction
c2	The TCL field indicates class 2 transaction

Parameters like Abort reason and Error codes are written in clear text, and so are TPIs. For Transaction Identifiers N* is N with the high order bit set; if N = 0x0000 then N* = 0x8000.

10.2 Class 0 Transaction

10.2.1 Basic Transaction

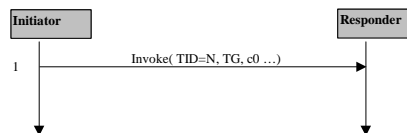


Figure 3 Basic Class 0 Transaction

1. The Initiator initiates a class 0 transaction (c0).

10.3 Class 1 Transaction

10.3.1 Basic Transaction




Figure 4 Basic Class 1 Transaction

1. The Initiator initiates a class 1 transaction (c1).
2. The Responder acknowledges the received invoke message.

10.4 Class 2 Transaction

10.4.1 Basic Transaction




Figure 5 Basic Class 2 Transaction

1. The Initiator initiates a class 2 transaction (c2).
2. The Responder waits for the invoke message to be processed and implicitly acknowledges the invoke message with the Result.
3. The Initiator acknowledges the received result message.

10.4.2 Transaction with “Hold on” Acknowledgement

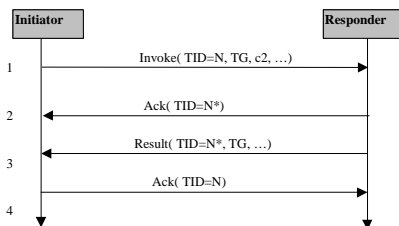


Figure 6 Class 2 Transaction with "hold on" acknowledgement.

1. The Initiator initiates a class 2 transaction (c2).
2. The Responder waits for the invoke message to be processed. The acknowledgement timer at the Responder expires and an "hold-on" acknowledgement is sent to prevent the Initiator from re-transmitting the invoke message.
3. The result is sent to the Initiator
4. The Initiator acknowledges the received result message.

10.5 Transaction Identifier Verification

10.5.1 Verification Succeeds

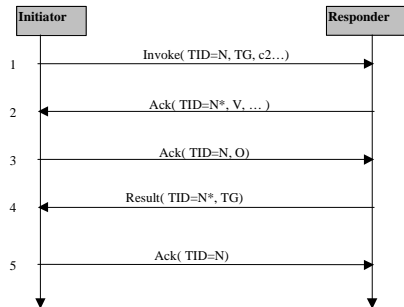


Figure 7 Verification Succeeds

The Responder receives a new invoke message and the TID test fails, this causes the Verification procedure to be invoked. The Responder returns an acknowledgement to the Initiator for a verification of whether it has an outstanding transaction with this TID. In this example, the Initiator has an outstanding transaction with the TID and acknowledges the verification.

10.5.2 Verification Fails




Figure 8 Verification Fails

The invoke message with TID=N is duplicated in the network, or has been delayed. When it arrives, transaction N has already been terminated and the Responder asks the Initiator to verify the transaction. The Initiator aborts the transaction by sending an Abort.

10.5.3 Transaction with Out-of-Order Invoke

An invoke message is delayed in the network. When the message finally arrives to the Responder, the Responder has cached a higher TID value. The Responder initiates a Verification in order to check whether the Initiator still has an invoke message with TID=N outstanding.

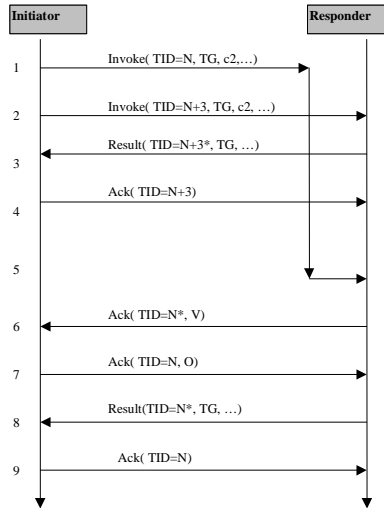


Figure 9 Delayed Invoke Message

Note that the Responder must not replace its cached TID value (N+3) with the lower TID value (N). If the cached TID is moved backwards, old duplicates with higher TID values will erroneously get accepted.

10.6 Segmentation and Re-assembly

This example illustrates a Class 2 transaction using segmentation. The Invoke is segmented and sent in five packets in two packet groups.

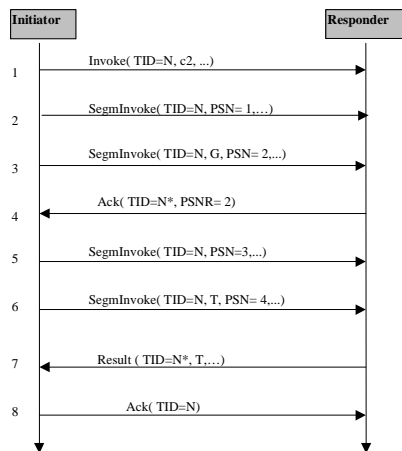


Figure 10 Segmentation of invoke message

The Initiator starts off by sending the first three packets in one batch. The last packet has the GTR flag to trigger an acknowledgement from the Responder. Once the acknowledgement is received by the Initiator the last two packets of the message are sent. The final message has the TTR flag set. After some time, the Responder sends back the result to the Initiator. The Initiator acknowledges the result and the transaction is finished.

Note that the PSN TPI is used for the Packet Sequence Number in the Ack PDU.

10.6.1 Selective Re-transmission

This example illustrates a Class 1 transaction using segmentation. One of the packets in the only packet group is lost and the Responder has to request the packet to be re-transmitted. In the case of SAR, if the TTR flag is set in the last segment, then the GTR flag must be ignored.

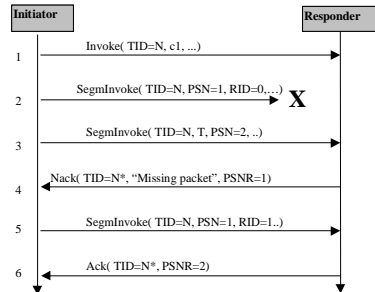


Figure 11 Selective re-transmission

The Initiator starts off by sending the first three packets. The second packet is lost. When the Responder receives the packet with the GTR flag set, it attempts to re-assemble the packet group but fails due to the one missing packet. The Responder returns a Nack to request the missing packet. The Initiator re-transmit the missing packet. The re-transmitted packet has the RID flag set. Once the missing packet has been received by the Responder the message is acknowledged and the transaction is finished.

Note that the PSN TPI is used for the Packet Sequence Number in the Ack PDU.

10.6.2 Re-transmission of the GTR/TTR Packet

This example illustrates a class 1 transaction using segmentation. The last packet is and the initiator has to re-transmit it.




Figure 12 GTR/TTR packet retransmission

The initiator starts off by sending three packets. The third packet is lost. Since the responder does not receive the packet with the GTR or TTR flag set, it can't determine whether the whole packet group has been transmitted or not. After a certain time without receiving any acknowledgement for this group, the initiator re-transmits the last packet of the group. The re-transmitted packet has the RID flag set. Once the responder has received the missing packet, the message is acknowledged and the transaction is finished.

10.6.3 SAR and TID Verification

If TID verification is necessary when receiving a segmented invoke message, it has to take place when the Responder successfully received the first packet group. TID verification acknowledges the successful transmission of the first packet group.

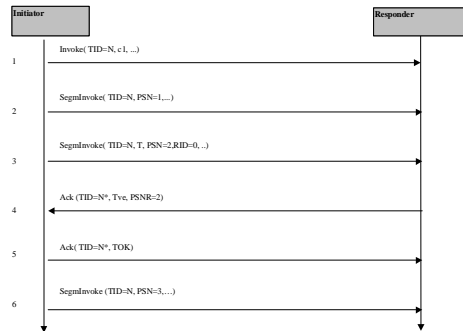


Figure 13 SAR and TID Verification, no missing packets

If some packets were lost during the transmission, first a negative acknowledgement (Nack) MUST be sent. TID verification takes place only after the successful transmission of the whole group.

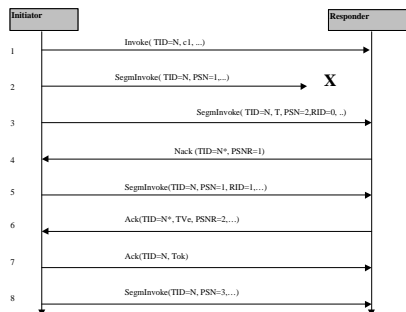


Figure 14 SAR and TID Verification, some packets were lost

If the TID is invalid (e.g. Responder received a network duplicate of a group trailer packet), the Responder will send a negative acknowledgement (Figure 14). The initiator MUST ignore Nack PDUs with invalid TID. The Responder MUST NOT retransmit the Nack PDU. After a reasonable amount of time, the Responder SHOULD remove all data related to the non-existing transaction.

10.6.4 Flow Control Using Option TPI (Maximum Group) Conjointly with SAR

Example situation

- Initiator can receive up to 200 bytes
- Responder can receive up to 100 bytes
- Initiator can send up to 300 bytes in a group
- The size of each packet in a group is 150 bytes

Under these conditions, the initiator would like to send 700 bytes of user data in an Invoke PDU. However the maximum value the initiator can send in a group is 300 bytes, so it sends only two packets in a group containing 300 bytes of data. After the responder receives the two packets, it can send an Ack PDU with the option TPI(Maximum Group) whose maximum group value is 100 bytes. If the initiator receives an Ack PDU with option TPI (Maximum Group), it must re-assign group

size to 100 bytes and sends Invoke PDU whose size if 100 bytes. If the responder sends and Ack PDU with option TPI (Maximum Group) whose value is 300 bytes, the initiator can send 3 packets in a group.

In this way, user data is segmented using option TPI (Maximum Group) and their flow can be controlled by option TPI (Maximum Group).

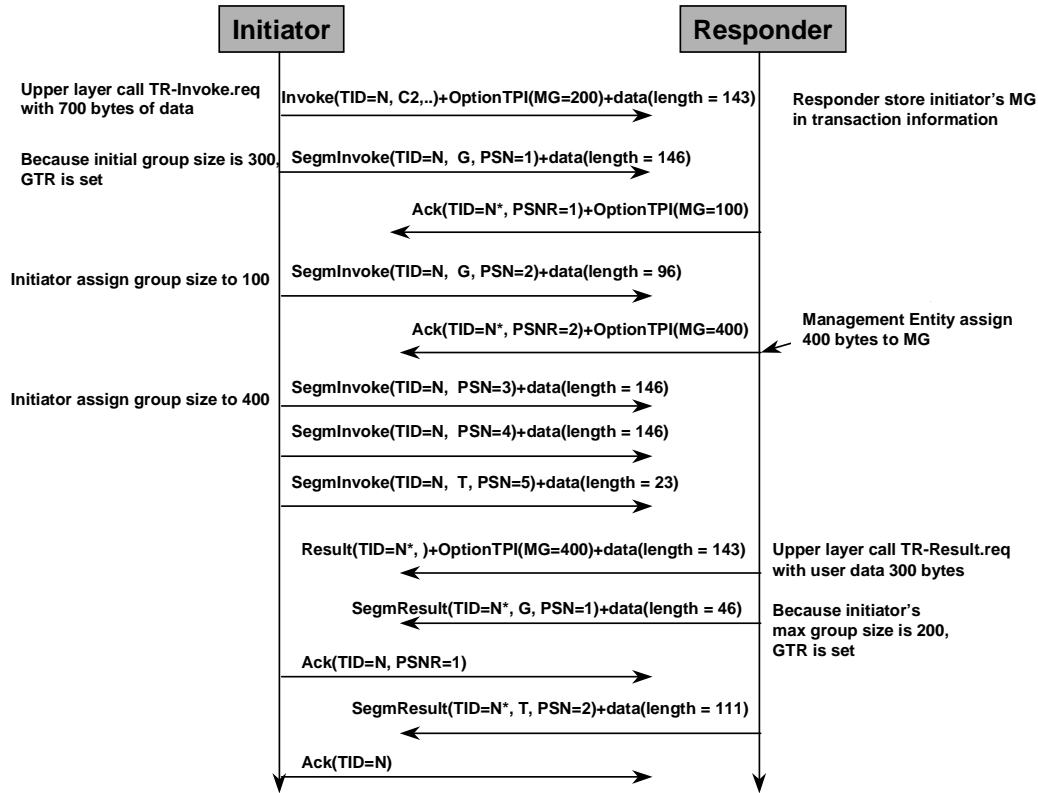


Figure 15 Option TPI(Maximum Group) usage

10.6.5 Basic Extended SAR

The example illustrates a class 2 transaction using extended SAR for the transmission of a large result. The example shows the beginning of the transaction. At the beginning of the transaction the negotiation of the extended SAR takes place by using the NumGroups Option TPI. Each packet group contains 3 packets. Packet 4 gets lost so the receiver sends a Nack containing this missing packet and the sender retransmits it.




Figure 16 Sample Flow in a Transaction using Extended SAR

10.6.6 Example of Re-transmission Hold-off

For example, assume a sender transmits 2 groups of 2 packets each, numbered 1, 2, 3, and 4 (packets 2 and 4 will have GTR set). If packet 1 is lost, the receiver will first see packet 2, with GTR set, and will start a Nack delay timer (in case packet 1 shows up out of order). When the Nack delay timer expires, the receiver will send a Nack for packet 1 (note that if packets 3 and/or 4 arrive prior to Nack delay timer expiration, the Nack sent in response to packet 2 should not include Nack information about packets 3 or 4). When the sender receives this Nack, it will immediately retransmit packet 1, and mark packet 1 as being in retransmission hold-off with a timestamp. When packets 3 and 4 arrive at the receiver, since the entire second group has been received and no reordering needs to be accounted for, it will immediately send a Nack for packet 1 again (an Ack would only be sent if there were no missing packets at all). When the sender receives this Nack, it will ignore the packet 1 information as long as packet 1 is still in retransmission hold-off. When packet 1 arrives at the receiver, it will send an Ack for packet 4.

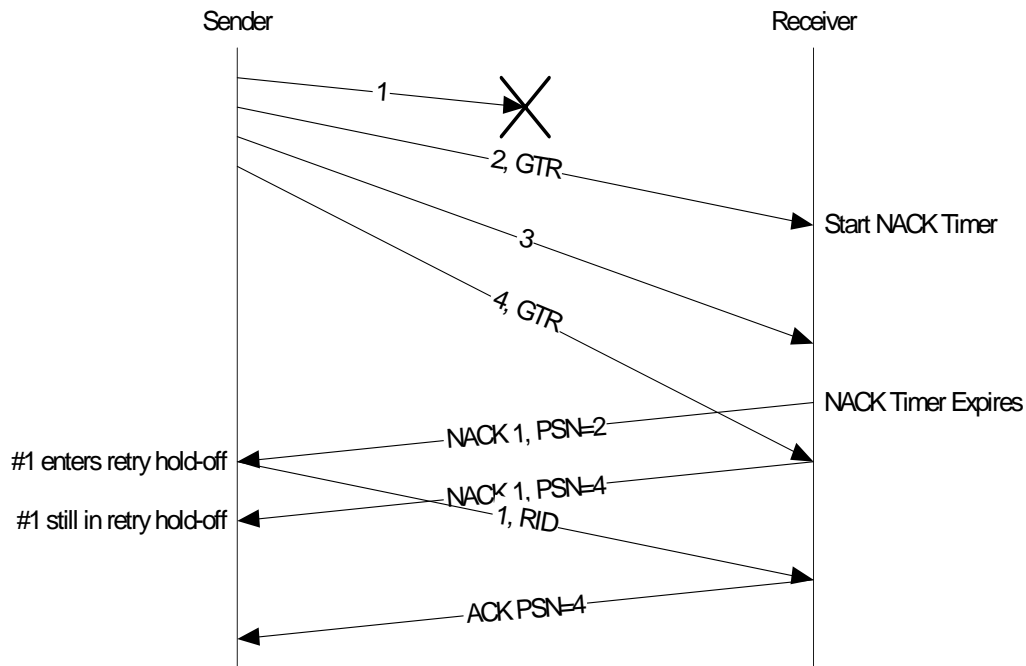


Figure 17 Elimination of Retransmission using Hold-off Timer

Note that if the retransmission hold-off period were too short, or not implemented, packet 1 would have been retransmitted twice, once more than necessary. Note also that if packets 3 and 4 had been delivered out of order, and the first Nack sent included information about packets 3 and 4, packet 3 would have been retransmitted unnecessarily as well. This is why a Nack generated for a group of packets should only include information about packets missing from that and previous groups.

10.6.7 Another Example of Re-transmission Hold-off

The following example demonstrates how the hold off period works in a class 2 transaction when the initiator is the sender.

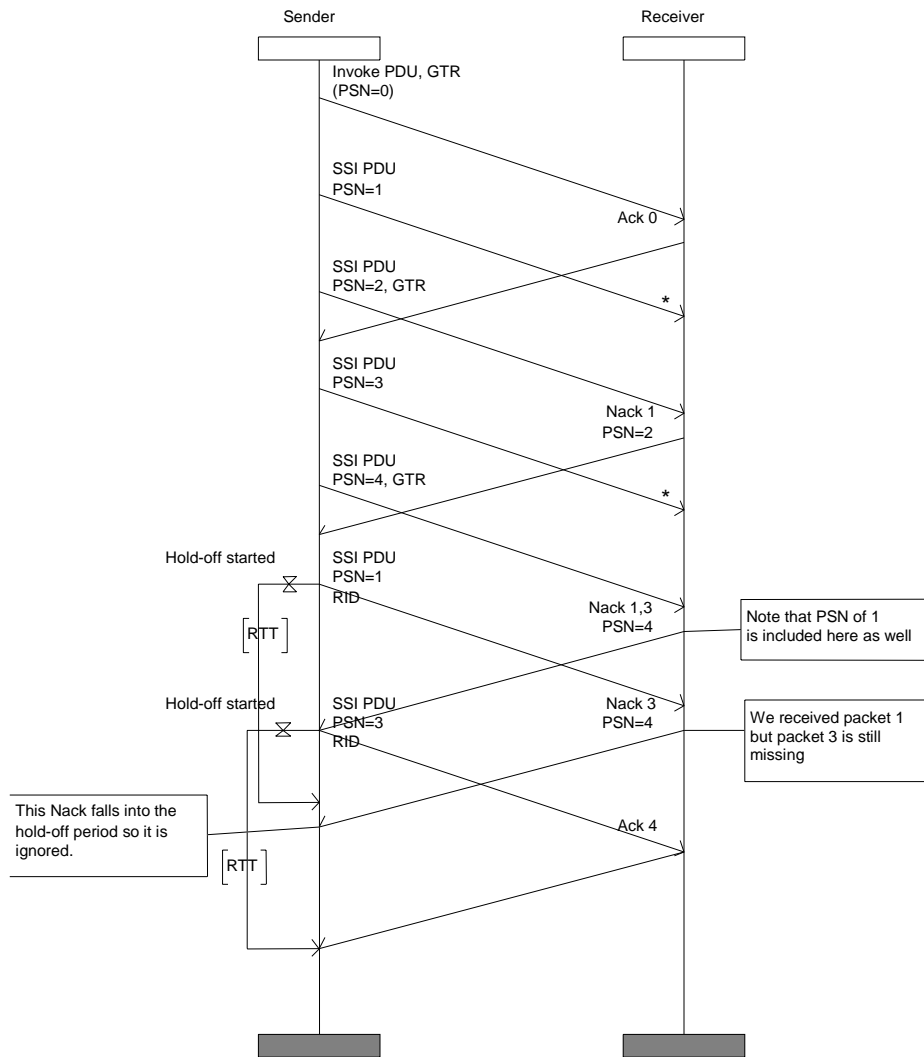


Figure 18 Filtering of Nacks during Hold-off Period

Packet 1 and 3 are lost. Due to the first Nack packet 1 is started to be retransmitted. At this point the hold-off period for approximately one RTT (see implementation above) is started. The Nack that lists packets 1 and 3 falls into the hold-off period of packet 1 so only packet 3 is retransmitted and thus a hold-off period for packet 3 is started. The last Nack falls into the hold-off period of packet 3 so no retransmission occurs due to that Nack. If the second Nack listing 1 and 3 had been lost the last Nack would have been very useful because it would have carried the information that packet 3 is missing.

Note that hold-off periods are started only for **retransmitted** packets. The first Nack for a packet must not be ignored by the sender.

10.6.8 SAR, NON-SAR Interactions

The following examples illustrate the interactions that can occur when a SAR capable initiator interacts with a non-SAR capable responder; in the case where SAR is supported but not used in the transaction.

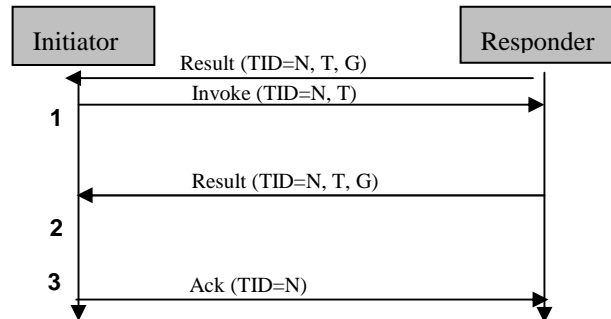


Figure 19: SAR Capable Initiator Communicating to Non SAR Capable Responder

1. The initiator indicates its support for SAR by using the settings for common transmission and group trailer flags (TTR and GTR) as described in table 12, section 8.2.2.
2. The responder indicates its lack of SAR support by using the appropriate settings for GTR and TTR.
3. The initiator acknowledges the received result message.

In the scenario in figure 19 it is also possible that the Responder may Abort the transaction using the reason code NOTIMPLEMENTEDSAR but, as pointed out in section 4.6 Note (1), the responder should consider whether it is necessary to abort. The implication in this notation is that only when a message uses SAR, the NON-SAR Responder should abort (e.g. [TTR=0 GTR=1] or [TTR=0 GTR=0]).

The following illustrates the scenario where a NON-SAR capable is communicating with a SAR capable gateway.

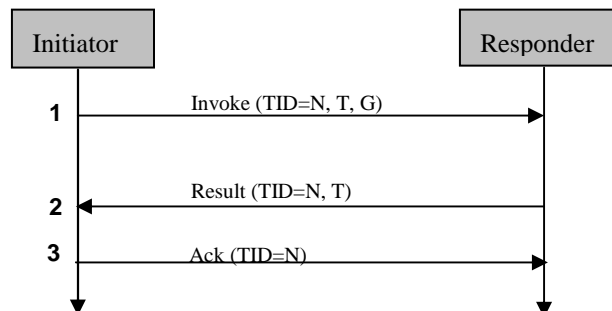


Figure 20: Non-SAR Capable Initiator Communicating to SAR Capable Responder

1. The initiator indicates its lack of support for SAR by using the settings for common transmission and group trailer flags (TTR and GTR) as described in table 12, section 8.2.2.

2. The responder indicates its support of SAR support by using the appropriate settings for GTR and TTR.
3. The initiator acknowledges the received result message.

Appendix A. Default Timer and Counter Values (Normative)

The timers are initial estimates and have not yet been verified.

The timer values in the tables below are expressed in seconds. The counters are expressed in times an event happens.

GSM SMS

The maximum round-trip time is assumed to be 40 seconds (while a median round-trip time is about 10 seconds), and the timer values are thus suggested to be:

Timer interval	Type	Without User Ack.	With User Ack.
Base Acknow. interval (A)	B_A	10	10
- Short	S_A	0	5
- Long	L_A	20	20
Base Retry interval (R)	B_R	60	60
- Short	S_R	35	40
- Long	L_R	70	70
- Group	G_R	45	45
Wait timeout interval (W)	W	300	300

Counter name	Value for stack acks	Value for user acks
Max Retransmissions	4	4
Max Ack timer Expiration	4	4

GSM USSD

The maximum round-trip time is assumed to be 5 seconds, and the timer values are thus suggested to be:

Timer interval	Type	Without User Ack.	With User Ack.
Base Acknow. Interval (A)	B_A	10	10
- Short	S_A	0	5
- Long	L_A	10	10
Base Retry Interval;	B_R	20	20
- Short	S_R	14	14
- Long	L_R	20	20
- Group	G_R	10	10
Wait timeout interval (W)	W	60	60

Counter name	Value for stack acks	Value for user acks
Max Retransmissions	4	4
Max Ack timer Expiration	4	4

Bearers supporting IP (Circuit switched, CDPD...)

The maximum round-trip time is assumed to be 3 seconds, and the timer values are thus suggested to be:

Timer interval	Type	Without User Ack.	With User Ack.
Base Acknow. interval (A)	B_A	2	2
- Short	S_A	0	1
- Long	L_A	4	4
Base Retry interval (R)	B_R	5	5
- Short	S_R	3	4
- Long	L_R	7	7
- Group	G_R	3	3
Wait timeout interval (W)	W	40	40

Counter name	Value for stack acks	Value for user acks
Max Retransmissions	8	8
Max Ack timer Expiration	6	6

Timer Usage

There are a number of timer interval with similar behaviour, but different values. These timers are defined to enable an optimal use of the available bandwidth. Chapter 10 refers to the abstract timer intervals A and R. These are mapped to the real interval values as defined in this section.

Message type	Class 2	Class 1
Invoke message	B_R	S_R
Hold on acknowledgement	B_A	-
Result message	L_R	-
Last acknowledgement	L_A	S_A
Last packet of packet group	G_R	G_R

For Class 0 no timer values are applicable.

Appendix B. Implementation Notes (Informative)

The following implementation notes are provided to identify areas where implementation choices may impact the performance and effectiveness of the WTP protocol. These notes provide guidance to implementers of the protocol.

B.1. Extended Timers for Large Messages

The Wireless Transaction Protocol is using Retransmission timers both to ensure reliable delivery of data to the receiver as well as to create a predictable behaviour in a lossy environment. The default timers are defined for relatively small transmissions. The protocol can manage also large data amounts, but then the timer values need to be adjusted. This implementation note suggests a scheme for recalculation of the timer intervals.

When large messages are transmitted from sender to receiver the transmission time for the complete message (or group) can become larger than the default value for the retransmission timer. The value thus has to be recalculated in order to avoid unnecessary retransmissions.

The new timer value to be used with both large datagrams as well as segmented transactions can be (implementation dependent) calculated according to the algorithm below:

$$rt = T + n * M$$

where

rt – recalculated retransmission timer value

T – Original timer value

n – estimated number of fragments (estimate made by protocol layer implementation)

M – bearer dependent value of ½ median roundtrip (estimate defined below)

The same algorithm is used both for large datagrams that will be fragmented at a lower protocol layer, such as in the WDP protocol, as well as for segmented messages where the WTP layer divides a group/message into multiple segments. The factor n is estimated in the former case and calculated in the latter case.

The following values for M can be used

Bearer	Median ½ roundtrip in seconds
SMS	5
USSD	3
IP	0.2

The value for IP represents a unit of 1 Kbyte. However, the timer value can be rounded to the nearest smaller integer.

B.2. Data Handling with Extended SAR

The WTP Invoke and Result primitives have a Frame Boundary parameter in order to provide applications with the ability to define their own data framing. This means that a single application data ‘frame’ may span across multiple primitives. Thus some WTP implementations may accept partial data frames from the user for transmission; similarly a WTP receiver may choose to buffer data till a complete frame is received or may pass it up to the user in smaller pieces.

B.3. Nack Generation and Interpretation

When multiple groups of packets are outstanding, more than one NACK may be received containing similar information. This could cause the sender to unnecessarily retransmit a packet. The retransmission hold-off period and NACK delay timer features are intended to mitigate this possibility.

B.3.1. Recommendations for Nack Generation

Whenever the receiver must generate a Nack PDU, a Nack delay timer should be started. The value of this timer should be either $\frac{1}{2}$ of the median round trip time, or twice the estimated round trip time mean deviation (RTTMDEV). A Nack packet should be generated for a group of packets when this Nack delay timer expires, or when the GTR packet for a following group of packets arrives. This will allow an implementation to use a single Nack delay timer, but requires the group size parameter to be large enough to account for the expected reordering. Note that if a Nack is sent before expiration of the Nack delay timer (due to reception of a GTR packet for a following group of packets), the Nack should not include information about packets that arrived later than the timer has been started unless there are no missing packets in the following group.

B.3.2. Recommendations for Nack Interpretation

Round trip times should be measured and used to drive a round trip time estimate (RTT_E), and round trip time mean deviation (RTT_MDEV). Each packet buffered by a sender for possible retransmission should have a retransmission hold-off timestamp (RHO_STAMP) associated with it.

When a Nack packet arrives, for each packet not acknowledged by the Nack the following procedure should be followed:

1. If the packet has not been retransmitted once already (i.e., RID clear), it should be retransmitted immediately with RID set, and the current time of day plus RTT_E plus 2 times RTT_MDEV should be recorded in the packet's retransmission hold-off timestamp, RHO_STAMP.
2. If the packet has been retransmitted once already (i.e., RID set), and the current time of day is earlier than the packet's RHO_STAMP, no action should be taken for this packet.
3. If the packet has been retransmitted once already (i.e., RID set), and the current time of day is later than the packet's RHO_STAMP, it should be retransmitted immediately, and the current time of day plus RTTE plus 2 times RTTMDEV should be recorded in the packet's RHO_STAMP.

Appendix C. Static Conformance Requirements (Normative)

The notation used in this appendix is specified in [IOPProc].

Item	Function	Reference	Status	Requirement
WTP-C-001	Transaction Class 0 Initiator	6.1.3	M	WDP:MCF
WTP-C-002	Transaction Class 0 Responder	6.1.3	M	WDP:MCF
WTP-C-003	Transaction Class 1 Initiator	6.2.4	M	WTP-C-022 AND WTP-C-015
WTP-C-004	Transaction Class 1 Responder	6.2.4	M	WTP-C-022
WTP-C-005	Transaction Class 2 Initiator	6.3.4	M	WTP-C-022 AND WTP-C-015
WTP-C-006	Transaction Class 2 Responder	6.3.4	O	WTP-C-022
WTP-C-007	User Acknowledgement	7.3	O	
WTP-C-008	Concatenation	4.1, 7.5	O	WDP:MCF
WTP-C-009	Separation	7.5	M	WDP:MCF
WTP-C-010	Retransmission until Acknowledgement	7.2	O	
WTP-C-011	Transaction Abort	4.6, 7.7, 7.12	O	WDP:MCF
WTP-C-012	Error Handling	7.12	O	WTP-C-011
WTP-C-013	Information in Last Acknowledgement	7.4, 7.10	O	WTP-C-018
WTP-C-014	Asynchronous Transactions	7.6	O	
WTP-C-015	Initiator response to TID Verification	7.1.5.2	O	WDP:MCF
WTP-C-016	Initiation of TID Verification by Responder	7.1.5.2, 7.8.1	O	WDP:MCF
WTP-C-017	Error Transport Information Item	7.10, 8.4.2	O	
WTP-C-018	Info Transport Information Item	7.10, 8.4.3	O	
WTP-C-019	Option Transport Information Item	7.10, 8.4.4	O	
WTP-C-020	PSN Transport Information Item	7.10, 8.4.5	O	
WTP-C-021	Segmentation and Re-assembly with Selective Retransmission and Packet Groups	7.14	O	WTP-C-020
WTP-C-022	Reliable transaction	7	O	WDP:MCF AND WTP-C-007 AND WTP-C-010 AND WTP-C-011 AND WTP-C-012
WTP-C-023	Extended Segmentation and Re-assembly with Selective Retransmission and Packet Groups	7.15	O	WTP-C-020 AND WTP-C-019 AND WTP-C-024 AND WTP-C-025
WTP-C-024	Frame Boundary Transport Information Item	7.15.2, 8.4.7	O	
WTP-C-025	SDU Boundary Transport Information Item	7.15.2, 8.4.6	O	
WTP-C-026	Support sliding window with ESAR	7.15.3, 7.15.4	O	

Item	Function	Reference	Status	Requirement
WTP-S-001	Transaction Class 0 Initiator	6.1.3	M	WDP:MSF
WTP-S-002	Transaction Class 0 Responder	6.1.3	M	WDP:MSF
WTP-S-003	Transaction Class 1 Initiator	6.2.4	M	WTP-S-022 AND WTP-S-015
WTP-S-004	Transaction Class 1 Responder	6.2.4	M	WTP-S-022
WTP-S-005	Transaction Class 2 Initiator	6.3.4	O	WTP-S-022 AND WTP-S-015
WTP-S-006	Transaction Class 2 Responder	6.3.4	M	WTP-S-022
WTP-S-007	User Acknowledgement	7.3	O	
WTP-S-008	Concatenation	4.1, 7.5	O	WDP:MSF
WTP-S-009	Separation	7.5	M	WDP:MSF
WTP-S-010	Retransmission until Acknowledgement	7.2	O	
WTP-S-011	Transaction Abort	4.6, 7.7, 7.12	O	WDP:MSF
WTP-S-012	Error Handling	7.12	O	WTP-S-011
WTP-S-013	Information in Last Acknowledgement	7.4, 7.10	O	WTP-S-018
WTP-S-014	Asynchronous Transactions	7.6	O	
WTP-S-015	Initiator response to TID Verification	7.1.5.2	O	WDP:MSF
WTP-S-016	Initiation of TID Verification by Responder	7.1.5.2, 7.8.1	O	WDP:MSF
WTP-S-017	Error Transport Information Item	7.10, 8.4.2	O	
WTP-S-018	Info Transport Information Item	7.10, 8.4.3	O	
WTP-S-019	Option Transport Information Item	7.10, 8.4.4	O	
WTP-S-020	PSN Transport Information Item	7.10, 8.4.5	O	
WTP-S-021	Segmentation and Re-assembly with Selective Retransmission and Packet Groups	7.14	O	WTP-S-020
WTP-S-022	Reliable transaction	7	O	WDP:MSF AND WTP-S-007 AND WTP-S-010 AND WTP-S-011 AND WTP-S-012
WTP-S-023	Extended Segmentation and Re-assembly with Selective Retransmission and Packet Groups	7.15	O	WTP-S-020 AND WTP-S-019 AND WTP-S-024 AND WTP-S-025
WTP-S-024	Frame Boundary Transport Information Item	7.15.2, 8.4.7	O	
WTP-S-025	SDU Boundary Transport Information Item	7.15.2, 8.4.6	O	
WTP-S-026	Support sliding window with ESAR	7.15.3, 7.15.4	O	

Appendix D. Change History

(Informative)

D.1 Approved Version History

Reference	Date	Description
WAP-201-WTP-20000219-a	19-Feb-2000	The used baseline specification
WAP-201_001-WTP-20001212-a	12-Dec-2000	10.6 a Small corrections to TID handling in responder state tables.
WAP-201_002-WTP-20001213-a	13-Dec-2000	2.1, 4.6, App A New SCR table format. Class 0 08-Feb-2001 Mechanisms for large data transfer
WAP-224-WTP-20010208-p	08-Feb-2001	Published as proposed version
WAP-224_001-WTP-20010710-p	10-Jul-2001	2.1, 7.15.3, 7.15.4, 8.4.4, 9.4.1, App A, App C Update CREQ reference. Clarify meaning of window closing and exponential back-off. Define default for Maximum Group option. Permit use of long Option TPIs. Fix garbled sentence. Expose sliding window in SCR tables.
OMA-WAP-WTP--20031612-a	16 Dec 2003	Sections 8.2.2, 10.6.8; WAP-224_002-WTP-20020827-a; clarifications on WTP-SAR ; update CREQ reference to IOP Process. Updated to use OMA template
WAP-224-WTP-20020827-a	10-Jan-2005	Revert document numbering to WAP Forum format. Correct date to reflect WAP-224_002-20020827-a which was inadvertently wrongly cited in the entry of 16 Dec 2003. Update ref for WDP