

Expires: August 2003

Marker PDU Aligned Framing for TCP Specification

1 Status of this Memo

This document is an Internet-Draft and is subject to all provisions of Section 10 of RFC2026.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/lid-abstracts.html>. The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

2 Abstract

A framing protocol is defined for TCP that is fully compliant with applicable TCP RFCs and fully interoperable with existing TCP implementations. The framing mechanism is designed to work as an "adaptation layer" between TCP and the Direct Data Placement [DDP] protocol, preserving the reliable, in-order delivery of TCP, while adding the preservation of higher-level protocol record boundaries that DDP requires.

Table of Contents

1	Status of this Memo.....	1
2	Abstract.....	1
3	Introduction.....	4
3.1	Motivation.....	4
3.2	Protocol Overview.....	5
4	Glossary.....	7
5	LLP and DDP requirements.....	8
5.1	TCP implementation Requirements to support MPA.....	8
5.1.1	TCP Transmit side.....	8
5.1.2	TCP Receive side.....	8
5.2	MPA's interactions with DDP.....	9
6	FPDU Formats.....	11
6.1	Marker Format.....	13
7	Data Transfer Semantics.....	14
7.1	MPA Markers.....	14
7.2	CRC Calculation.....	16
7.3	MPA on TCP Sender Segmentation.....	19
7.3.1	Effects of MPA on TCP Segmentation.....	19
7.3.2	FPDU Size Considerations.....	20
7.4	MPA Receiver FPDU Identification.....	21
7.4.1	Re-segmenting Middle boxes and non MPA-aware TCP senders....	22
8	Connection Semantics.....	24
8.1	Connection setup.....	24
8.2	Normal Connection Teardown.....	25
9	Error Semantics.....	26
10	Security Considerations.....	27
10.1	Protocol-specific Security Considerations.....	27
10.2	Using IPsec With MPA.....	27
11	IANA Considerations.....	28
12	References.....	29
12.1	Normative References.....	29
12.2	Informative References.....	29
13	Appendix.....	31
13.1	Receiver implementation.....	31
13.1.1	Transport & Network Layer Reassembly Buffers	31
14	Author's Addresses.....	33
15	Acknowledgments.....	34
16	Full Copyright Statement.....	37

Table of Figures

Figure 1 ULP MPA TCP Layering.....	6
Figure 2 FPDU Format.....	11
Figure 3 Marker Format.....	13
Figure 4 Example FPDU Format with Marker.....	15
Figure 5 Annotated Hex Dump of an FPDU.....	18
Figure 6 Annotated Hex Dump of an FPDU with Marker.....	18
Figure 7: Example Startup negotiation.....	25

Revision history

[02] Enhanced descriptions of how MPA is used over an unmodified TCP.

[02] Removed "No Packing" text.

[02] Made MPA an adaptation layer for DDP, instead of a generalized framing solution.

[02] Added clarifications of the MPA/TCP interaction for optimized implementations and that any such optimizations are to be used only when requested by MPA.

Note: a discussion of reasons for these changes can be found in [ELZER-MPA].

3 Introduction

This section discusses the reason for creating MPA on TCP and a general overview of the protocol. Later sections show the MPA headers (see section 6 on page 11), and detailed protocol requirements and characteristics (see section 7 on page 14), as well as Connection Semantics (section 8 on page 22), Error Semantics (section 9 on page 26), and Security Considerations (section 10 on page 27).

3.1 Motivation

The Direct Data Placement protocol [DDP], when used with TCP [RFC793] requires a mechanism to detect record boundaries. The DDP records are referred to as Upper Layer Protocol Data Units by this document. The ability to locate the Upper Layer Protocol Data Unit (ULPDU) boundary is useful to a hardware network adapter that uses DDP to directly place the data in the application buffer based on the control information carried in the ULPDU header. This may be done without requiring that the packets arrive in order. Potential benefits of this capability are the avoidance of the memory copy overhead and a smaller memory requirement for handling out of order or dropped packets.

Many approaches have been proposed for a generalized framing mechanism. Some are probabilistic in nature and others are deterministic. A probabilistic approach is characterized by a detectable value embedded in the octet stream. It is probabilistic because under some conditions the receiver may incorrectly interpret application data as the detectable value. Under these conditions, the protocol may fail with unacceptable frequency. A deterministic approach is characterized by embedded controls at known locations in the octet stream. Because the receiver can guarantee it will only examine the data stream at locations that are known to contain the embedded control, the protocol can never misinterpret application data as being embedded control data. For unambiguous handling of an out of order packet, the deterministic approach is preferred.

The MPA protocol provides a framing mechanism for DDP running over TCP using the deterministic approach. It allows the location of the ULPDU to be determined in the TCP stream even if the TCP segments arrive out of order.

3.2 Protocol Overview

MPA is described as an extra layer above TCP and below DDP. The end-to-end data flow is:

1. The DDP's ULP negotiates the use of DDP and MPA at both ends of a connection.
2. DDP determines the Maximum ULDPDU (MULPDU) size by querying MPA for this value. MPA derives this information from TCP, when it is available, or chooses a reasonable value. This information is already supported on many TCP implementations, including all modern flavors of BSD networking, through the TCP_MAXSEG socket option.
3. DDP creates ULPDUs of MULPDU size or smaller, and hands them to MPA at the sender.
4. MPA creates a Framed Protocol Data Unit (FPDU) by pre-pending a header, inserting markers, and appending a CRC after the ULPDU and PAD (if any). MPA delivers the FPDU to TCP.
5. The TCP sender puts the FPDUs into the TCP stream. If the TCP Sender is MPA-aware, it segments the TCP stream in such a way that a TCP Segment boundary is also the boundary of an FPDU. TCP then passes each segment to the IP layer for transmission.
6. The TCP receiver may be MPA-aware or may not be MPA-aware. If it is MPA-aware, it may separate passing the TCP payload to MPA from passing the TCP payload ordering information to MPA. In either case, RFC compliant TCP wire behavior is observed at both the sender and receiver.
7. The MPA receiver locates and assembles complete FPDUs within the stream, verifies their integrity, and removes MPA markers, ULPDU_Length, PAD and CRC.
8. MPA then provides the complete ULPDUs to DDP. MPA may also separate passing MPA payload to DDP from passing the MPA payload ordering information.

The layering of PDUs with MPA is shown in Figure 1, below.

MPA-aware TCP is a TCP layer which potentially contains some additional semantics as defined in this document. MPA is implemented as a data stream ULP for TCP and is therefore RFC compliant. MPA-aware TCP is RFC compliant.

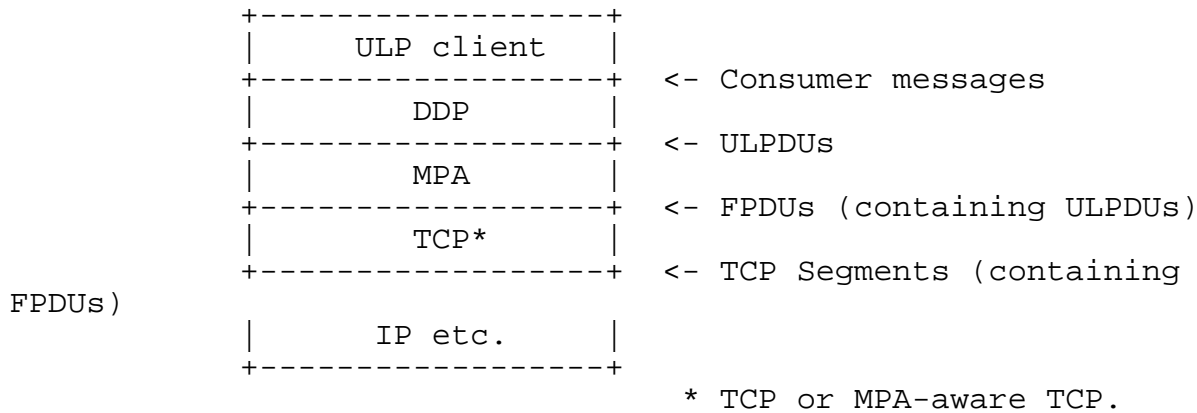


Figure 1 ULP MPA TCP Layering

An MPA-aware TCP sender is able to segment the data stream such that TCP segments begin with FPDUs (FPDU Alignment). This has significant advantages for receivers. When segments arrive with aligned FPDUs the receiver usually need not buffer any portion of the segment, allowing DDP to place it in its destination memory immediately, thus avoiding copies from intermediate buffers (DDP's reason for existence).

MPA with an MPA-aware TCP receiver allows a DDP on MPA implementation to recover ULPDUs that may be received out of order. This enables a DDP on MPA implementation to save a significant amount of intermediate storage by placing the ULPDUs in the right locations in the application buffers when they arrive, rather than waiting until full ordering can be restored.

MPA implementations that support recovery of out of order ULPDUs MUST support a mechanism to indicate the ordering of ULPDUs as the sender transmitted them and indicate when missing intermediate segments arrive. These mechanisms allow DDP to reestablish record ordering and report Delivery of complete messages (groups of records).

MPA also addresses enhanced data integrity. Many users of TCP have noted that the TCP checksum is not as strong as could be desired [CRCTCP]. Studies have shown that the TCP checksum indicates segments in error at a much higher rate than the underlying link characteristics would indicate. With these higher error rates, the chance that an error will escape detection, when using only the TCP checksum for data integrity, becomes a concern. A stronger integrity check can reduce the chance of data errors being missed.

MPA includes a CRC check to increase the ULPDU data integrity to the level provided by other modern protocols, such as SCTP [RFC2960].

4 Glossary

Delivery - (Delivered, Delivers) - For MPA, Delivery is defined as the process of informing DDP that a particular PDU is ordered for use. This is specifically different from "passing the PDU to DDP", which may generally occur in any order, while the order of "Delivery" is strictly defined.

EMSS - Effective Maximum Segment Size. EMSS is the smaller of the TCP maximum segment size (MSS) as defined in RFC 793 [RFC793], and the current path Maximum Transfer Unit (MTU) [RFC1191].

FPDU - Framing Protocol Data Unit. The unit of data created by an MPA sender.

FPDU Alignment - the property that a TCP segment begins with an FPDU.

PDU - protocol data unit

MPA - Marker-based ULP PDU Aligned Framing for TCP protocol. This document defines the MPA protocol.

MULPDU - Maximum ULPDU. The current maximum size of the record that is acceptable for DDP to pass to MPA for transmission.

Node - A computing device attached to one or more links of a Network. A Node in this context does not refer to a specific application or protocol instantiation running on the computer. A Node may consist of one or more MPA on TCP devices installed in a host computer.

Remote Peer - The MPA protocol implementation on the opposite end of the connection. Used to refer to the remote entity when describing protocol exchanges or other interactions between two Nodes.

ULP - Upper Layer Protocol. The protocol layer above the protocol layer currently being referenced. The ULP for MPA is DDP [DDP].

ULPDU - Upper Layer Protocol Data Unit. The data record defined by the layer above MPA (DDP). ULPDU corresponds to DDP's "DDP Segment".

5 LLP and DDP requirements

5.1 TCP implementation Requirements to support MPA

The TCP implementation MUST inform MPA when the TCP connection is closed or has begun closing the connection (e.g. received a FIN).

5.1.1 TCP Transmit side

To provide optimum performance, an MPA-aware transmit side TCP implementation SHOULD be enabled to:

- * With an EMSS large enough to contain the FPDU(s), segment the outgoing TCP stream such that the first octet of every TCP Segment begins with an FPDU. Multiple FPDUs MAY be packed into a single TCP segment as long as they are entirely contained in the TCP segment.
- * Report the current EMSS to the MPA transmit layer.

An MPA-aware TCP transmit side implementation MUST continue to use the method of segmentation expected by non-MPA applications (and described in TCP RFCs) when MPA is not enabled on the connection. When MPA is enabled above an MPA-aware TCP, it SHOULD specifically enable the segmentation rules described above for the DDP segments (FPDUs) posted for transmission.

If the transmit side TCP implementation is not able to segment the TCP stream as indicated above, MPA should make a best effort to achieve that result. For example, using the TCP_NODELAY socket option to disable the Nagle algorithm will usually result in many of the segments starting with an FPDU.

If the transmit side TCP implementation is not able to report the EMSS, MPA may assume that TCP will use 1460 octet segments in creating FPDUs. If the implementation has reason to believe that the TCP segment size is actually smaller than 1460, it may instead use a 536 octet FPDU.

5.1.2 TCP Receive side

When an MPA receive implementation and the MPA-aware receive side TCP implementation supports handling out of order ULPDUs, the TCP receive implementation SHOULD be enabled to:

- * Pass incoming TCP segments to MPA as soon as they have been received and validated, even if not received in order. The TCP layer MUST have committed to keeping each segment before it can be passed to the MPA. This means that the segment must have passed the TCP, IP, and lower layer data integrity validation (i.e., checksum), must be in the receive window, must not be a duplicate, must be part of the same epoch (if timestamps are

used to verify this) and any other checks required by TCP RFCs. The segment MUST NOT be passed to MPA more than once unless explicitly requested (see Section 9).

This is not to imply that the data must be completely ordered before use. An implementation may accept out of order segments, SACK them [RFC2018], and pass them to DDP when the reception of the segments needed to fill in the gaps arrive. Such an implementation can "commit" to the data early on, and will not overwrite it even if (or when) duplicate data arrives. MPA expects to utilize this "commit" to allow the passing of ULPDUs to DDP when they arrive, independent of ordering.

- * Provide a mechanism to indicate the ordering of TCP segments as the sender transmitted them. One possible mechanism might be attaching the TCP sequence number to each segment.
- * Provide a mechanism to indicate when a given TCP segment (and the prior TCP stream) is complete. One possible mechanism might be to utilize the leading (left) edge of the TCP Receive Window.

DDP on MPA MUST utilize these two mechanisms to establish the Delivery semantics that DDP's consumers agree to. These semantics are described fully in [DDP]. These include requirements on DDP's consumer to respect ownership of buffers prior to the time that DDP delivers them to the consumer.

An MPA-aware TCP receive side implementation MUST continue to buffer TCP segments until completely ordered and then deliver them as expected by non-MPA applications (and described in TCP RFCs) when MPA is not enabled on the connection. When MPA is enabled above an MPA-aware TCP, TCP SHOULD enable the in and out of order passing of data, and the separate ordering information as described above.

When an MPA receive implementation is coupled with a TCP receive implementation that does not support the preceding mechanisms, TCP passes and Delivers incoming stream data to MPA in order.

5.2 MPA's interactions with DDP

DDP requires MPA to maintain DDP record boundaries from the sender to the receiver. When using MPA on TCP to send data, DDP provides records (ULPDUs) to MPA. MPA will use the reliable transmission abilities of TCP to transmit the data, and will insert appropriate additional information into the TCP stream to allow the MPA receiver to locate the record boundary information.

As such, MPA accepts complete records (ULPDUs) from DDP at the sender and returns them to DDP at the receiver.

MPA combined with an MPA-aware TCP can only ensure FPDU Alignment with the TCP Header if the FPDU is less than or equal to TCP's EMSS. Since FPDU alignment is generally desired by the receiver, DDP must cooperate with MPA to ensure FPDUs' lengths do not exceed the EMSS under normal conditions. This is done with the MULPDU mechanism.

MPA provides information to DDP on the current maximum size of the record that is acceptable to send (MULPDU). DDP SHOULD limit each record size to MULPDU. The range of MULPDU values MUST be between 128 octets and 64768 octets, inclusive.

The sending DDP MUST NOT post a ULPDU larger than 64768 octets to MPA. DDP MAY post a ULPDU of any size between one and 64768 octets, however MPA is NOT REQUIRED to support a ULPDU length that is greater than the current MULPDU.

While the maximum theoretical length supported by the MPA header ULPDU_Length field is 65535, TCP over IP requires the IP datagram maximum length to be 65535 octets. To enable MPA to support FPDU Alignment, the maximum size of the FPDU must fit within an IP datagram. Thus the ULPDU limit of 64768 octets was derived by taking the maximum IP datagram length, subtracting from it the maximum total length of the sum of the IPv4 header, TCP header, IPv4 options, TCP options, and the worst case MPA overhead, and then rounding the result down to a 128 octet boundary.

On receive, MPA MUST pass each ULPDU with its length to DDP when it has been validated.

If an MPA implementation supports passing out of order ULPDUs to DDP, the MPA implementation SHOULD:

- * Pass each ULPDU with its length to DDP as soon as it has been fully received and validated.
- * Provide a mechanism to indicate the ordering of ULPDUs as the sender transmitted them. One possible mechanism might be providing the TCP sequence number for each ULPDU.
- * Provide a mechanism to indicate when a given ULPDU (and prior ULPDUs) are complete. One possible mechanism might be to allow DDP to see the current outgoing TCP Ack sequence number.
- * Provide an indication to DDP that the TCP has closed or has begun to close the connection (e.g. received a FIN).

6 FPDU Formats

MPA senders create FPDUs out of ULPDUs. The format of an FPDU shown below MUST be used for all MPA FPDUs. For purposes of clarity, markers are not shown in Figure 2.




Figure 2 FPDU Format

ULPDU_Length: 16 bits (unsigned integer). This is the number of octets of the contained ULPDU. It does not include the length of the FPDU header itself, the pad, the CRC, or of any markers that fall within the ULPDU. The 16-bit ULPDU Length field is large enough to support the largest IP datagrams for IPv4 or IPv6.

PAD: The PAD field trails the ULPDU and contains between zero and three octets of data. The pad data MUST be set to zero by the sender and ignored by the receiver (except for CRC checking). The length of the pad is set so as to make the size of the FPDU an integral multiple of four.

CRC: 32 bits, this CRC is used to verify the entire contents of the FPDU, using CRC32C See section 7.2 CRC Calculation on page 16.

The FPDU adds a minimum of 6 octets to the length of the ULPDU. In addition, the total length of the FPDU will include the length of

any markers and from 0 to 3 pad octets added to round-up the ULDPDU size.

6.1 Marker Format

The format of a marker MUST be as specified in Figure 3:




Figure 3 Marker Format

RESERVED: The Reserved field MUST be set to zero on transmit and ignored on receive (except for CRC calculation).

FPDUPTR: The FPDU Pointer is a relative pointer, 16-bits long, interpreted as an unsigned integer, that indicates the number of octets in the TCP stream from the beginning of the FPDU to the first octet of the entire marker.

7 Data Transfer Semantics

This section discusses some characteristics and behavior of the MPA protocol as well as implications of that protocol.

7.1 MPA Markers

MPA senders MUST insert a marker into the data stream at a 512 octet periodic interval in the TCP Sequence Number Space. The marker contains a 16 bit unsigned integer referred to as the FPDUPTR (FPDU Pointer).

If the FPDUPTR's value is non-zero, the FPDU Pointer is a 16 bit relative back-pointer. FPDUPTR MUST contain the number of octets in the TCP stream from the beginning of the current FPDU to the first octet of the marker, unless the marker falls between FPDUs. Thus the location of the first octet of the previous FPDU header can be determined by subtracting the value of the given marker from the current octet-stream sequence number (i.e. TCP sequence number) of the first octet of the marker. Note that this computation must take into account that the TCP sequence number could have wrapped between the marker and the header.

An FPDUPTR value of 0x0000 is a special case - it is used when the marker falls exactly between FPDUs. In this case, the marker MUST be placed in the following FPDU and viewed as being part of that FPDU (e.g. for CRC calculation). Thus an FPDUPTR value of 0x0000 means that immediately following the marker is an FPDU header.

Since all FPDUs are integral multiples of 4 octets, the bottom two bits of the FPDUPTR as calculated by the sender are zero. MPA reserves these bits so they MUST be treated as zero for computation at the receiver.

The MPA markers MUST be inserted immediately following MPA connection establishment, and at every 512th octet of the TCP octet stream thereafter. As a result, the first marker has an FPDUPTR value of 0x0000. If the first marker begins at octet sequence number SeqStart, then markers are inserted such that the first octet of the marker is at octet sequence number SeqNum if the remainder of $(SeqNum - SeqStart) \bmod 512$ is zero. Note that SeqNum can wrap.

For example, if the TCP sequence number were used to calculate the insertion point of the marker, the starting TCP sequence number is unlikely to be zero, and 512 octet multiples are unlikely to fall on a modulo 512 of zero. If the MPA connection is started at TCP sequence number 11, then the 1st marker will begin at 11, and subsequent markers will begin at 523, 1035, etc.

7.2 CRC Calculation

When sending an FPDU, the sender MUST include a valid CRC field. The CRC field in the MPA FPDU MUST be computed using the CRC32C polynomial in the manner described in the iSCSI Protocol [iSCSI] document for Header and Data Digests.

The fields which MUST be included in the CRC calculation when sending an FPDU are as follows:

- 1) If the first octet of the FPDU is the "ULPDU Length" field, the CRC-32c is calculated from the first octet of the "ULPDU Length" header, through all the ULPDU and markers (if present), to the last octet of the PAD (if present), inclusive. If there is a marker immediately following the PAD, the marker is included in the CRC calculation for this FPDU.
- 2) If the first octet of the FPDU is a marker, (i.e. the marker fell between FPDUs, and thus is required to be included in the second FPDU), the CRC-32c is calculated from the first octet of the marker, through the "ULPDU Length" header, through all the ULPDU and markers (if present), to the last octet of the PAD (if present), inclusive.
- 3) After calculating the CRC-32c, the resultant value is placed into the CRC field at the end of the FPDU.

When an FPDU is received, the receiver MUST first perform the following:

- 1) Calculate the CRC of the incoming FPDU in the same fashion as defined above.
- 2) Verify that the calculated CRC-32c value is the same as the received CRC-32c value found in the FPDU CRC field. If not, the receiver MUST treat the FPDU as an invalid FPDU.

The procedure for handling invalid FPDUs is covered in the Error Section (see section 9 on page 26)

The following is an annotated hex dump of an example FPDU sent as the first FPDU on the stream. As such, it starts with a marker. The FPDU contains 24 octets of the contained ULPDU, which are all zeros. The CRC32c has been correctly calculated and can be used as a reference. See the [DDP] and [RDMA] specification for definitions of the DDP Control field, Queue, MSN, MO, and Send Data.

Octet Count	Contents	Annotation
0000	00 00	Marker: Reserved
0002	00 00	FPDUPTR
0004	00 2a	Length
0006	40 03	DDP Control Field, Send with Last flag set
0008	00 00	Reserved (STag position with no STag)
000a	00 00	
000c	00 00	Queue = 0
000e	00 00	
0010	00 00	MSN = 1
0012	00 01	
0014	00 00	MO = 0
0016	00 00	
0018	00 00	
		Send Data (24 octets of zeros)
002e	00 00	
0030	4C 86	CRC32c
0032	B3 84	

Figure 5 Annotated Hex Dump of an FPDU

The following is an example sent as the second FPDU of the stream where the first FPDU (which is not shown here) had a length of 492 octets and was also a Send to Queue 0 with Last Flag set. This example contains a marker.

Octet Count	Contents	Annotation
01ec	00 2a	Length
01ee	40 03	DDP Control Field: Send with Last Flag set
01f0	00 00	Reserved (STag position with no STag)
01f2	00 00	
01f4	00 00	Queue = 0
01f6	00 00	
01f8	00 00	MSN = 2
01fa	00 02	
01fc	00 00	MO = 0
01fe	00 00	
0200	00 00	Marker: Reserved
0202	00 14	FPDUPTR
0204	00 00	
		Send Data (24 octets of zeros)
021a	00 00	
021c	A1 9C	CRC32c
021e	D1 03	

Figure 6 Annotated Hex Dump of an FPDU with Marker

7.3 MPA on TCP Sender Segmentation

The various TCP RFCs allow considerable choice in segmenting a TCP stream. In order to optimize FPDU recovery at the MPA receiver, MPA specifies additional segmentation rules.

MPA MUST encapsulate the ULDPDU such that there is exactly one ULDPDU contained in one FPDU.

An MPA-aware TCP sender SHOULD, when enabled for MPA, on TCP implementations that support this, and with an EMSS large enough to contain at least one FPDU, segment the outbound TCP stream such that each TCP segment begins with an FPDU, and fully contains all included FPDUs.

Implementation note: To achieve the previous segmentation rule, TCP's Nagle [RFC0896] algorithm SHOULD be disabled.

There are exceptions to the above rule. Once an ULDPDU is provided to MPA, the MPA on TCP sender MUST transmit it or fail the connection; it cannot be repudiated. As a result, during changes in MTU and EMSS, or when TCP's Receive Window size (RWIN) becomes too small, it may be necessary to send FPDUs that do not conform to the segmentation rule above.

A possible, but less desirable, alternative is to use IP fragmentation on accepted FPDUs to deal with MTU reductions or extremely small EMSS.

The sender MUST still format the FPDU according to FPDU format as shown in Figure 2.

On a retransmission, TCP does not necessarily preserve original TCP segmentation boundaries. This can lead to the loss of FPDU alignment and containment within a TCP segment during TCP retransmissions. An MPA-aware TCP sender SHOULD try to preserve original TCP segmentation boundaries on a retransmission.

7.3.1 Effects of MPA on TCP Segmentation

Applications expected to see strong advantages from Direct Data Placement include transaction-based applications and throughput applications. Request/response protocols typically send one FPDU per TCP segment and then wait for a response. Therefore, the application is expected to set TCP parameters such that it can trade off latency and wire efficiency. This is accomplished by setting the TCP_NODELAY socket option.

When latency is not critical, and the application provides data in chunks larger than EMSS at one time, the TCP implementation may "pack" any available stream data into TCP segments so that the segments are filled to the EMSS. If the amount of data available

is not enough to fill the TCP segment when it is prepared for transmission, TCP can send the segment partly filled, or use the Nagle algorithm to wait for the ULP to post more data (discussed below).

DDP/MPA senders will fill TCP segments to the EMSS with a single FPDU when a DDP message is large enough. Since the DDP message may not exactly fit into TCP segments, a "message tail" often occurs that results in an FPDU that is smaller than a single TCP segment. If a "message tail", small DDP messages, or the start of a larger DDP message are available, MPA MAY "pack" the resulting FPDUs into TCP segments. When this is done, the TCP segments can be more fully utilized, but, due to the size constraints of FPDUs, segments may not be filled to the EMSS.

Note that MPA receivers must do more processing of a TCP segment that contains multiple FPDUs, this may affect the performance of some receiver implementations.

TCP implementations often utilize the "Nagle" [RFC0896] algorithm to ensure that segments are filled to the EMSS whenever the round trip latency is large enough that the source stream can fully fill segments before Acks arrive. The algorithm does this by delaying the transmission of TCP segments until a ULP can fill a segment, or until an ACK arrives from the far side. The algorithm thus allows for smaller segments when latencies are shorter to keep the ULP's end to end latency to reasonable levels.

The Nagle algorithm is not mandatory to use [RFC1122].

It is up to the ULP to decide if Nagle is useful with DDP/MPA. Note that many of the applications expected to take advantage of MPA/DDP prefer to avoid the extra delays caused by Nagle. In such scenarios it is anticipated there will be minimal opportunity for packing at the transmitter and receivers may choose to optimize their performance for this anticipated behavior.

7.3.2 FPDU Size Considerations

MPA defines the Maximum Upper Layer Protocol Data Unit (MULPDU) as the size of the largest ULPDU fitting in an FPDU. For an empty TCP Segment, MULPDU is EMSS minus the FPDU overhead (6 octets) minus space for markers and pad octets.

The maximum ULPDU Length for a single ULPDU MUST be computed as:

$$\text{MULPDU} = \text{EMSS} - (6 + 4 * \text{Ceiling}(\text{EMSS} / 512) + \text{EMSS} \bmod 4)$$

The formula above accounts for the worst-case number of markers.

As a further optimization of the wire efficiency an MPA implementation MAY dynamically adjust the MULPDU (see section

7.3.1. for latency and wire efficiency trade-offs). When one or more FPDUs are already packed into a TCP Segment, MULPDU MAY be reduced accordingly.

DDP SHOULD provide ULPDUs that are as large as possible, but less than or equal to MULPDU.

If the TCP implementation needs to adjust EMSS to support MTU changes, the MULPDU value is changed accordingly.

In certain rare situations, the EMSS may shrink to very small sizes. If this occurs, the MPA on TCP sender MUST NOT shrink the MULPDU below 128 octets and is not required to follow the segmentation rules in Section 7.3 MPA on TCP Sender Segmentation on page 19.

If one or more FPDUs are already packed into a TCP segment, such that the remaining room is less than 128 octets, MPA MUST NOT provide a MULPDU smaller than 128. In this case, MPA would typically provide a MULPDU for the next full sized segment, but may still pack the next FPDU into the small remaining room, provide that the next FPDU is small enough to fit.

The value 128 is chosen as to allow DDP designers room for the DDP Header and some user data.

7.4 MPA Receiver FPDU Identification

An MPA receiver MUST first verify the FPDU before passing the ULPDU to DDP. To do this, the receiver MUST:

- * locate the start of the FPDU unambiguously,
- * verify its CRC.

If the above conditions are true, the MPA receiver passes the ULPDU to DDP.

To detect the start of the FPDU unambiguously one of the following MUST be used:

- 1: In an ordered TCP stream, the ULPDU Length field in the current FPDU when FPDU has a valid CRC, can be used to identify the beginning of the next FPDU.
- 2: A Marker can always be used to locate the beginning of an FPDU (in FPDUs with valid CRCs). Since the location of the marker is known in the octet stream (sequence number space), the marker can always be found.

- 3: Having found an FPDU by means of a Marker, following contiguous FPDUs can be found by using the ULPDU Lengths (from FPDUs with valid CRCs) to establish the next FPDU boundary.

The ULPDU Length field (see section 6) MUST be used to determine if the entire FPDU is present before forwarding the ULPDU to DDP.

CRC calculation is discussed in section 7.2 on page 16 above.

7.4.1 Re-segmenting Middle boxes and non MPA-aware TCP senders

Since MPA on MPA-aware TCP senders start FPDUs on TCP segment boundaries, a receiving DDP on MPA on TCP implementation may be able to optimize the reception of data in various ways.

However, MPA receivers MUST NOT depend on FPDU Alignment on TCP segment boundaries.

Some MPA senders may be unable to conform to the sender requirements because their implementation of TCP is not designed with MPA in mind. Even if the sender is MPA-aware, the network may contain "middle boxes" which modify the TCP stream by changing the segmentation. This is generally interoperable with TCP and its users and MPA must be no exception.

The presence of markers in MPA allows an MPA receiver to recover the FPDUs despite these obstacles, although it may be necessary to utilize additional buffering at the receiver to do so.

Some of the cases that a receiver may have to contend with are listed below as a reminder to the implementer:

- * A single Aligned and complete FPDU, either in order, or out of order: This can be passed to DDP as soon as validated, and Delivered when ordering is established.
- * Multiple FPDUs in a TCP segment, aligned and fully contained, either in order, or out of order: These can be passed to DDP as soon as validated, and Delivered when ordering is established.
- * Incomplete FPDU: The receiver should buffer until the remainder of the FPDU arrives. If the remainder of the FPDU is already available, this can be passed to DDP as soon as validated, and Delivered when ordering is established.
- * Unaligned FPDU start: The partial FPDU must be combined with its preceding portion(s). If the preceding parts are already available, and the whole FPDU is present, this can be passed to DDP as soon as validated, and Delivered when ordering is established. If the whole FPDU is not available, the receiver should buffer until the remainder of the FPDU arrives.

- * Combinations of Unaligned or incomplete FPDUs (and potentially other complete FPDUs) in the same TCP segment: If any FPDU is present in its entirety, or can be completed with portions already available, it can be passed to DDP as soon as validated, and Delivered when ordering is established.

8 Connection Semantics

8.1 Connection setup

DDP on MPA requires that DDP's consumer MUST activate DDP, MPA, and any TCP enhancements for MPA, on a TCP half connection at the same location in the octet stream at both the sender and the receiver. This is required in order for the marker scheme to correctly locate the markers.

DDP, MPA, and any TCP enhancements for MPA, MAY be started separately in each direction, or enabled in both directions at once.

This can be accomplished several ways, and is left up to DDP's ULP:

- * DDP's ULP MAY require DDP on MPA startup immediately after TCP connection setup. This has the advantage that no additional negotiation is needed (at least for MPA). In this case the marker MUST be the first four octets sent (this marker has the special value 0x0000, meaning it belongs to the FPDU that follows).

This may be accomplished by using a well-known port, or a service locator protocol to locate an appropriate port on which DDP on MPA is expected to operate.

- * DDP's ULP MAY negotiate the start of DDP on MPA sometime after a normal TCP startup, using TCP streaming data exchanges on the same connection. The exchange establishes that DDP on MPA (as well as other ULPs) will be used, and exactly locates the point in the octet stream where MPA is to begin operation. Again, the marker is the first four octets sent when operation begins (this marker has the special value 0x0000, meaning it belongs to the FPDU that follows). Note that such a negotiation protocol is outside the scope of this specification. A simplified example of such a protocol is shown below.

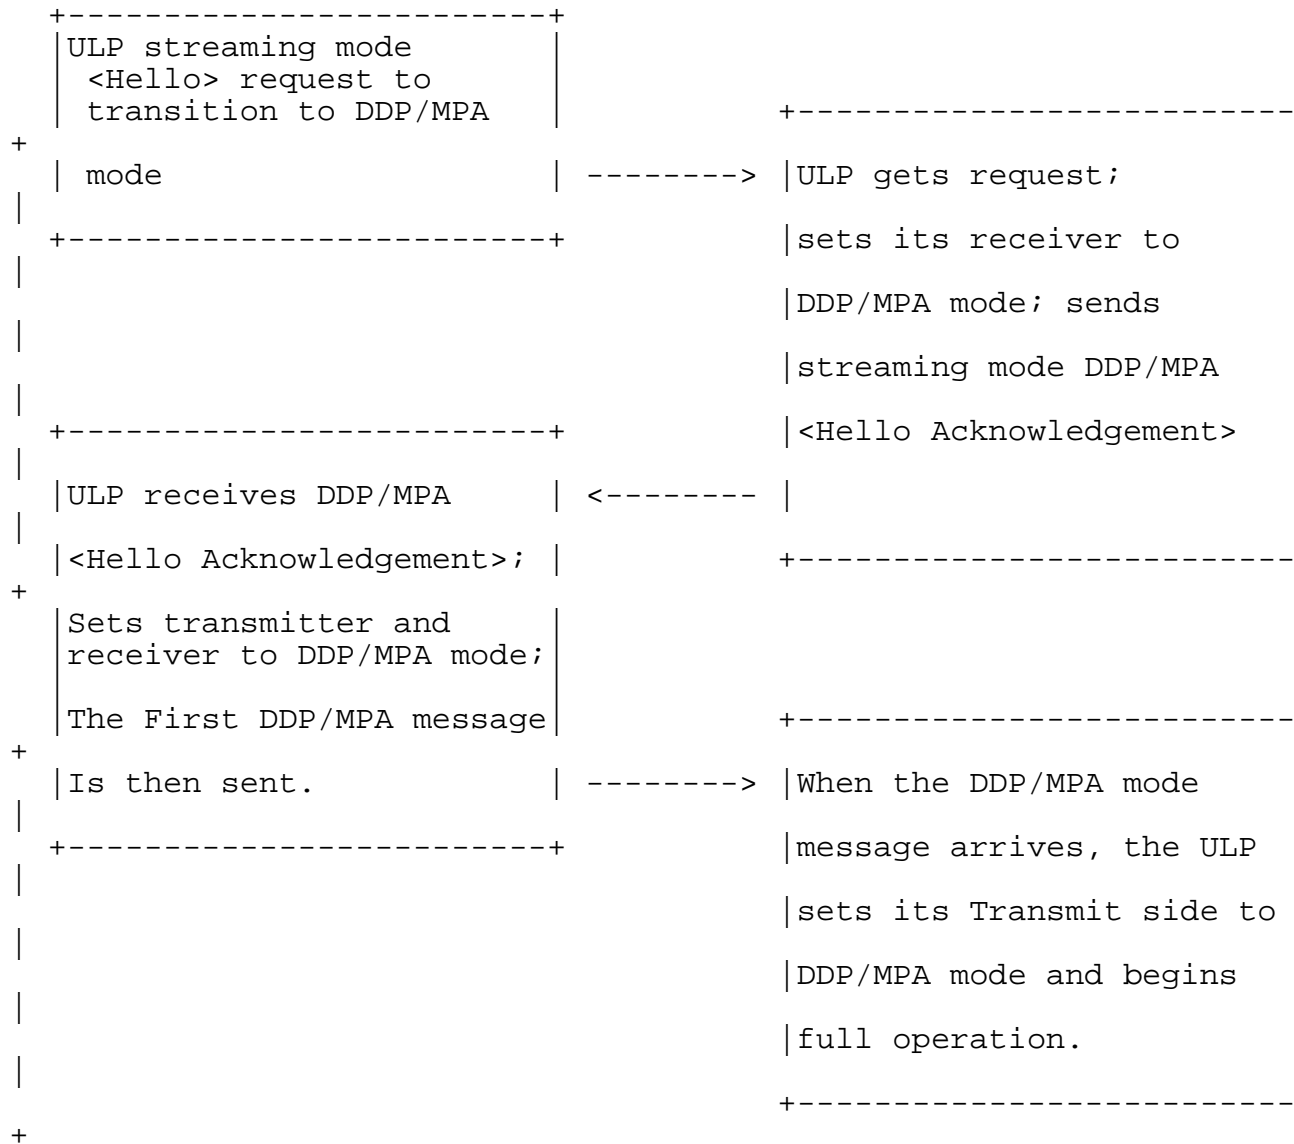


Figure 7: Example Startup negotiation

8.2 Normal Connection Teardown

Each half connection of MPA terminates when DDP closes the corresponding TCP half connection.

A mechanism SHOULD be provided by MPA to DDP for DDP to be made aware that a graceful close of the LLP connection has been received by the LLP (e.g. FIN is received).

9 Error Semantics

The following errors MUST be detected by MPA and the codes SHOULD be provided to DDP:

Code Error

- 1 TCP connection closed, terminated or lost. This includes lost by timeout, too many retries, RST received or FIN received.
- 2 Received MPA CRC does not match the calculated value for the FPDU.
- 3 In the event that the CRC is valid, received MPA marker and 'ULPDU Length' fields do not agree on the start of a FPDU. If the FPDU start determined from previous ULPDU Length fields does not match with the MPA marker position, MPA SHOULD deliver an error to DDP. It may not be possible to make this check as a segment arrives, but the check SHOULD be made when a gap creating an out of order sequence is closed and any time a marker points to an already identified FPDU. It is OPTIONAL for a receiver to check each marker, if multiple markers are present in an FPDU, or if the segment is received in order.

When conditions 2 or 3 above are detected, an MPA-aware TCP implementation MAY choose to silently drop the TCP segment rather than reporting the error to DDP. In this case, the sending TCP will retry the segment, usually correcting the error, unless the problem was at the source. In that case, the source will usually exceed the number of retries and terminate the connection.

Once MPA delivers an error of any type, it MUST NOT pass or deliver any additional FPDUs on that half connection.

MPA MUST NOT close the TCP connection following a reported error. Closing the connection is the responsibility of DDP's ULP.

Note that since MPA will not deliver any FPDUs on a half connection following an error detected on the receive side of that connection, DDP's ULP is expected to tear down the connection. This may not occur until after one or more last messages are transmitted on the opposite half connection. This allows a diagnostic error message to be sent.

10 Security Considerations

This section discusses the security considerations for MPA.

10.1 Protocol-specific Security Considerations

The vulnerabilities of MPA to third-party attacks are no greater than any other protocol running over TCP. A third party, by sending packets into the network that are delivered to an MPA receiver, could launch a variety of attacks that take advantage of how MPA operates. For example, a third party could send random packets that are valid for TCP, but contain no FPDU headers. An MPA receiver reports an error to DDP when any packet arrives that cannot be validated as an FPDU when properly located on an FPDU boundary. This would have a severe impact on performance. Communication security mechanisms such as IPsec [RFC2401] may be used to prevent such attacks. Independent of how MPA operates, a third party could use ICMP messages to reduce the path MTU to such a small size that performance would likewise be severely impacted. Range checking on path MTU sizes in ICMP packets may be used to prevent such attacks.

10.2 Using IPsec With MPA

IPsec can be used to protect against the packet injection attacks outlined above. Because IPsec is designed to secure individual IP packets, MPA can run above IPsec without change. IPsec packets are processed (e.g., integrity checked and decrypted) in the order they are received, and an MPA receiver will process the decrypted FPDUs contained in these packets in the same manner as FPDUs contained in unsecured IP packets.

11 IANA Considerations

If a well-known port is chosen as the mechanism to identify a DDP on MPA on TCP, the well-known port must be registered with IANA. Because the use of the port is DDP specific, registration of the port with IANA is left to DDP.

12 References

12.1 Normative References

- [iSCSI] Satran, J., "iSCSI", draft-ietf-ips-iscsi-20.txt (work in progress), January 2003.
- [RFC1191] Mogul, J., and Deering, S., "Path MTU Discovery", RFC 1191, November 1990.
- [RFC2018] Mathis, M., Mahdavi, J., Floyd, S., Romanow, A., "TCP Selective Acknowledgment Options", RFC 2018, October 1996.
- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.
- [RFC793] Postel, J., "Transmission Control Protocol - DARPA Internet Program Protocol Specification", RFC 793, September 1981.

12.2 Informative References

- [CRCTCP] Stone J., Partridge, C., "When the CRC and TCP checksum disagree", ACM Sigcomm, Sept. 2000.
- [DDP] H. Shah et al., "Direct Data Placement over Reliable Transports", draft-shah-iwarp-ddp-00.txt (Work in progress), October 2002
- [RFC2401] Atkinson, R., Kent, S., "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [RFC0896] J. Nagle, "Congestion Control in IP/TCP Internetworks", RFC 896, January 1984.
- [NagleDAck] Minshall G., Mogul, J., Saito, Y., Verghese, B., "Application performance pitfalls and TCP's Nagle algorithm", Workshop on Internet Server Performance, May 1999.
- [RDMA] R. Recio et al., "RDMA Protocol Specification", draft-recio-iwarp-rdmap-00.txt, October 2002
- [RFC2960] R. Stewart et al., "Stream Control Transmission Protocol", RFC 2960, October 2000.
- [RFC792] Postel, J., "Internet Control Message Protocol". September 1981
- [RFC1122] Braden, R.T., "Requirements for Internet hosts - communication layers". October 1989.

[ELZUR-MPA] Elzur, U., "Analysis of MPA over TCP Operations" draft-elzur-iwarp-mpa-tcp-analysis-00.txt, February 2003.

13 Appendix

This appendix is for information only and is NOT part of the standard.

13.1 Receiver implementation

13.1.1 Transport & Network Layer Reassembly Buffers

The use of reassembly buffers (either TCP reassembly buffers or IP fragmentation reassembly buffers) is implementation dependent. When MPA is enabled, reassembly buffers are needed if FPDU Alignment is lost or if IP fragmentation occurs. This is because the incoming out of order segment may not contain enough information for MPA to process all of the FPDU. For cases where a re-segmenting middle box is present, or where the TCP sender is not MPA-aware, the presence of markers significantly reduces the amount of buffering needed.

Recovery from IP Fragmentation must be transparent to the MPA Consumers.

13.1.1.1 Network Layer Reassembly Buffers

Most IP implementations set the IP Don't Fragment bit. Thus upon a path MTU change, intermediate devices drop the IP datagram if it is too large and reply with an ICMP message which tells the source TCP that the path MTU has changed. This causes TCP to emit segments conformant with the new path MTU size. Thus IP fragments under most conditions should never occur at the receiver. But it is possible.

There are several options for implementation of network layer reassembly buffers:

1. drop any IP fragments, and reply with an ICMP message according to [RFC792] (fragmentation needed and DF set) to tell the Remote Peer to resize its TCP segment
2. support an IP reassembly buffer, but have it of limited size (possibly the same size as the local link's MTU). The end Node would normally never advertise a path MTU larger than the local link MTU. It is recommended that a dropped IP fragment cause an ICMP message to be generated according to RFC792.
3. multiple IP reassembly buffers, of effectively unlimited size.
4. support an IP reassembly buffer for the largest IP datagram (64 KB).
5. support for a large IP reassembly buffer which could span multiple IP datagrams.

An implementation should support at least 2 or 3 above, to avoid dropping packets that have traversed the entire fabric.

There is no end-to-end ACK for IP reassembly buffers, so there is no flow control on the buffer. The only end-to-end ACK is a TCP ACK, which can only occur when a complete IP datagram is delivered to TCP. Because of this, under worst case, pathological scenarios, the largest IP reassembly buffer is the TCP receive window (to buffer multiple IP datagrams that have all been fragmented).

Note that if the Remote Peer does not implement re-segmentation of the data stream upon receiving the ICMP reply updating the path MTU, it is possible to halt forward progress because the opposite peer would continue to retransmit using a transport segment size that is too large. This deadlock scenario is no different than if the fabric MTU (not last hop MTU) was reduced after connection setup, and the remote Node's behavior is not compliant with [RFC1122].

13.1.1.2 TCP Reassembly buffers

A TCP reassembly buffer is also needed. TCP reassembly buffers are needed if FPDU Alignment is lost when using TCP with MPA or when the MPA FPDU spans multiple TCP segments.

Since lost FPDU Alignment often means that FPDUs are incomplete, an MPA on TCP implementation must have a reassembly buffer large enough to recover an FPDU that is less than or equal to the MTU of the locally attached link (this should be the largest possible advertised TCP path MTU). If the MTU is smaller than 140 octets, the buffer MUST be at least 140 octets long to support the minimum FPDU size. The 140 octets allows for the minimum MULPDU of 128, 2 octets of pad, 2 of ULPDU_Length, 4 of CRC, and space for a possible marker. As usual, additional buffering may provide better performance.

Note that if the TCP segment were not stored, it is possible to deadlock the MPA algorithm. If the path MTU is reduced, FPDU Alignment requires the source TCP to re-segment the data stream to the new path MTU. The source MPA will detect this condition and reduce the MPA segment size, but any FPDUs already posted to the source TCP will be re-segmented and lose FPDU Alignment. If the destination does not support a TCP reassembly buffer, these segments can never be successfully transmitted and the protocol deadlocks.

When a complete FPDU is received, processing continues normally.

14 Author's Addresses

Stephen Bailey
Sandburst Corporation
600 Federal Street
Andover, MA 01810 USA
Phone: +1 978 689 1614
Email: steph@sandburst.com

Paul R. Culley
Hewlett-Packard Company
20555 SH 249
Houston, Tx. USA 77070-2698
Phone: 281-514-5543
Email: paul.culley@hp.com

Uri Elzur
Broadcom
16215 Alton Parkway
CA, 92618
Phone: 949.585.6432
Email: uri@broadcom.com

Renato J Recio
IBM
Internal Zip 9043
11400 Burnett Road
Austin, Texas 78759
Phone: 512-838-3685
Email: recio@us.ibm.com

John Carrier
Adaptec Inc.
691 South Milpitas Blvd.
Milpitas, CA 95035
Phone: 360-378-8526
Email: John_Carrier@adaptec.com

15 Acknowledgments

Dwight Barron

Hewlett-Packard Company
20555 SH 249
Houston, Tx. USA 77070-2698
Phone: 281-514-2769
Email: dwight.barron@hp.com

Jeff Chase

Department of Computer Science
Duke University
Durham, NC 27708-0129 USA
Phone: +1 919 660 6559
Email: chase@cs.duke.edu

Ted Compton

EMC Corporation
Research Triangle Park, NC 27709, USA
Phone: 919-248-6075
Email: compton_ted@emc.com

Dave Garcia

Hewlett-Packard Company
19333 Vallco Parkway
Cupertino, Ca. USA 95014
Phone: 408.285.6116
Email: dave.garcia@hp.com

Hari Ghadia

Adaptec, Inc.
691 S. Milpitas Blvd.,
Milpitas, CA 95035 USA
Phone: +1 (408) 957-5608
Email: hari_ghadia@adaptec.com

Howard C. Herbert

Intel Corporation
MS CH7-404
5000 West Chandler Blvd.
Chandler, Arizona 85226
Phone: 480-554-3116
Email: howard.c.herbert@intel.com

Jeff Hilland
Hewlett-Packard Company
20555 SH 249
Houston, Tx. USA 77070-2698
Phone: 281-514-9489
Email: jeff.hilland@hp.com

Mike Ko
IBM
650 Harry Rd.
San Jose, CA 95120
Phone: (408) 927-2085
Email: mako@us.ibm.com

Mike Krause
Hewlett-Packard Corporation, 43LN
19410 Homestead Road
Cupertino, CA 95014 USA
Phone: +1 (408) 447-3191
Email: krause@cup.hp.com

Dave Minturn
Intel Corporation
MS JF1-210
5200 North East Elam Young Parkway
Hillsboro, Oregon 97124
Phone: 503-712-4106
Email: dave.b.minturn@intel.com

Jim Pinkerton
Microsoft, Inc.
One Microsoft Way
Redmond, WA, USA 98052
Email: jpink@microsoft.com

Hemal Shah
Intel Corporation
MS PTL1
1501 South Mopac Expressway, #400
Austin, Texas 78746
Phone: 512-732-3963
Email: hemal.shah@intel.com

Allyn Romanow
Cisco Systems
170 W Tasman Drive
San Jose, CA 95134 USA
Phone: +1 408 525 8836
Email: allyn@cisco.com

Tom Talpey
Network Appliance
375 Totten Pond Road
Waltham, MA 02451 USA
Phone: +1 (781) 768-5329
EMail: thomas.talpey@netapp.com

Patricia Thaler
Agilent Technologies, Inc.
1101 Creekside Ridge Drive, #100
M/S-RG10
Roseville, CA 95678
Phone: +1-916-788-5662
email: pat_thaler@agilent.com

Jim Wendt
Hewlett Packard Corporation
8000 Foothills Boulevard MS 5668
Roseville, CA 95747-5668 USA
Phone: +1 916 785 5198
Email: jim_wendt@hp.com

Jim Williams
Emulex Corporation
580 Main Street
Bolton, MA 01740 USA
Phone: +1 978 779 7224
Email: jim.williams@emulex.com

16 Full Copyright Statement

This document and the information contained herein is provided on an "AS IS" basis and ADAPTEC INC., AGILENT TECHNOLOGIES INC., BROADCOM CORPORATION, CISCO SYSTEMS INC., DUKE UNIVERSITY, EMC CORPORATION, EMULEX CORPORATION, HEWLETT-PACKARD COMPANY, INTERNATIONAL BUSINESS MACHINES CORPORATION, INTEL CORPORATION, MICROSOFT CORPORATION, NETWORK APPLIANCE INC., SANDBURST CORPORATION, THE INTERNET SOCIETY, AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright (c) 2002 ADAPTEC INC., BROADCOM CORPORATION, CISCO SYSTEMS INC., EMC CORPORATION, HEWLETT-PACKARD COMPANY, INTERNATIONAL BUSINESS MACHINES CORPORATION, INTEL CORPORATION, MICROSOFT CORPORATION, NETWORK APPLIANCE INC., All Rights Reserved