Document: T10/01-230r0 Date: 18 July 2001
To: T10 Committee Membership
From: Edward A. Gardner, Ophidian Designs
Subject: SRP Buffer Descriptor Subclause Rewrite

My notes from the June 19-20 SRP working group include substantial changes requested to the description of data buffer descriptors. However I can find no mention of these in the official minutes of that meeting (01-195). Since the minutes do spell out much more trivial changes, I hesitate to include these changes without approval.

The remainder of this document is the affected portions of SRP with change bars from revision 07.

4.4.2 Data buffer descriptors

4.4.2.1 Overview

An SRP_CMD request (see 5.7) may contain a data-out buffer descriptor, a data-in buffer descriptor, both or neither, depending upon the data transfer(s) requested by the SCSI command. For each data buffer descriptor an indirect bit indicates whether the buffer descriptor is direct or indirect. A count field indicates the number of memory descriptors present in the data buffer descriptor. The indirect bit and count field together determine the length and format of the data buffer descriptor. Table 2 lists the combinations of the indirect bit and count field and the resulting data buffer descriptor length and format.

Table 2 - Data buffer descriptor formats

indirect bit[1]	count field[2]	Length (bytes)[3]	Data buffer descriptor format
0	0	0	Buffer descriptor not present.
0	1	16	Direct data buffer descriptor (see 4.4.2.2).
0	m > 1		Reserved.
1	0		Reserved.
1	1	20	Indirect data buffer descriptor without a PARTIAL MEMORY DESCRIPTOR LIST (see 4.4.2.3).
1	m > 1	16*m+4	Indirect data buffer descriptor with a PARTIAL MEMORY DESCRIPTOR LIST CONTAINING m-1 memory descriptors (see 4.4.2.3).

Note 1 The indirect bit for a data-out buffer descriptor is bit DOIND of the SRP_CMD request. The indirect bit for a data-in buffer descriptor is bit DIIND of the SRP_CMD request. See 5.7.
Note 2 The count field for a data-out buffer descriptor is field DOCOUNT of the SRP_CMD request. The count field for a data-in buffer descriptor is field DICOUNT of the SRP_CMD request. See 5.7.
Note 3 The length of a data buffer descriptor is determined from its indirect bit and count field values.

Support for indirect data buffers is optional. SRP targets indicate whether they support indirect data buffers during RDMA channel establishment.

MAXIMUM INITIATOR TO TARGET IU, negotiated during RDMA channel establishment (see 5.3), imposes an implicit limit on the length of an indirect data buffer descriptor. The actual limit for any particular indirect data buffer descriptor also depends upon the CDB length and the number of buffer descriptors required by the command.

4.4.2.2 Direct data buffers

The presence of a direct data buffer descriptor is indicated by the indirect bit containing 0 and the count field containing 1. A direct data buffer descriptor contains a single memory descriptor (see table 1). The memory descriptor identifies the data buffer, which is a single memory segment within a memory region's virtual address space. If a direct data buffer descriptor defines a data-out buffer, the SRP target shall only issue RDMA Read operations using the memory descriptor contained in the direct data buffer descriptor. If a direct data buffer descriptor defines a data-in buffer, the SRP target shall only issue RDMA Write operations using the memory descriptor contained in the direct data buffer descriptor. The SRP target shall use the DATA LENGTH field of the memory descriptor as the length of the data-out buffer or data-in buffer.

4.4.2.3 Indirect data buffers

The presence of an indirect data buffer descriptor is indicated by the indirect bit containing 1 and the count field containing 1 or larger. An indirect data buffer contains one or more memory segments, which may or may not be contiguous. The memory segments may be in a single memory region or spread among several memory regions. The data buffer is the concatenation of the memory segments. Each memory segment may have any length, including a length of zero bytes. Table 3 shows the format of an indirect data buffer descriptor.

Table 3 - Indirect data buffer descriptor

Byte \ Bit	7	6	5	4	3	2	1	0
0	INDIRECT TABLE MEMORY DESCRIPTOR							
• • •								
15								
16	(MSB) TOTAL LENGTH							
• • •								
19								(LSB)
20	PARTIAL MEMORY DESCRIPTOR LIST							
• • •								
3+16*m								
Note 1 m is the value contained in the data buffer descriptor's count field.								

INDIRECT TABLE MEMORY DESCRIPTOR identifies a memory segment that contains an indirect table. An indirect table is a list of one or more memory descriptors. The memory segments identified by the memory descriptors in the indirect table, concatenated together, comprise the indirect data buffer. The DATA LENGTH field of INDIRECT TABLE MEMORY DESCRIPTOR contains the number of memory descriptors in the indirect table times 16. Target behavior when INDIRECT TABLE MEMORY DESCRIPTOR contains any other value is vendor specific.

TOTAL LENGTH contains the sum of the contents of the DATA LENGTH fields of the memory descriptors in the indirect table. Target behavior when TOTAL LENGTH contains any other value is vendor specific. The target shall use either the TOTAL LENGTH field or the sum of the DATA LENGTH fields as the length of the data-out buffer or data-in buffer.

PARTIAL MEMORY DESCRIPTOR LIST is only present when the data buffer descriptor's count field contains a value larger than 1. It is a list of m-1 memory descriptors, where m is the value contained in the count field. PARTIAL MEMORY DESCRIPTOR LIST contains copies of the first m-1 memory descriptors in the indirect table. Target behavior when PARTIAL MEMORY DESCRIPTOR LIST contains any other value is vendor specific.

An SRP target shall only issue RDMA Read operations to the indirect table.

If an indirect data buffer descriptor defines a data-out buffer, the SRP target shall only issue RDMA Read operations using the memory descriptors contained in the indirect table or the PARTIAL MEMORY DESCRIPTOR LIST.

If an indirect data buffer defines a data-in buffer, the SRP target shall only issue RDMA Write operations using the memory descriptors contained in the indirect table or the PARTIAL MEMORY DESCRIPTOR LIST.

Figure 3 illustrates an indirect data buffer without a PARTIAL MEMORY DESCRIPTOR LIST. Memory is shown containing four memory segments: the indirect table, memory segment 1, memory segment 2 and memory segment 3. The mapping of each memory descriptor to its memory segment has been shown as a single arrow. For details of this mapping see 4.4.1 and figure 2. Figure 3 does not show the memory regions in which the memory segments reside. All four segments might be in a single memory region, each might be in a separate memory region, or several might be in one memory region and the remainder in one or more other memory regions.

In this example the indirect flag is 1 and the count field is 1. The data buffer descriptor is 20 bytes long. The data buffer is comprised of three memory segments: memory segment 1, memory segment 2 and memory segment 3. A separate memory segment contains the indirect table, a list of three memory descriptors identifying memory segments 1 through 3. The INDIRECT TABLE MEMORY DESCRIPTOR field of the data buffer descriptor identifies the memory segment containing the indirect table. The DATA LENGTH field of INDIRECT TABLE MEMORY DESCRIPTOR would contain 48, the length of the indirect table. The TOTAL LENGTH field of the data buffer descriptor would contain the sum of the DATA LENGTH fields of the memory descriptors in the indirect table (the sum of DATA LENGTH 1, DATA LENGTH 2 and DATA LENGTH 3). This sum is the total length of the data buffer.

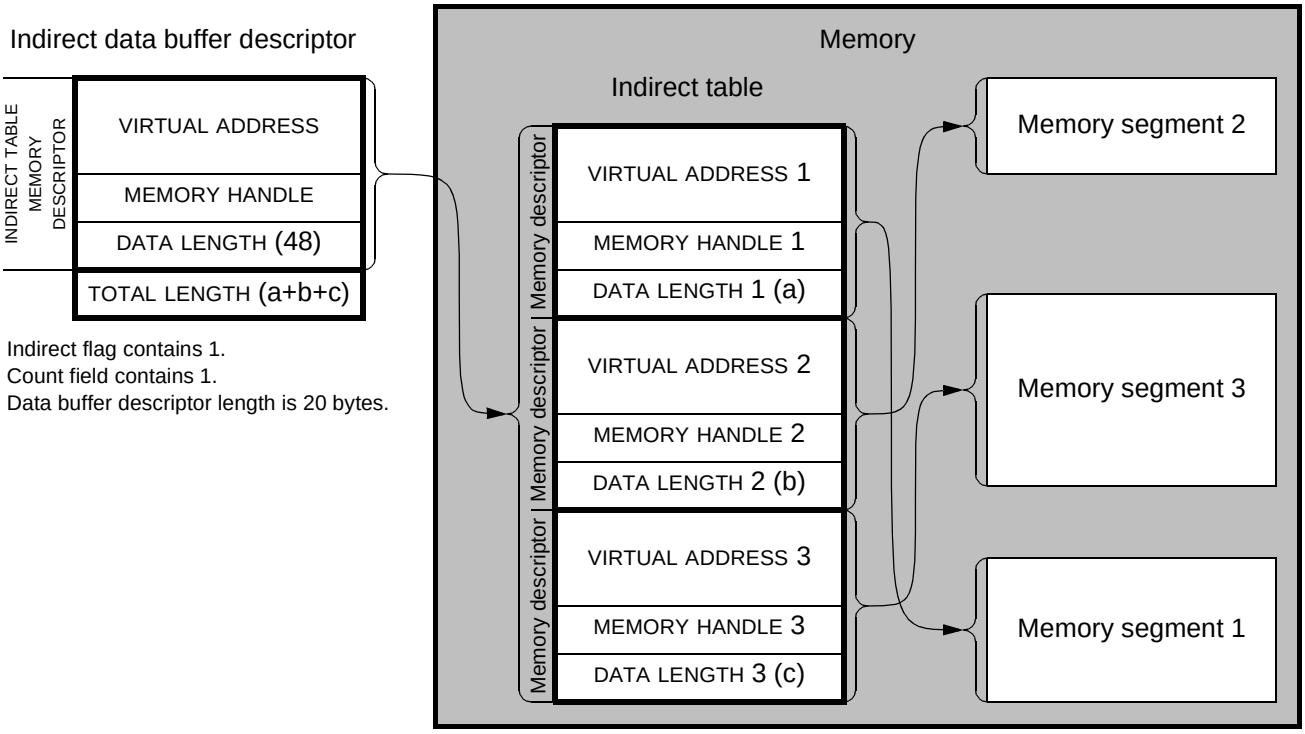

Figure 3 - Indirect data buffer mapping without cached memory descriptors

Figure 4 illustrates the same example except with a PARTIAL MEMORY DESCRIPTOR LIST. The data buffer, indirect table, INDIRECT TABLE MEMORY DESCRIPTOR and TOTAL LENGTH are all identical to the previous example. The indirect bit is still 1, but the count field is now 3, indicating that the PARTIAL MEMORY DESCRIPTOR LIST contains two memory descriptors. These are identical to the first two memory descriptors in the indirect table. The third memory descriptor is only present in the indirect table. The data buffer descriptor is 52 bytes long.

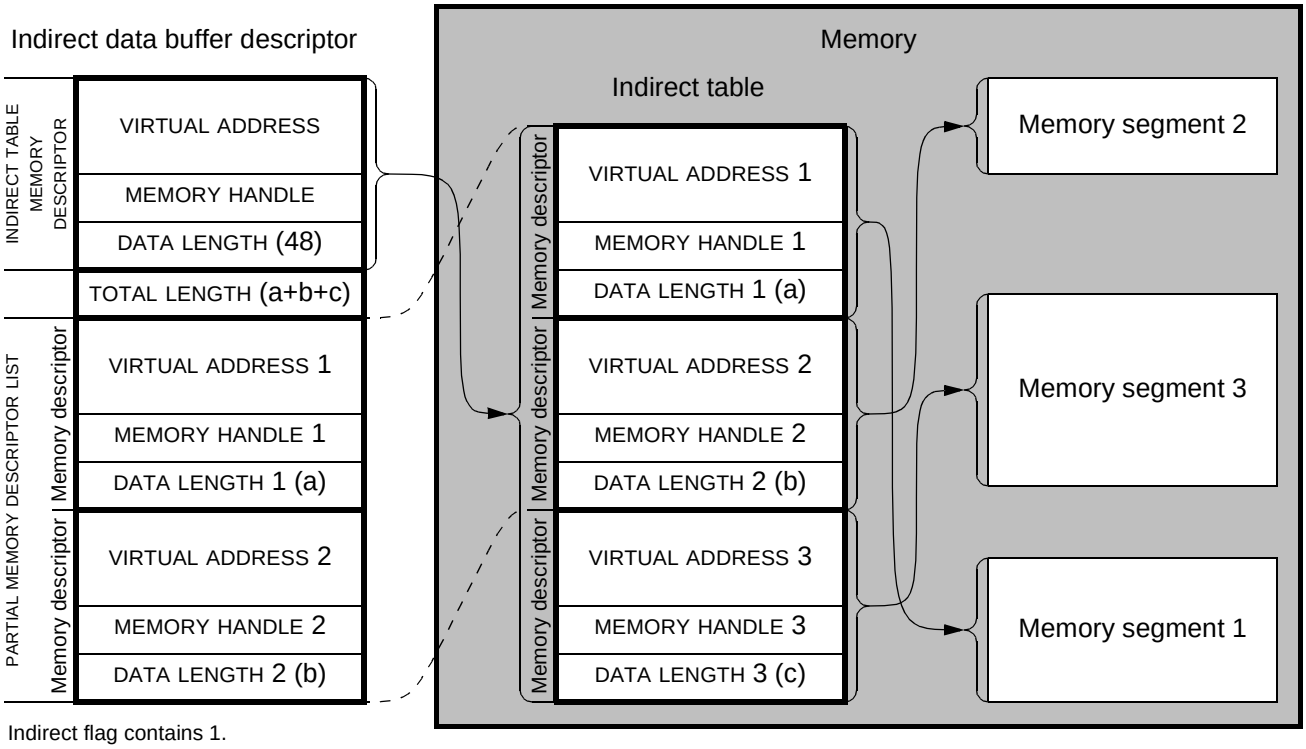

Indirect flag contains 1.
Count field contains 3.
Data buffer descriptor length is 52 bytes.

Figure 4 - Indirect data buffer mapping with cached memory descriptors

5.7 SRP_CMD request

...

The DOIND bit contains the indirect bit and the DOCOUNT field contains the count field for the DATA-OUT BUFFER DESCRIPTOR field (see 4.4.2). DOIND shall contain zero if the value of INDSUP was zero in the SRP_LOGIN_RSP.

The DIIND bit contains the indirect bit and the DICOUNT field contains the count field for the DATA-IN BUFFER DESCRIPTOR field (see 4.4.2). DIIND shall contain zero if the value of INDSUP was zero in the SRP_LOGIN_RSP.