

# *iSCSI – a SCSI over TCP mapping*


## *IETF – 50 – New Format*




Barry Reinhold

Julian Satran


IBM Research Lab in Haifa






## *New Formats*

- ◆ Format considerations
  - ◆ Do not expand header
  - ◆ Read Header in maximum 2 steps
  - ◆ Do not use unchecked data length
- ◆ Format versions
  - ◆ Format 1
    - ◆ 1-Lengths always double checked (parity)
    - ◆ Any header length
  - ◆ Format 2
    - ◆ Header Length has to used unchecked
    - ◆ Headers total limited a 255 – 4-byte words



## *New Format-1*

QL	Data or Additional Header Length	
BHS		
DQ	Data Length	
Type	Length	AHS field
AHS fields		
Type	Length	AHS field
Digest		




## *New Format-1 coding*

### ◆ QL

- ◆ B7&6-0 Only data, 1 Only AHS, 2 Both
  - ◆ B5-4 Reserved
  - ◆ B3-0 Parity


### ◆ DQ

- ◆ B7-4 Reserved
- ◆ B3-0 Parity



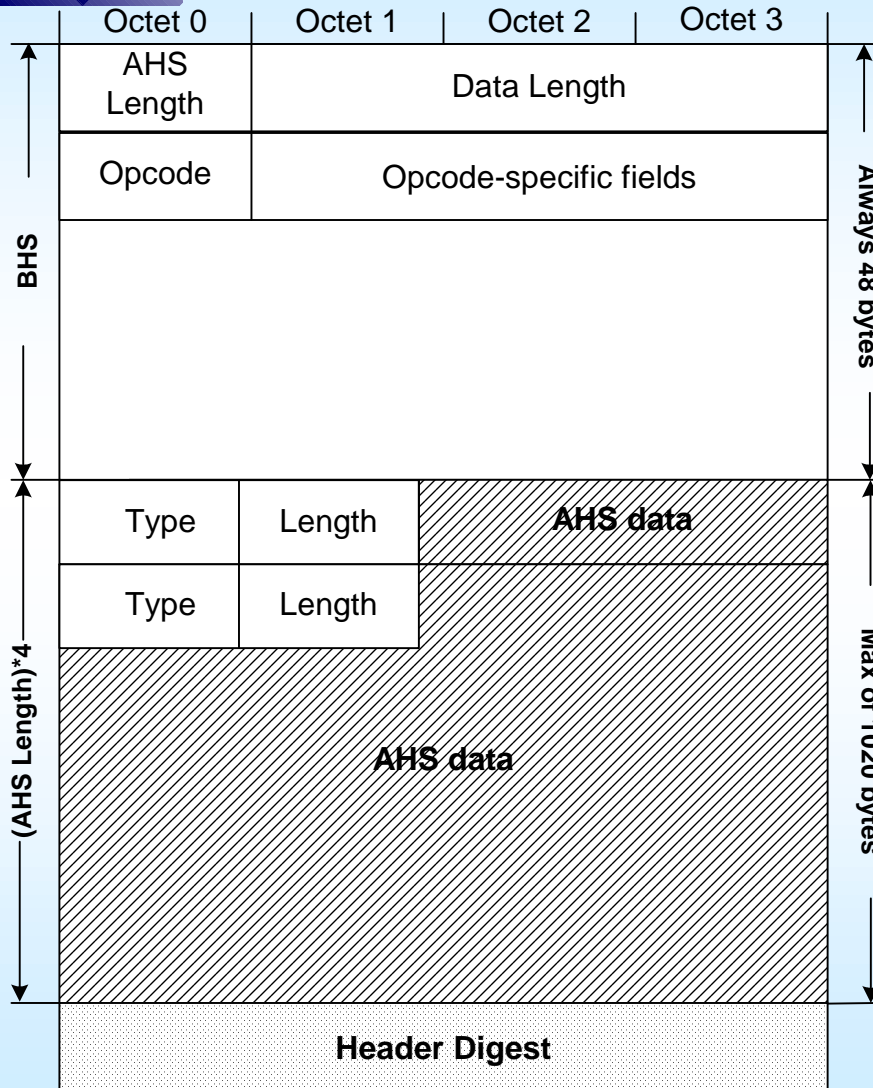
## *New Format-2*

AHS T-length	Data Length	
BHS		
Type	Length	AHS field
AHS fields		
Type	Length	AHS Field
Digest		



# iSCSI PDU Header

## Notes




1. AHS\_length is given in 4 byte words allowing a maximum size of 1020 bytes for all combined AHSEs.


2. Data Length is measured in bytes, allowing a maximum of 16 Mbytes of data in a single iSCSI PDU.

3. The AHS\_length field is used before the header is checked against the digest. If the AHS field is corrupt it can cause blocking on the TCP stream.

4. The length of an individual ANS is limited to 1020 bytes. AHSEs are word aligned through the use of padding.




## AHS Format



### Fields in Additional Header Segment

1. **AHS\_type** - First field in AHS, always starts on a 32 bit word boundary, 1 byte in length. The type field identifies how to interpret the AHS data field. It is broken into sub fields as follows:
  - bit 0 (MSB): Drop bit - If set this PDU must be dropped if the receiver does not understand how to interpret the AHS.
  - bit 1 : reserved - Set to zero by transmitter, and ignored by receiver.
  - bits 2-7: **AHS\_type** - The value contained in this field indicates how the bits in the **AHS\_data** field are to be used. The field is enumerated, and can take on any value from 0 to 63. Values from 0-62 are reserved for assignment by iSCSI. An enumerated value of 63 indicates that the **AHS\_data** field contains information not defined by iSCSI. In this case the first N bytes of the **AHS\_data** field shall contain an OUI. Size and layout TBD. [ 3 byte IEEE OUI in 6 byte space?]
2. **AHS\_length** - Second field in AHS, 1 byte in length, starting after the **AHS\_type** field. The **AHS\_length** field contains the number of 4 byte words allocated for the complete AHS. This number includes the word containing the **AHS\_type** and **AHS\_length** field itself. Hence the smallest possible AHS has a **AHS\_length** of 1. The largest value is 255 words or 1020 bytes.
3. **AHS\_data**- Third field in AHS, between 2-1018 bytes in length, starts after **AHS\_length** field. The information contained in this field is dependent on the value contained in the **AHS\_type** field. AHS formats defined by this document are given in section TBD.



**Example AHSes**

Possible format for  
ADDCDB

AHS_type			AHS_length	AHS_data
0	1	2-7		
1	0	000001 (ADDCDB)	0x02 (allocated space in 4 byte words)	Actual number of bytes of data (0x04)
AHS_data				
ADDCDB payload (4 bytes)				

Possible format  
for Small Integer


AHS_type			AHS_length	AHS_data
0	1	2-7		
0	0	000101 (FCP_CRN)	0x01 (allocated space in 4 byte words)	0x008c

Example format of an AHS containing  
an IEEE OUI

AHS_type			AHS_length	AHS_data
0	1	2-7		
0	0	111111 (not iSCSI)	0x03 (allocated space in 4 byte words)	0x0000
AHS_data				
0x0001ac4f (organizational OUI)				
AHS_data				
Information encoded as defined by OUI 0x01ac4f (4 bytes)				



I JUST LOVE HAPPY ENDINGS



MarkParisi@aol.com  
www.offthemark.com

ALL AN 10-11-11, 12-6-1993 MARK PARISI