

Douglas kept reminding everyone who can read of his view that this effort was nothing less than an attempt to rape and pillage TCP beyond recognition.

Toward the end of the month Joseph Touch (University of Southern California) joined Douglas and reprised his negative role from the Minneapolis meeting (March Happenings).

Andreas Jungmaier (University of Essen) also had plenty of negative things to say but after exchanging a couple of messages with Stephen, Andreas dug deep into the proposal and surfaced a few ideas for improving it.

The Transport Services working group leadership believes that a favorable consensus exists for the Framing proposal provided all the i's are dotted and t's crossed. This view gained some credence when Brian F. G. Bidulock, a name not seen before on the iSCSI reflector weighed in to support Stephen and David.



### Mode Pages

Long time SCSI practitioners will remember the Disconnect-Reconnect Mode Page as being the home for a set of knobs that control the performance of the underlying interconnect. Back in the days when the only interconnect was SCSI parallel, the pieces sort of fitted and T10 has done a passable job of keeping the picture coherent through the development of Fibre Channel, IEEE (Institute for Electrical and Electronics Engineers) 1394, SSA (Serial Storage Architecture), and other SCSI transports.

iSCSI adds a new wrinkle in the cloth by allowing most of those interesting Disconnect-Reconnect parameters to be negotiated during Login. It further compounds the turn from parallel SCSI days by prohibiting changes made via the Mode Select command.

Eddy Quicksall had a few choice words on that last idea.

"I don't like the idea of not letting the user of a SCSI utility be able to change some of the parameters for iSCSI."

"They may be relevant to him and there may not be a user interface to the iSCSI driver."

"Parallel SCSI sets these low level parameters via a standard Mode Select command, so why not iSCSI?"

The duty of replaying the history fell on Julian Satran.

"This issue was raised several times."

"Having parameters changed in two layers is not a good idea."

"Leaving only SCSI change them will not enable us to restrict some changes to login or initial login."

"The compromise we have does not break any existing software and allows us to do what we want."

A site license is granted to subscribers so that the ENDL Letter may be copied for wide distribution within the organization.

The ENDL Letter is to be treated with the same confidentiality as internal company memos. Your co-operation is appreciated.

# CONFIDENTIAL

## ENDL Letter

Page 88

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

Bob Griswold (Crossroads) was seeing two as better than one.

"I would argue that having two ways to change the same thing adds flexibility, rather than complexity."

"If the iSCSI Text Messages are allowed to change iSCSI only parameters, then the transport has access to what it needs."

"In addition, if Mode Pages work, the old SCSI utilities will work with iSCSI."

"As stated by Eddy, the legacy of SCSI is that Mode Pages, irrespective of what those parameters are manipulating (transport, commands, responses, timing, etc), are accessible through SCSI mode commands."

Bob Snively (Brocade) saw something more than an all-or-nothing deal.

"Both of you are right."

"iSCSI should stick to negotiating only parameters related with the TCP connections i.e. the iSCSI security parameters and the iSCSI session operational parameters other than EMDP (Enable Modify Data Pointers) and burst length."

"All parameters relating to the normal SCSI activities, including EMDP and burst length (which are required for SCSI but not iSCSI, buffer management) would be negotiated in the normal SCSI manner through the Mode Sense/Mode Select pages."

"Everybody would live happily ever after."

Julian was ready to meet Bob Snively half way.

"I think that for regularity it would be wise to completely remove EMDP and FirstBurstSize from iSCSI control."

"MaximumBurstSize defines the Data PDU length and is transport related."

"Its effects are more iSCSI than SCSI related and so it should stay under iSCSI control."

Charles Binford (Pirus Networks) was not buying a half solution.

"I agree with Julian about removing EMDP and FirstBurstSize from iSCSI control."

"I disagree on MaxBurstSize."

"MaxBurstSize defines Data PDU length only because we said it does and that probably is a bad idea."

"We could just as easily remove the relationship between Data PDU length and MaxBurstSize."

"I would prefer we let iSCSI control Data PDU length as it does today and change the meaning of MaxBurstSize under iSCSI to be more in line with what MaxBurstSize means under other SCSI protocols."

Charles pointed out that such a change would make iSCSI more natural for a generic SCSI utility to manipulate.

Charles hit the specifics hard.

A site license is granted to subscribers so that the ENDL Letter may be copied for wide distribution within the organization.

The ENDL Letter is to be treated with the same confidentiality as internal company memos. Your co-operation is appreciated.

# CONFIDENTIAL

## ENDL Letter

Copyright © ENDL Incorporated. Reproduction without permission is prohibited.

A site license is granted to subscribers so that the ENDL Letter may be copied for wide distribution within the organization.

"The meaning of MaxBurstSize under FCP (SCSI Fibre Channel Protocol) is the largest Sequence a Fibre Channel target can send."

"Since a Class 3 target can send multiple Read data sequences to the host with no handshake, the parameter is essentially a don't care for reads."

"For Writes it effectively governs the largest amount of data a target can ask for on an XFER\_RDY (Transfer Ready) which is R2T (Request To Transfer) in iSCSI."

Charles' next step was to propose two alternative for the definition of Max-BurstSize.

- o The largest amount of data a target can ask for under a single R2T, or
- o Ignored.

"My problem with the current definition is typical iSCSI Data PDU lengths are relatively small (default 8KB in the spec)."

"The typical MaxBurstSize for other protocols is much larger, 64-512KB."

Charles' position held sway with Julian and a week later revised iSCSI text appeared which instantiated the idea.

### Snapshots

**Error Recovery:** Mallikarjun Chadalapaka (Hewlett Packard) reviewed the current status of the iSCSI error recovery work in London. A major improvement is the definition of four hierarchical levels of error recovery ranging from the bare minimum to full power command replay. The group liked the hierarchical idea and sent Mallikarjun home with instructions to focus on the details.

The required implementation levels may fall at different places in the hierarchy for device types i.e. tape error recovery needs vary from those for disks, a truth well known to anyone who follows SCSI or Fibre Channel.

The elimination of SNACK was discussed but SNACK stayed in the draft because supporting tapes without SNACK was too horrifying to contemplate. SNACK is an iSCSI request PDU that asks for retransmission of another PDU using the Sequence Number that should have been in the requested PDU.

Note: If SNACK is an acronym for something then that something is far from obvious. Maybe, just maybe, it is Sequence Number ACKnowledge, but that does not make any sense unless it is based on historical usage of a term, as retransmission and acknowledgment are two completely different things.

**Security Gateways:** The FCIP side of the house tried hard this month to get a T11-like ruling on the security requirements that must be stated in the FCIP draft. The goal was to have FCIP document the protocol aspects of FCIP and not the security requirements. The defense for this approach was based on a layering concept with FCIP devices attaching something called a security gateway in order to be the really secure devices which the IETF wants everything built under its auspices to be.

The ENDL Letter is to be treated with the same confidentiality as internal company memos. Your co-operation is appreciated.