


iSCSI – a SCSI over TCP mapping


IETF52-Salt Lake City – December -2001

Julian Satran - IBM

Mallikarjun Chadalapaka - HP


IBM Research Lab in Haifa






Status

- ◆ Not completely closed “chapters”
 - ◆ Security
 - ◆ Tunnel vs. Transport IPsec
 - ◆ Strong
 - ◆ Framing ?
 - ◆ COWS as an alternative option
 - ◆ Abort & Clear Task set
 - ◆ OOO PDU handling
- ◆ Open items
 - ◆ NOP in discovery session




Miscellaneous items to fix

- ◆ Abort task set – finalize cleanup at target before TM response (instead of the current finalize cleanup at initiator and after TM response)
- ◆ Remove reference to ordering of requests from other initiators for Clear Task set



Abort task set - today

- ◆ Before handover to SCSI cleaning – barrier mechanism
- ◆ At target TM request given to SCSI followed by an immediate return of a TM response
- ◆ A response expected at initiator for every ITT outstanding
- ◆ ITTs ready for reuse only after SCSI response returned possibly long after TM response
- ◆ SCSI response can't be trusted




Abort task set after change

- ◆ Before handover to SCSI cleaning – barrier mechanism only at target
- ◆ At target TM request given to SCSI
- ◆ After SCSI answers at target TM wait for all commands to reach target and all outstanding responses on all connections to be acked and returns TM response after that
- ◆ ITTs ready for reuse TM response returned
- ◆ SCSI response can be trusted




Out of Order PDU handling

- ◆ Order within a single connection maintained after an item is assigned a number is a widely held assumption in all exception handling
- ◆ The overlapping of command shipping and DMA can be achieved by prefetching [OR] by
- ◆ the already provided by the multiple connection mechanism




NOP in discovery sessions

- ◆ Do we want to have long lived discovery sessions?
- ◆ A majority of e-mails seem to indicate that mechanism provided by other management mechanisms are considered sufficient



Framing

- ◆ Status of draft? Experimental ? Informational?
- ◆ Does it make sense to add a COWS scheme to iSCSI ?
 - ◆ It is simple in all hardware implementations – even for NICs offering only simple assists
 - ◆ For small marking intervals (marking at almost every PDU) involves almost the same overhead as marking in software implementations
 - ◆ It requires touching every word but if done within a CRC or Ipsec-ESP software implementation it can be even less expensive than markers



Editorial

- ◆ Chapter 3 becomes also an overview – detailed commands in later chapter
- ◆ Appendixes (except framing, recovery and examples) moved in main body