

iSCSI Initiator Experience

Kalman Meth

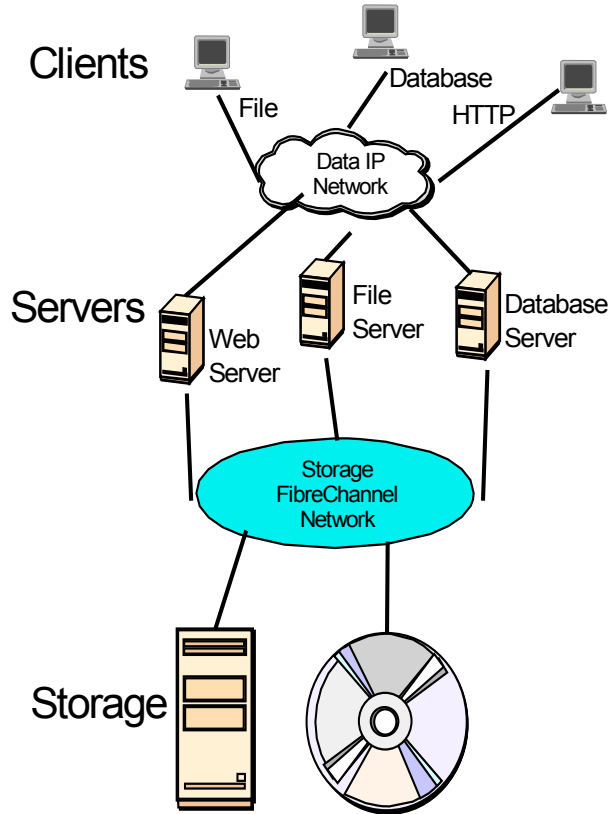
IBM Haifa Research Lab

April 2002

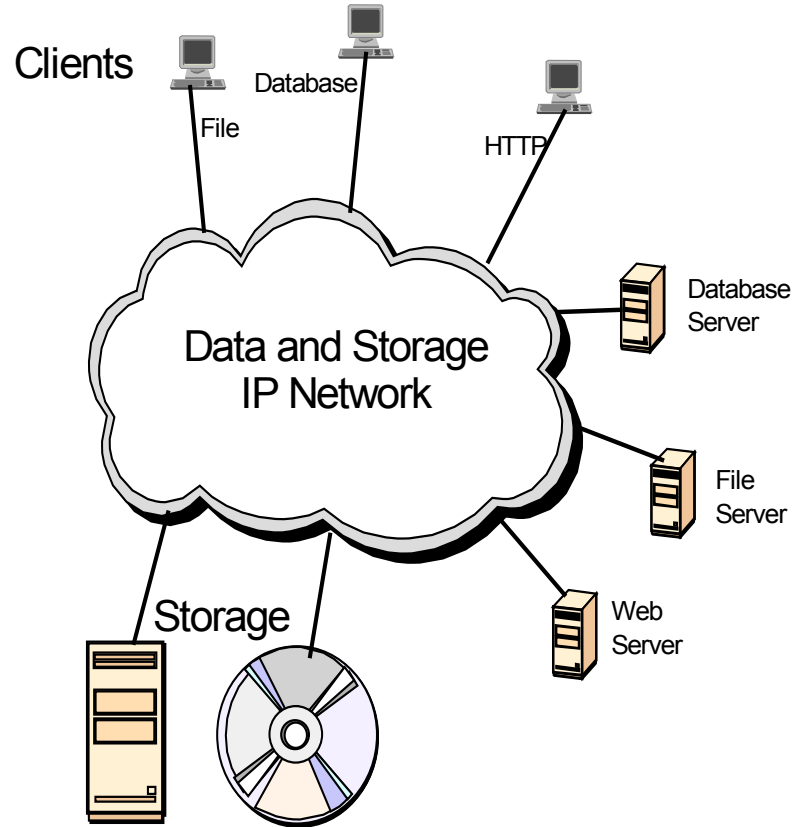
What is iSCSI?

- SCSI is a protocol for devices such as disk, tape, CD ROM
- iSCSI = Internet SCSI = SCSI over TCP/IP
 - send SCSI commands over an IP network
- Utilize features of TCP
 - reliable, ordered delivery
 - retransmission of dropped packets
 - congestion control

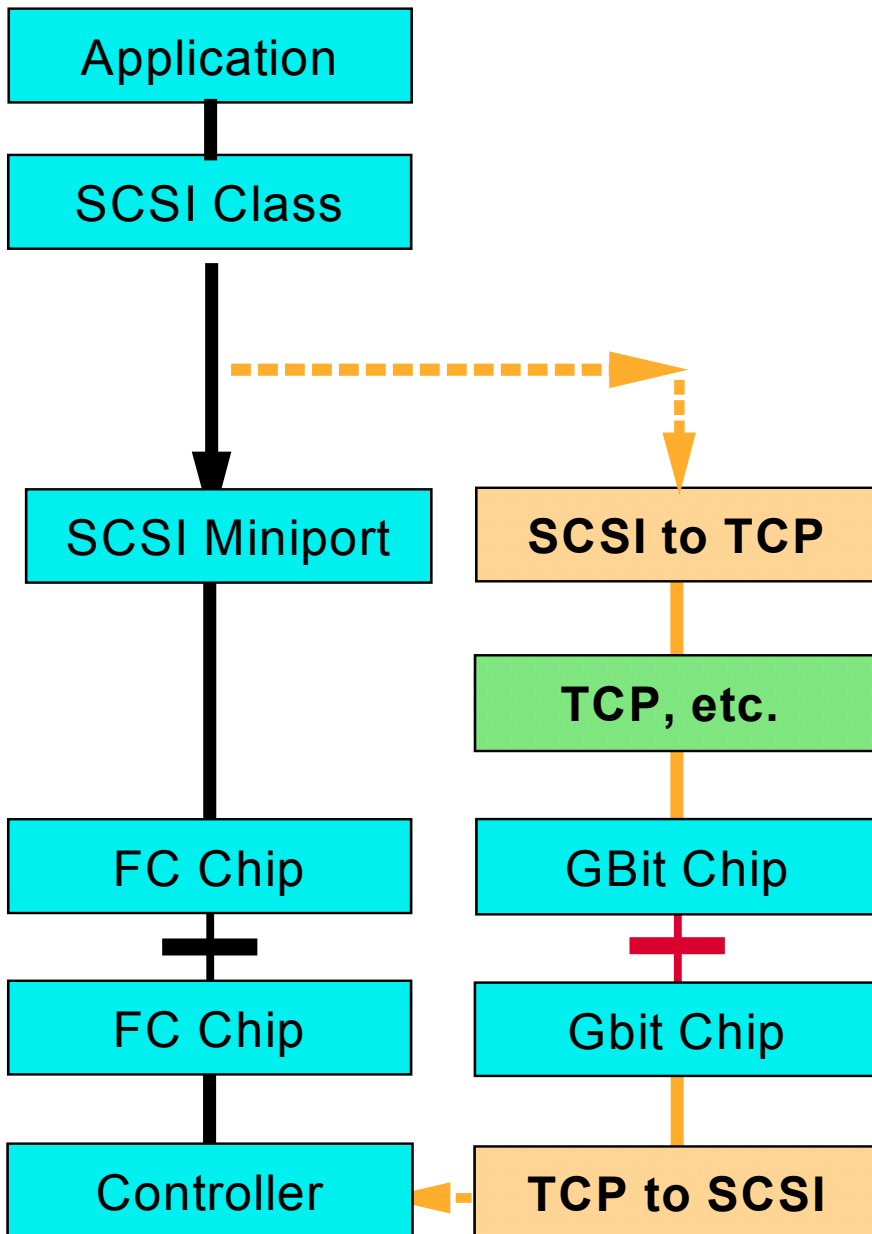
Classic SAN



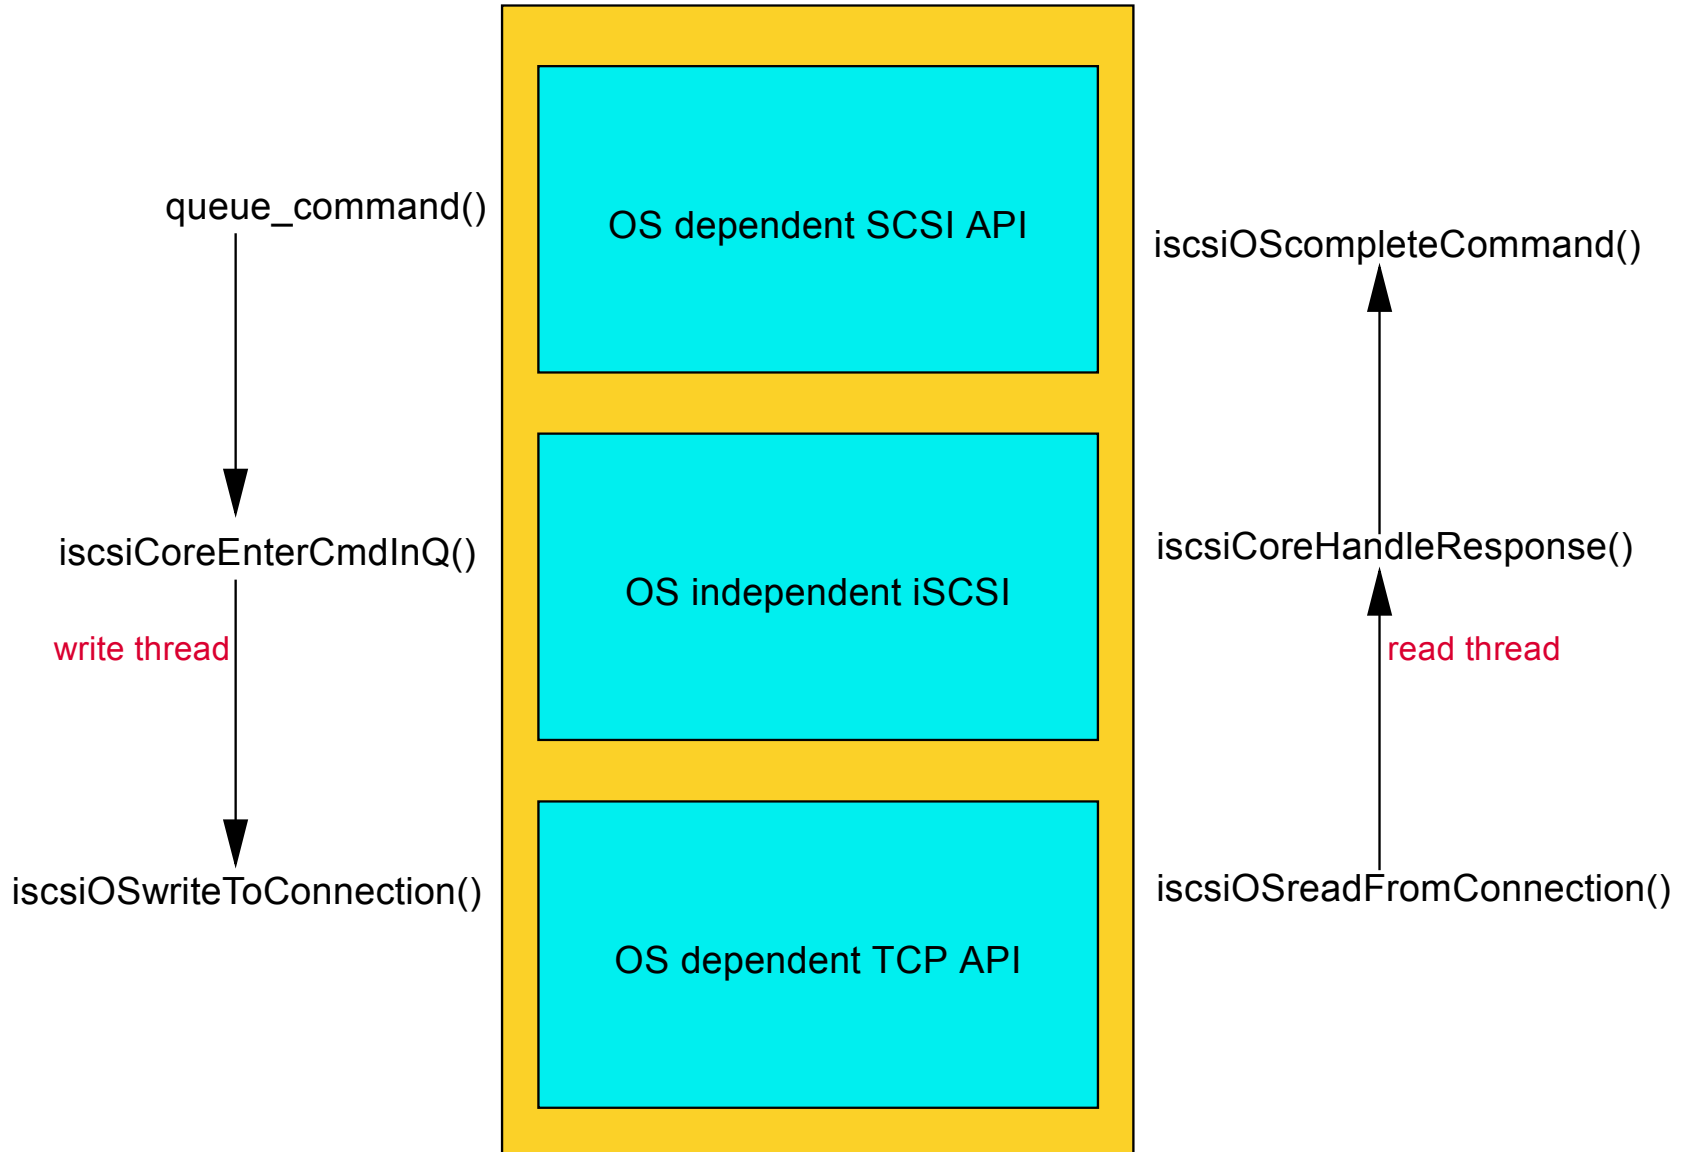
iSCSI



iSCSI layering



SCSI to TCP



Implementation Pointers

- Commands may be issued from within an interrupt handler
- Command completion expected to occur inside interrupt handler
- Problem to mount and unmount file system automatically upon reboot and shutdown

Comments on Performance

- Usually faster than local IDE
 - used IBM's TotalStorage 200i controller
 - SCSI RAID more than compensates the network overhead
 - biggest boost when performing multiple simultaneous IO operations
- Higher CPU utilization (for software TCP)
- Many network interrupts for large transfers
 - due to small MTU size