Remote Mirroring over Low-bandwidth WAN with iSCSI

Ming Zhang

University of Rhode Island

NASA/IEEE MSST 2004
12th NASA Goddard/21st IEEE Conference on
Mass Storage Systems & Technologies
The Inn and Conference Center
University of Maryland University College
Adelphi MD USA
April 13-16, 2004

Remote Mirroring over Low-bandwidth WAN with iSCSI

Ming Zhang (mingz@ele.uri.edu)
University of Rhode Island

NASA/IEEE MSST2004
12th NASA Goddard/21st IEEE Conference on Mass Storage Systems and Technologies The Inn and Conference Center
University of Maryland University College Adelphi MD USA
April 13-16, 2004

Why remote mirroring?

- Data is more valuable than others
  - Hardware: Mouse (~Free); PC (<$1,000); 1TB Storage (~$10K);
    - Bill Gates (03/29/04) - “hardware to be nearly free in 10 years”;
  - Human: Computer Architect (~$100K/year);
  - Data (almost priceless);
- Data error and loss are difficult to “fix” (recovery).
- Compliance Requirement.
  - HIPAA, E-mail archival;

1-Hour Unavailable Data Cost

Brokerage operations	$6,450,000
Credit card authorization	$2,600,000
Ebay (1 outage 22 hours)	$225,000
Amazon.com	$180,000
Package shipping services	$150,000
Home shopping channel	$113,000
Catalog sales center	$90,000
Airline reservation center	$89,000
Cellular service activation	$41,000
On-line network fees	$25,000
ATM service fees	$14,000

(Sources: InternetWeek 4/3/2000 + Fibre Channel: A Comprehensive Introduction, R. Kembel 2000, p.8. "...based on a survey done by Contingency Planning Research.")

Why remote mirroring over low-bandwidth WAN?

- Cost: 40 Mbps ATM ($60,000/year in CA) vs. 3 Mbps Cable ($1,200/year in RI).
- Data: 50 TB/year (40 Mbps) vs. 3.7 TB/year (3 Mbps) (assume the average throughput can only achieve 1/3 of maximum value).
- Actual data changes in smal business is far less than that.

Why remote mirroring using iSCSI?

- iSCSI is a new open IETF standard.
  - SCSI over TCP/IP;
  - Open Standard vs. proprietary techniques.

Questions?

- Is remote mirroring over iSCSI a feasible solution?
- Which mirror scheme should be used?
  - Synchronous, semi-synchronous, or asynchronous?
- Is there any special design consideration needed for remote iSCSI target?

Experimental Methodology

- LAN and WAN;
- iSCSI and DCD (disk caching disk) enhanced iSCSI;
- Hardware configurations:
  - S-S, S-iL, S-iW, S-iDL, S-iDW;
- Workloads:
  - Benchmark tools: PostMark, IoMeter;
  - Traces: Financial-1, Financial-2, TPC-C;

DCD –Disk Caching Disk

- A patented architecture using a small log disk, referred to as cache-disk, as a secondary disk cache to optimize write performance (ISCA96, USENIX ATC99).

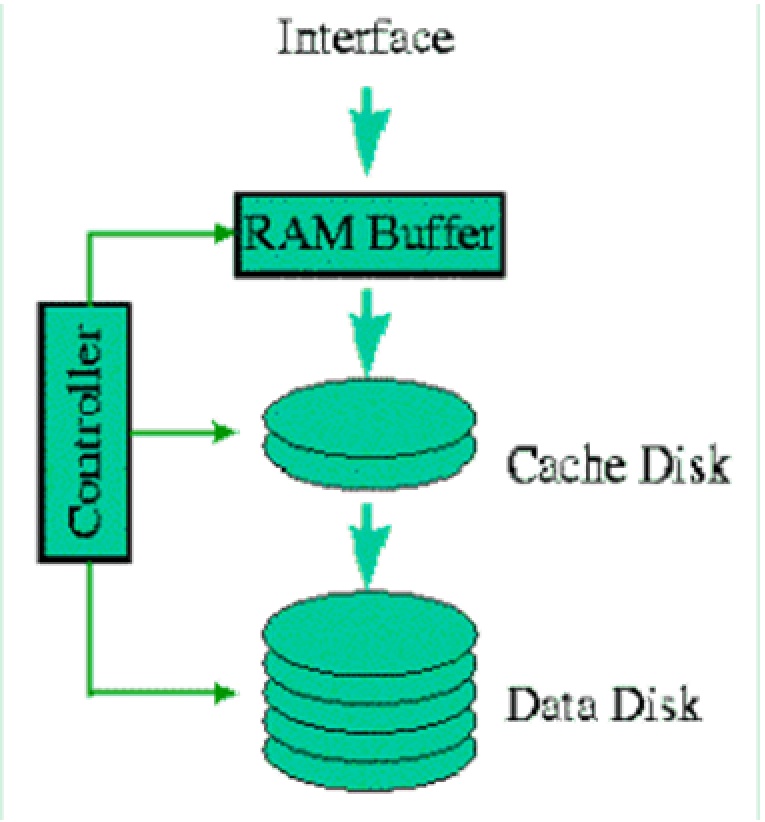

Environment Setting - LAN

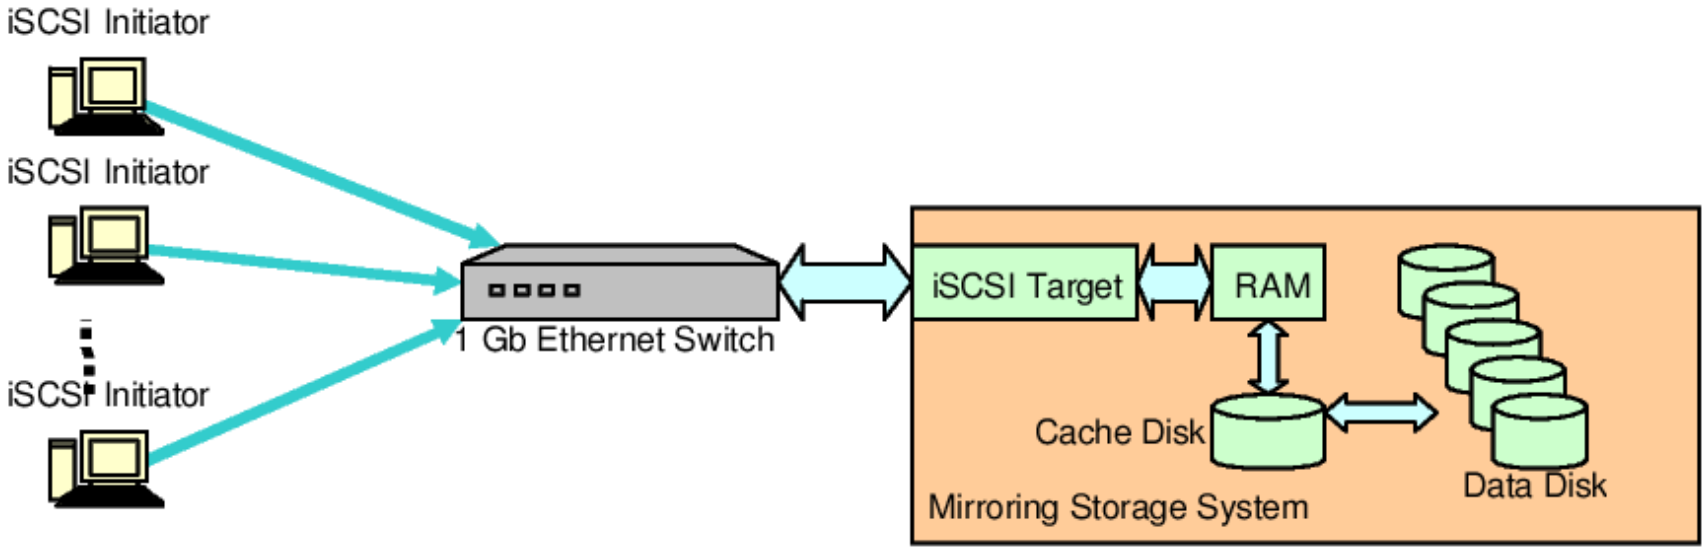

Environment Setting - WAN

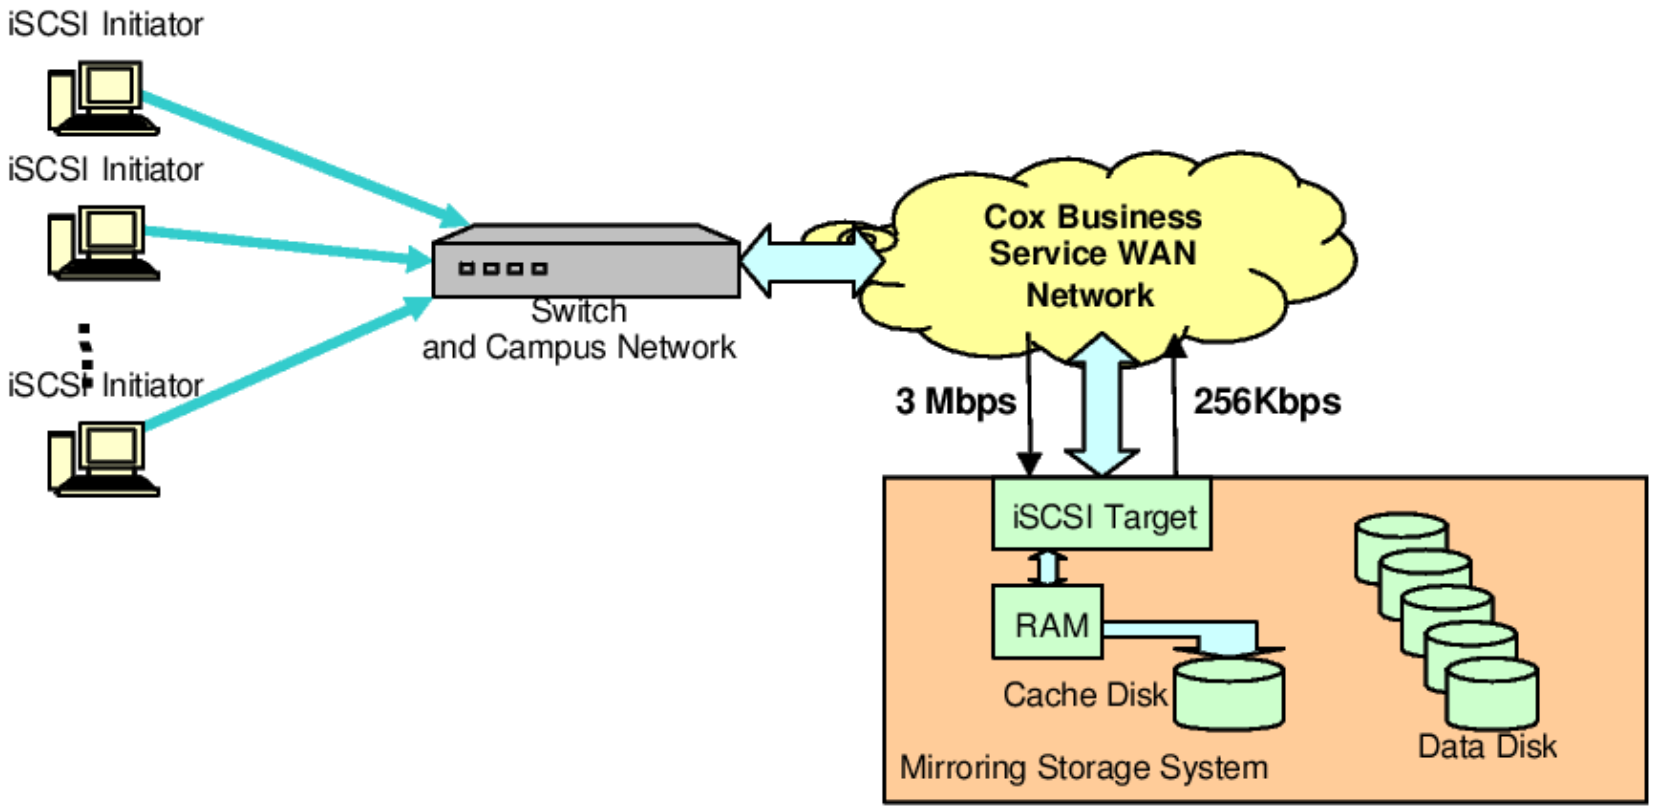

PostMark Results

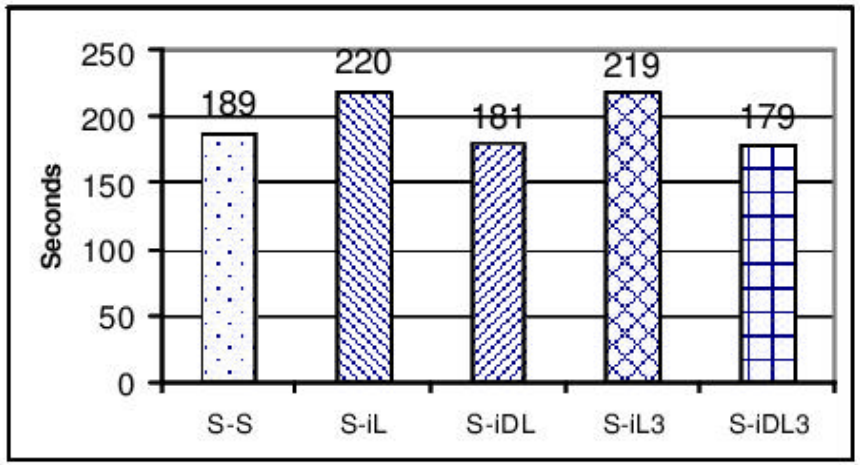

(a) 50% Read 50% Write over LAN

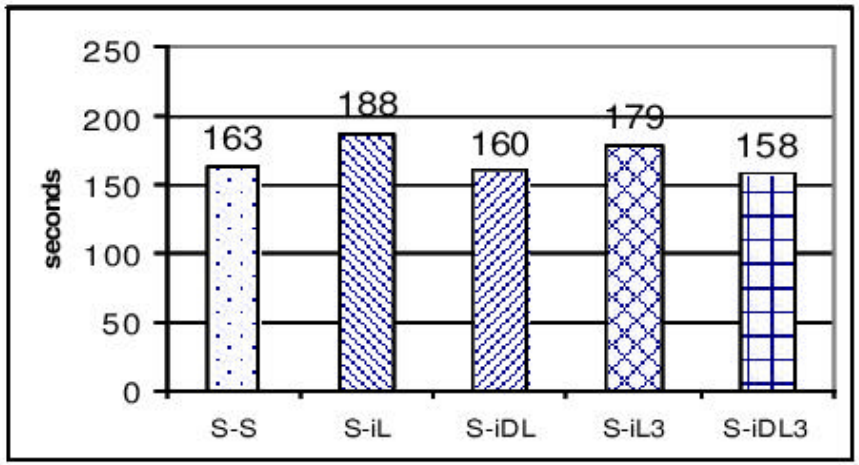

(b) 70% Read 30% Write over LAN

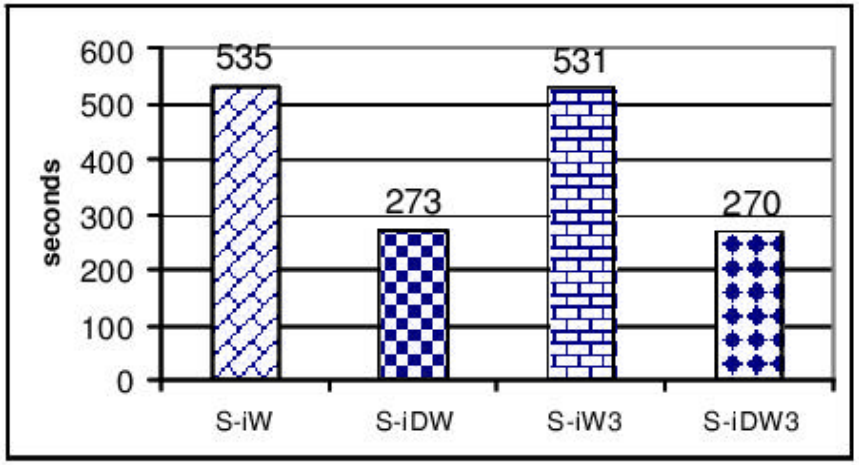

(c) 50% Read 50% Write over WAN

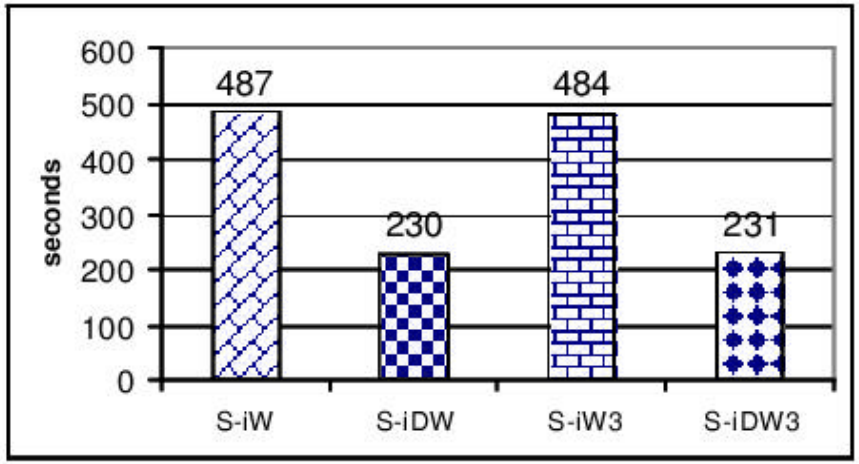

(d) 70% Read 30% Write over WAN

IoMeter Results

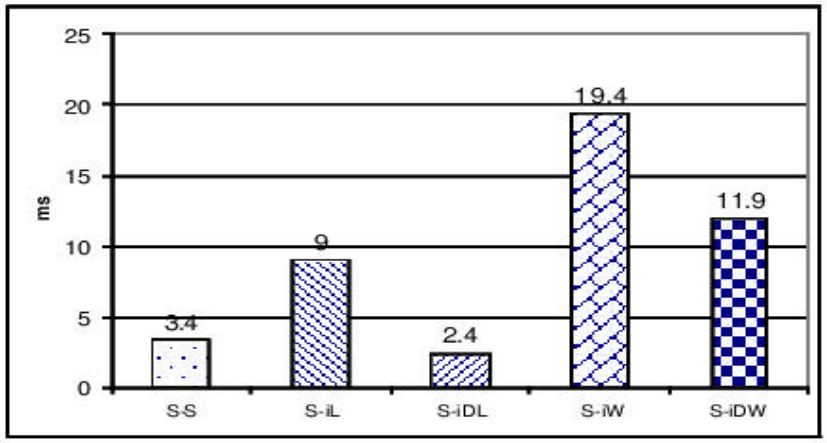

(a) Average write response time for random write only

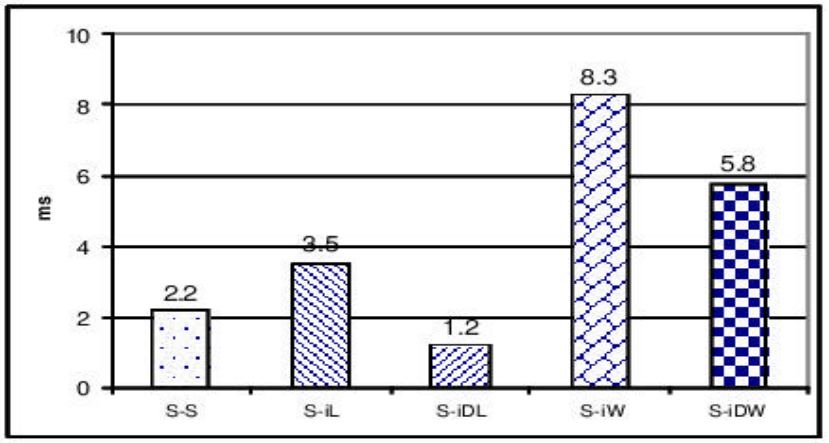

(b) Average write response time for 50% random write and 50% random read

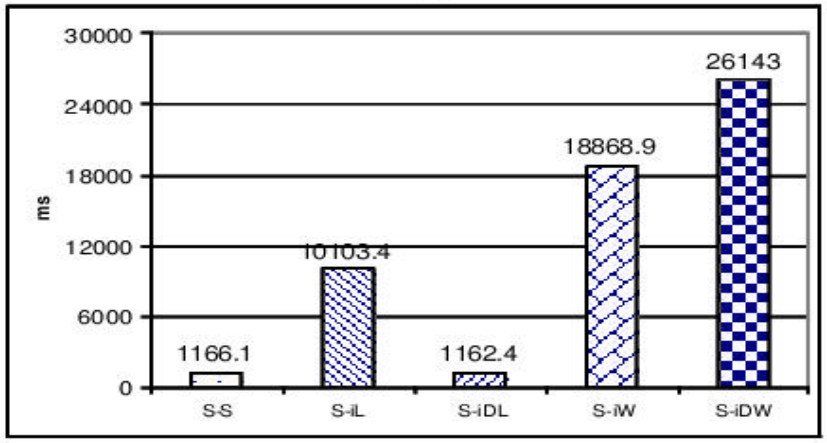

(c) Maximum write response time for random write only

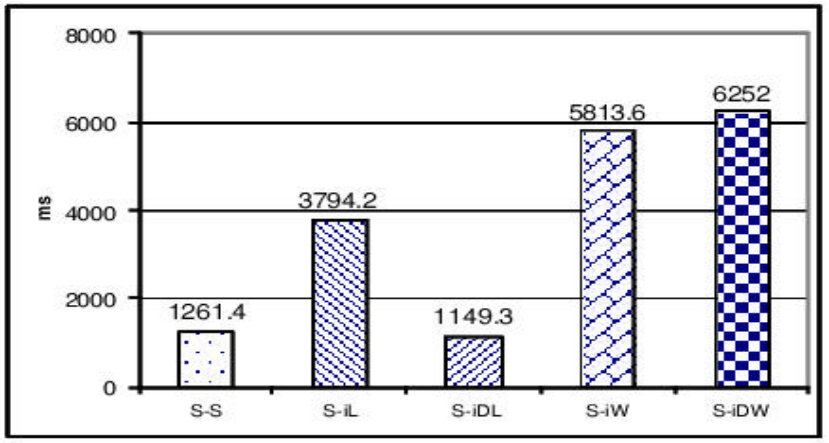

(d) Maximu m write response time for 50% random write and 50% random read

Trace Results –Average Response Time

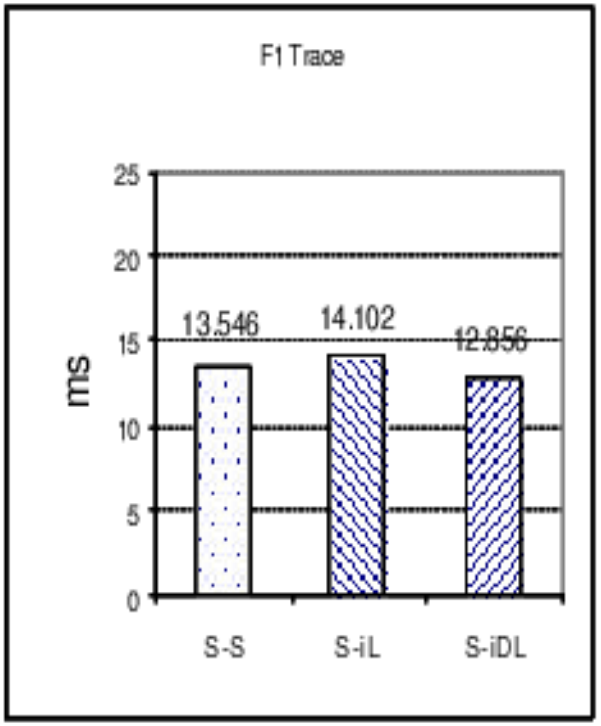

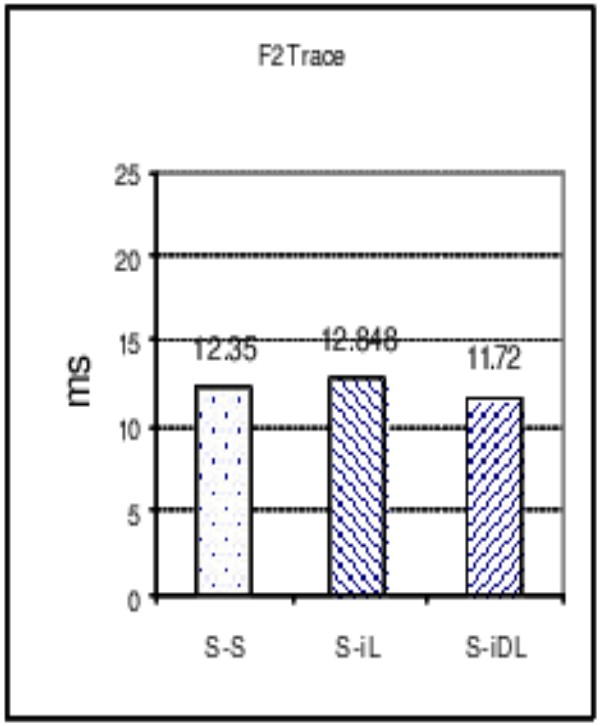

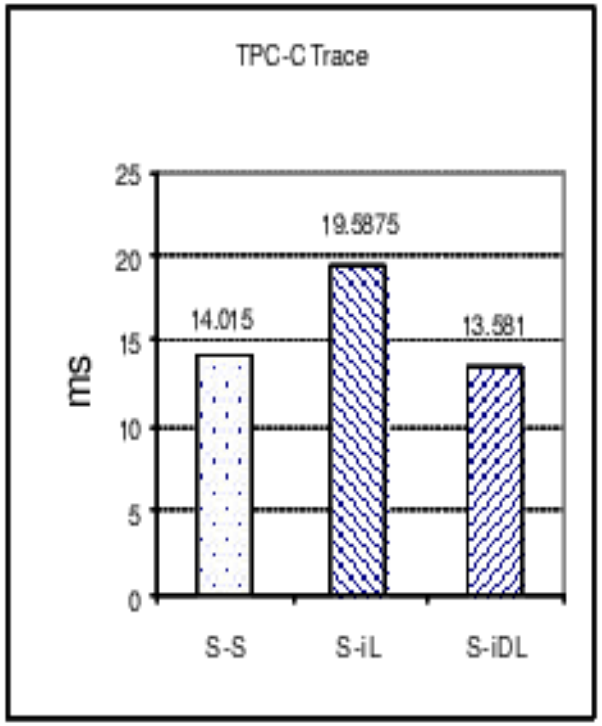

Trace Results –Minimum Response Time

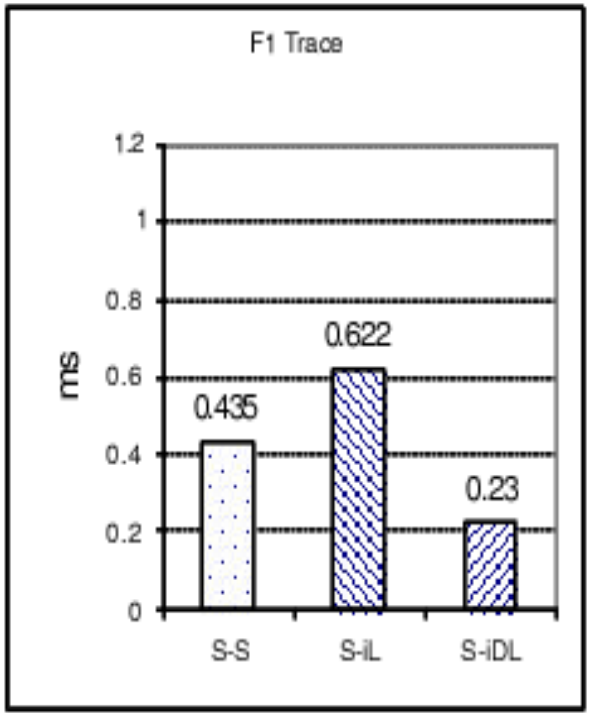

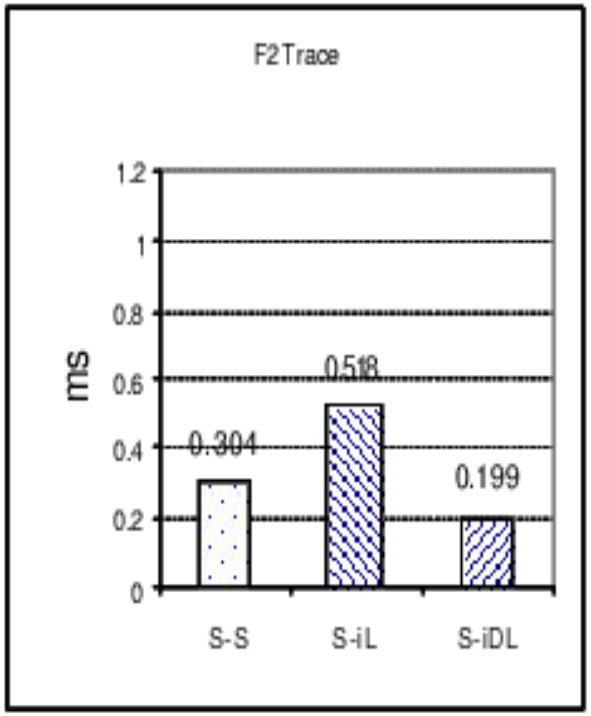

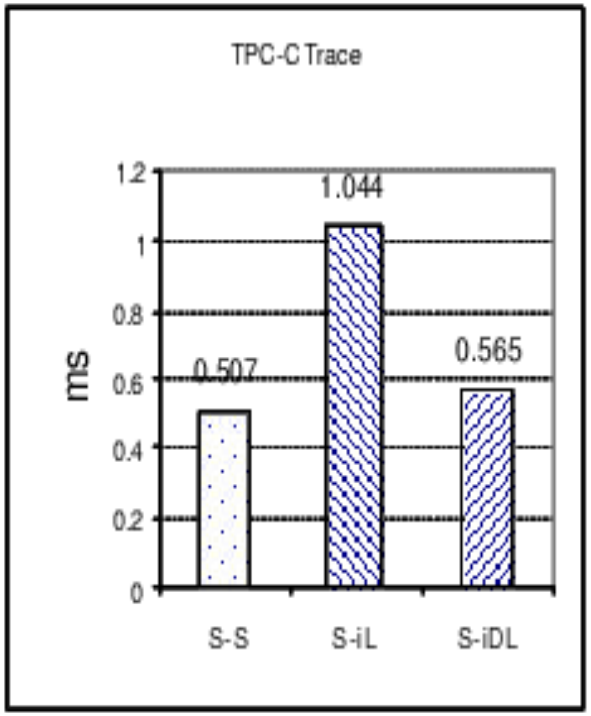

Trace Results –Maximum Response Time

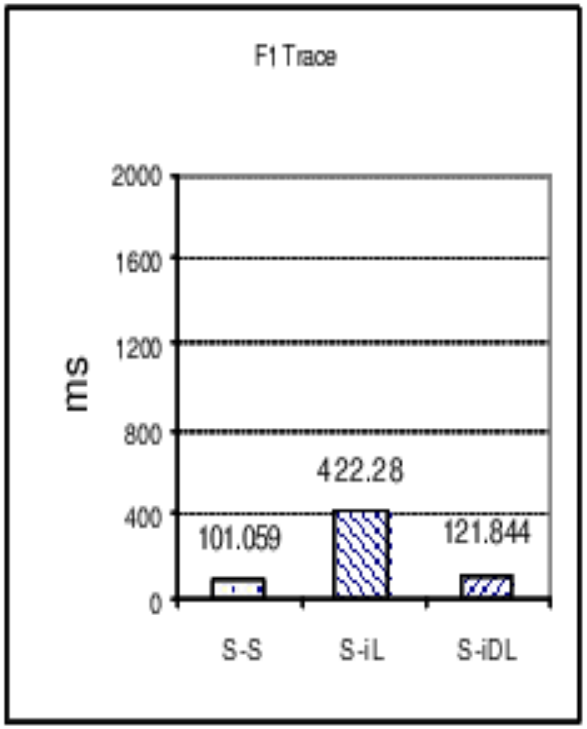

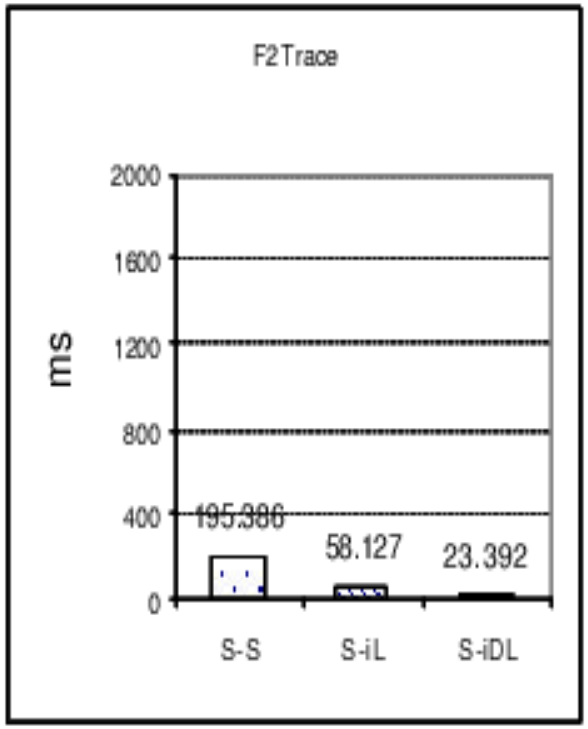

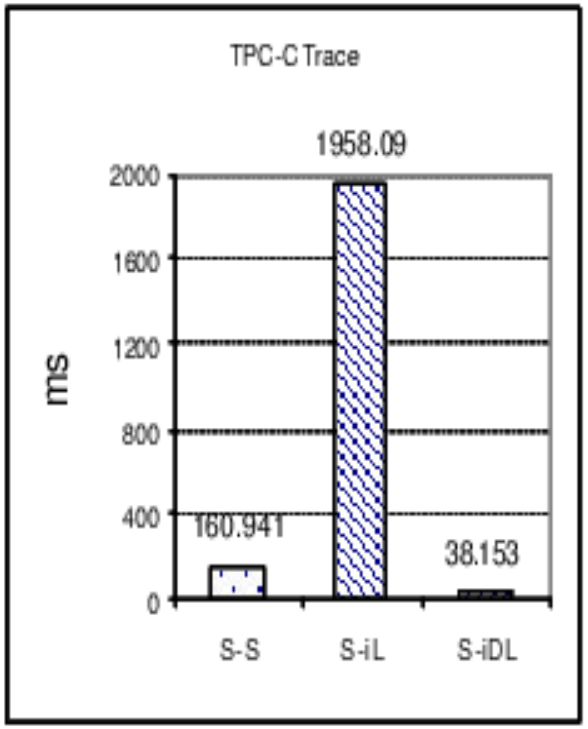

Trace Results –Response Time Plot

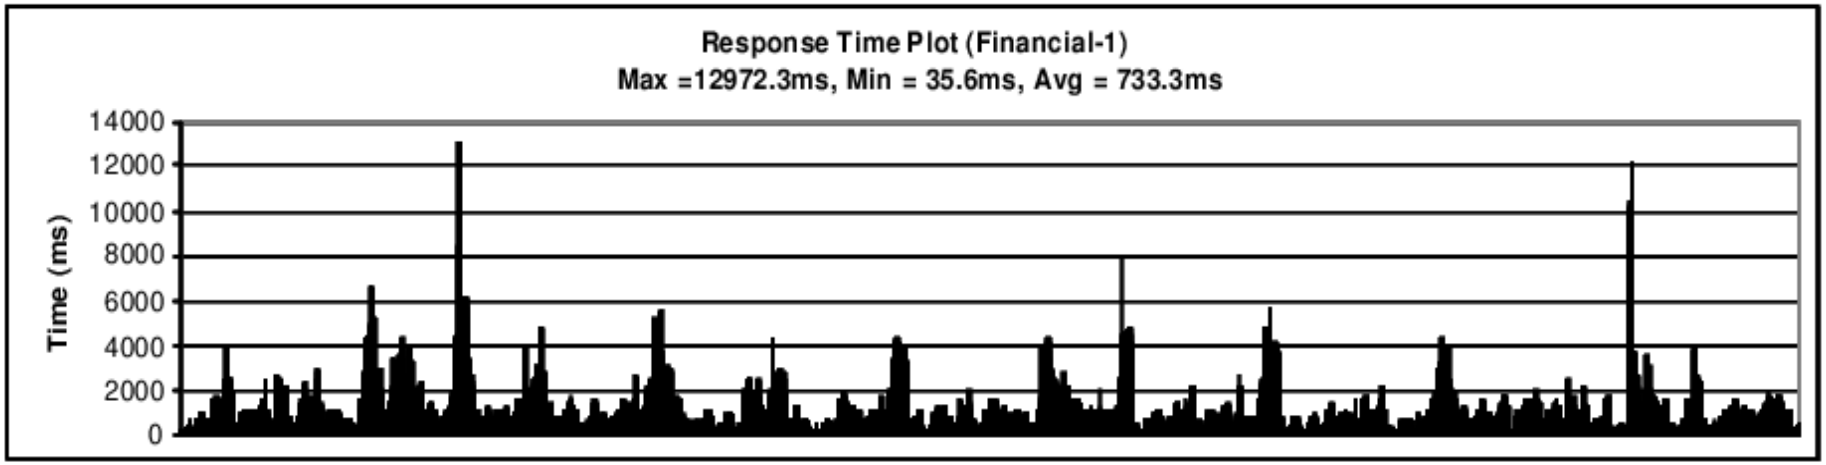

(a) Financial-1

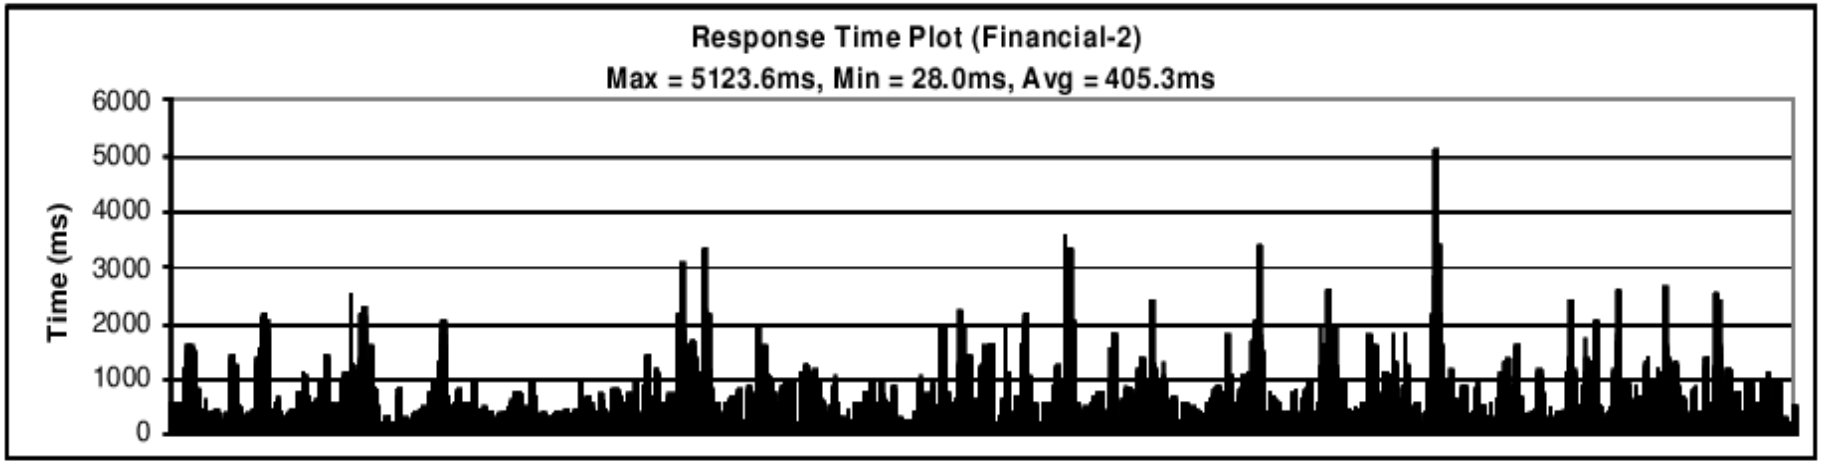

(b) Financial-2

Conclusion

- Remote mirroring over low-bandwidth WAN with iSCSI is a cost-effective and feasible solution;
- Asynchronous mirroring is needed for low-bandwidth WAN.
- Aggressive caching at iSCSI target side can greatly improve the performance.