

Accelerating Data Through the Fabric

*The Benefits of a Parallelized, Highly Scalable
Storage Architecture in Performance Driven and
Multi-Protocol HPC Environments*

Brad Winett
VP, Business Development
bwinett@datadirectnet.com

April 2002

High-Performance SAN Requirements

- High Throughput and I/O access with linear performance scalability
- Simplicity - easy management with fast deployment and reconfiguration
- Open systems architecture, supports O/S independent workstations to clusters to monolithic supercomputers
- Reduce raw disks and hosts (as much as possible) to commodities
- Multi-Protocol Support (FC, GigE...)

SAN Scalability Issues

Is anyone in the audience trying to scale a SAN with a shared file system?

- If not, you probably aren't aware of what will really limit your performance!

If the world was perfect, just adding RAID's and switches would linearly scale...

- Unfortunately, the world is not perfect.

WHY?

Because slow LUNs, host-based striping, non-deterministic I/O patterns, and fabric contention and switching latencies are evil and pernicious %@#& that conspire against you as you try to scale!

SAN Scalability Answers


The solution is an architecture that:

- Supports very high performance “PowerLUNs” (>>1GB/sec) as a single target
- Supports true parallel concurrent access to PowerLUNs
- Distributed, parallel processing RAID
 - Equal performance Read & Write
 - Excellent multi-initiator performance (non-determinism)
 - No performance degradation in crippled or rebuild modes
 - True duplex operation
 - High availability and fault tolerance
- Multi-Protocol Support (FC, GigE, OBD...)
- Cost Effective

The Test


TEST #1 No Acceleration

- Essentially 8 Striped RAID Arrays
- Eight Striped (8+1) LUNs



TEST #3 Full Acceleration

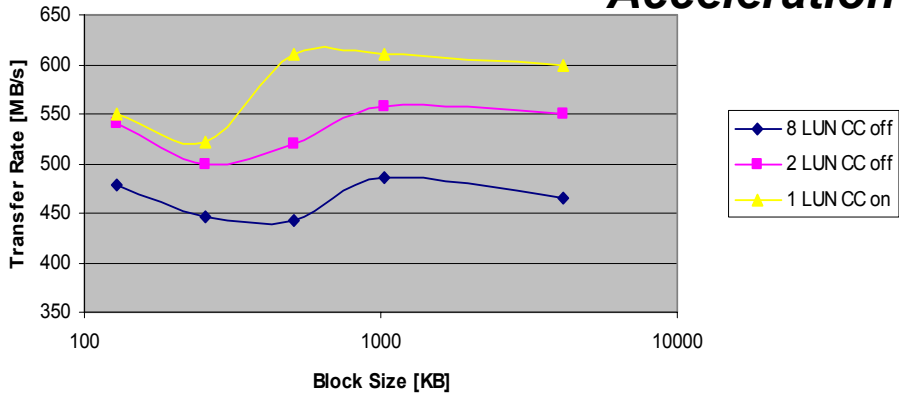
- One (64+8) PowerLUN
- No Striping



Test Comparisons

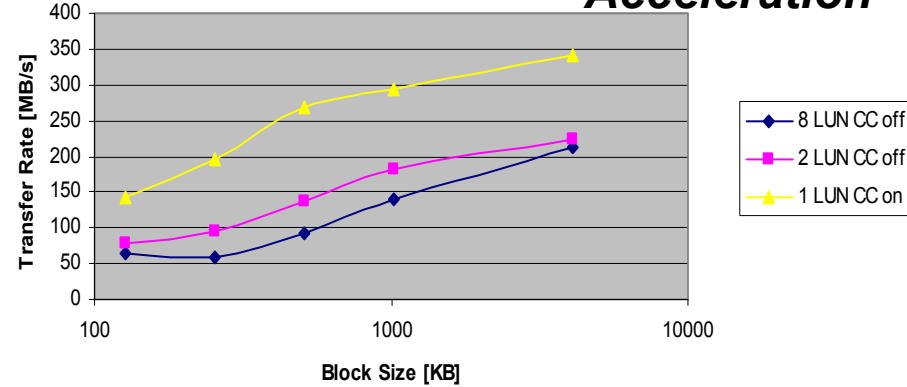
100% Sequential Reads

Up to 41% Acceleration




100% Random Reads

Up to 53% Acceleration




100% Sequential Writes

Up to 35% Acceleration



100% Random Writes

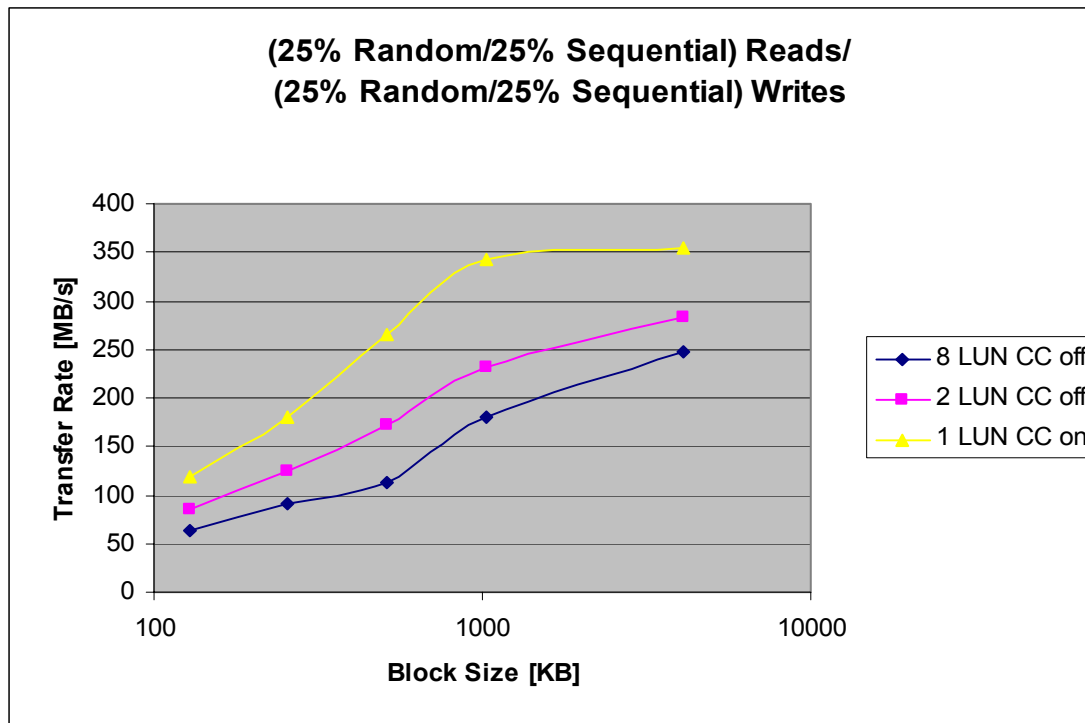
Up to 9% Acceleration



April 2002

	TEST #1 8 Striped LUNs
	TEST #2 2 Striped PowerLUNs
	TEST #3 1 PowerLUN

Test Comparisons



**Up to 40% Large-Block
Data Acceleration**

- TEST #1 8 Striped LUNs
- TEST #2 2 Striped LUNs
- TEST #3 1 LUN

What does it Cost?

- Fibre-Channel 2Gb
 - *100TB @ 10GByte/sec* <\$3M
 - *1PB @ 100GByte/sec* <\$30M
- Gigabit Ethernet
 - *50TB @ 5GByte/sec* <\$1.5M

Who Can Say it Works?

- **US Army Fort Schafter**
- **US Army ARL**
- **DKRZ Germany**
- **Lawrence Livermore**
(12 @ 90TB)
- **NASA Goddard**
(6 @ 200TB)
- **NCAR**
- **NCSA**
(7 @ 60TB)
- **NEC**
- **NOAA FSL**
- **Northrop Grumman**
- **Sandia Labs**
(8 @ 40TB)
- **White Sands Missile Range**

The Accelerators



Silicon Storage Appliances

*Only from
DataDirect
Networks!*