Synchronous Cache Mirroring Technology and Active/Active Controllers

Lenovo ThinkSystem DS2200, DS4200, DS6200

June 2017

David Vestal, WW Product Marketing

Lenovo.com/systems

Introduction

Lenovo's ThinkSystem DS2200, DS4200, and DS6200 SAN Storage Arrays deliver industry-leading cache performance through cache mirroring. Any time write data is sent by the host to the first RAID controller in an active-active pair, the data is instantaneously and simultaneously mirrored to the cache of the second controller.

This "cache coherency" (or, alternatively, "cache mirroring") process is performed without any additional involvement from the processors on either RAID controller. Plus, controller cache is synchronized using a broadcast write across a dedicated, high-performance interconnect. This means that the entire bandwidth of the backend interconnect is available for moving data, unlike slower, more traditional implementations from other storage vendors.

Additionally, the memory bus is used only once to write the data. In traditional implementations, the data must traverse the bus at least twice: once for the data movement, and a second time to mirror the data movement. With Lenovo's cache mirroring, the processors are not involved in this mirroring process. As a result, write latencies are significantly reduced.

Compared to older, conventional implementations, cache mirroring offers a performance improvement of approximately 50 percent!

Cache Coherency: The Old Approach

Conventional cache coherency (a.k.a. cache mirroring) between dual active RAID controllers is accomplished by mirroring host write data using host or disk channels. This approach introduces channel I/O command overhead on RAID

controllers as well as on the host and target. This method also requires additional memory bandwidth on the primary controller, since write data must be read back out of cache during mirror operations. The block diagram in Figure 1 below helps illustrate the conventional cache mirroring method.

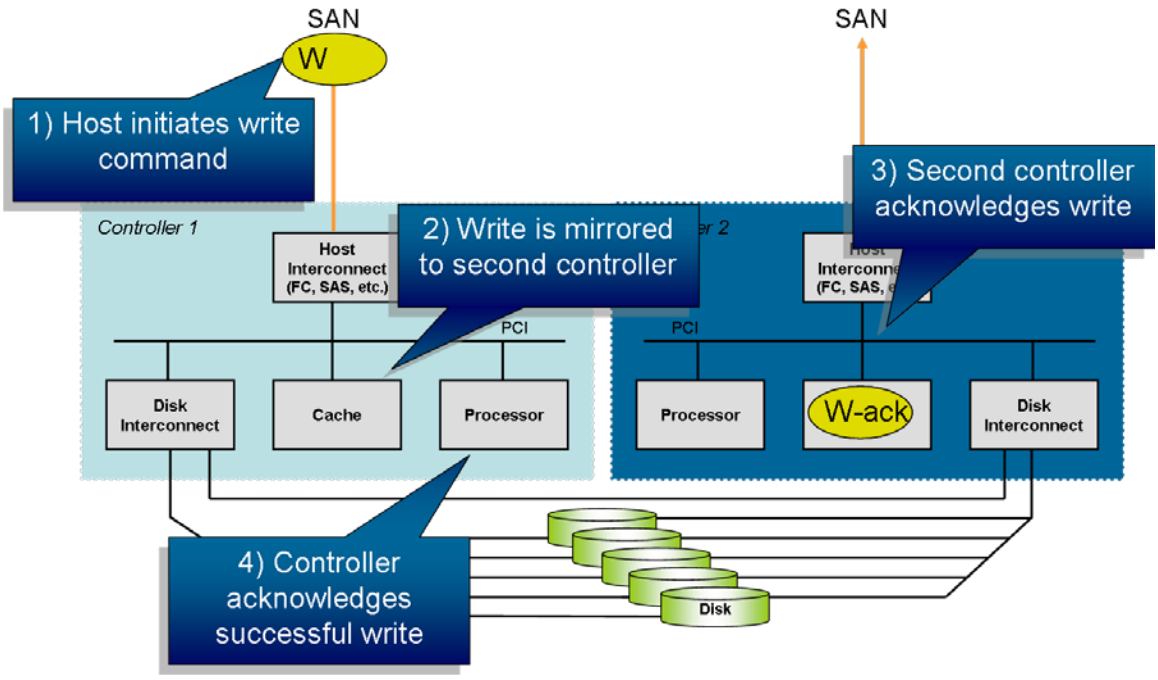

Figure 1 - Block Diagram Illustrating Conventional Cache Coherency

The High-Performance Approach

With Lenovo's cache mirroring, the cache memory controller eliminates the primary latency overhead and reduces bandwidth requirements on the primary cache by automatically "broadcasting" the write data to the other controller's cache. Plus, the broadcast does not require software intervention for data

input/output or to support SAS target mode or simultaneous target/initiator Fibre Channel mode on host channels. This significantly minimizes command overhead on both controllers. And, as stated above, there is no additional memory bandwidth required on the primary controller to mirror the write data.

The overall result is an exceptionally high bandwidth, low latency, active-active write-back cache capability, which approaches the performance of a dual independent cache mode operation. Figure 2 helps illustrate the cache mirroring method.

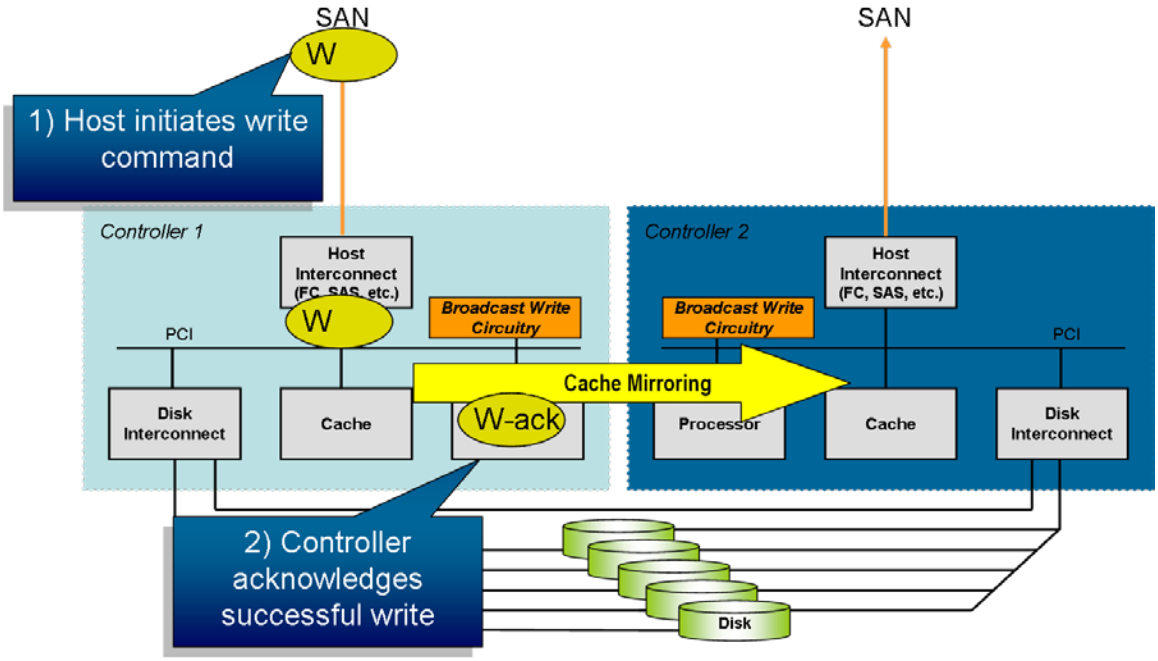

Figure 2 – Cache mirroring