Data Storage Process

Virtual Array

The term “Virtual Array” refers to the way the array manages the disks as a pool of data storage blocks instead of whole physical disks. Like other virtualization within computer systems, this virtualization greatly simplifies the management of the array. Internally, the array uses sophisticated data structures to manage the logical-to-physical address translation. These data structures, often referred to as the “maps”, are key to the operation of the array. See Figure 12.

Administrators’ manage the capacity of the array using Redundancy Groups and LUNs. Each disk belongs to a predefined Redundancy Group, and a LUN is created from the capacity of a Redundancy Group. This is similar to traditional arrays. The virtualization eliminates the need to manage the lower level details. Redundancy Groups can be constructed from any number or capacity of supported disks. Any number of disks can be added to a Redundancy Group at any time. LUNs can be of any size up to the available capacity of a RAID Group, or created and deleted without the knowledge of the underlying physical disk layout. The VA 7100 supports up to 128 LUNs; the VA 7400/7410 support up to 1024 LUNs.

Figure 12 Virtual Data Storage

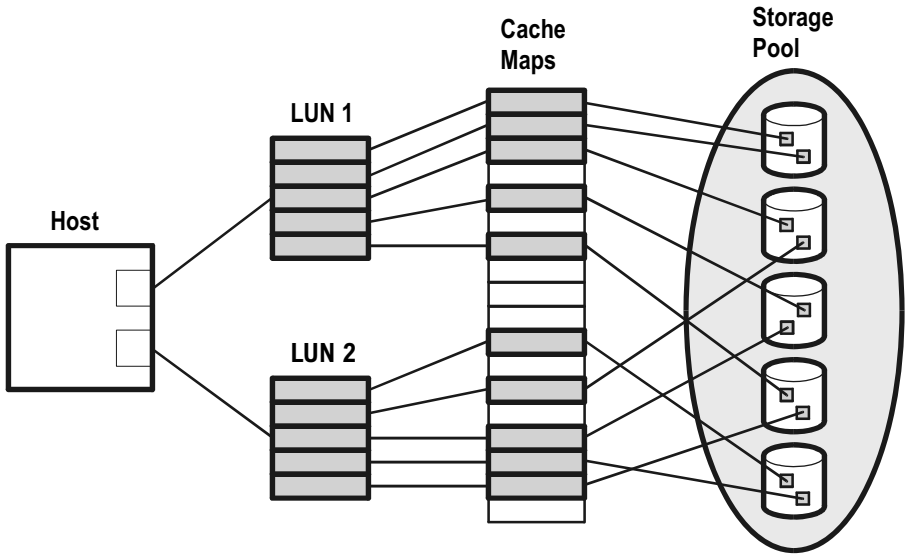

Redundancy Groups

Array physical capacity is divided into Redundancy Groups. A Redundancy Group (RG) can be thought of as an independent array. Each RG has its own set of disks, active hot spare, and controller. LUNs are created from capacity within a single RG. LUNs can be accessed simultaneously through either controller.

Multiple redundancy groups provide the following benefits:

- Fault isolation. Because each redundancy group has its own resources, a disk failure in one RG will not impact the other RG. This effectively increases the data availability of the array.
- Performance management. Applications can be assigned to different RGs, thus isolating their performance impact on each other.
- Greater configurability. Each RG can be constructed from different classes of disks. As an example, one RG could be constructed from a few, small,

high-performance disks, and the other RG from large, slower, high-capacity disks.

The VA 7100 and VA 7400/7410 differ in their implementation of redundancy groups.

VA 7100/7110 Redundancy Group

The VA 7100 and VA 7110 each have one redundancy group (RG1). See Figure 13 and Figure 14. All the disks in the array belong to RG1. LUNs created from RG1 are available through both controllers (in a dual controller configuration).

There are two internal fibre channel loops, one from each controller. The Fibre channel disks are dual ported; each fibre channel port is connected to a different controller. The controllers are connected via an internal high-performance bus, which allows the LUNs to be accessed through both controllers, and for loop or disk failover communication.

Figure 13 VA 7100 Redundancy Group

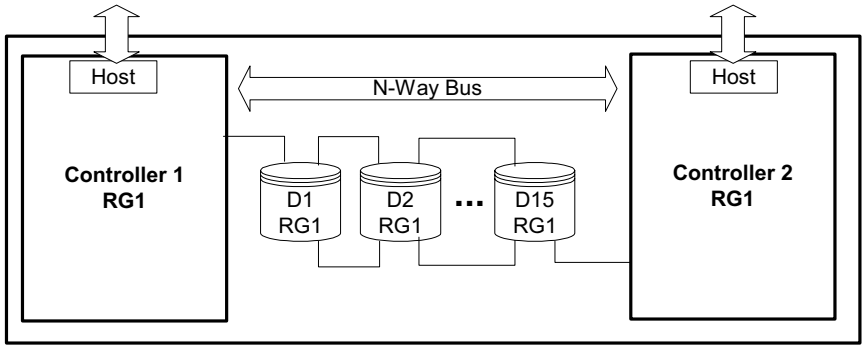

Figure 14 VA 7110 Redundancy Group

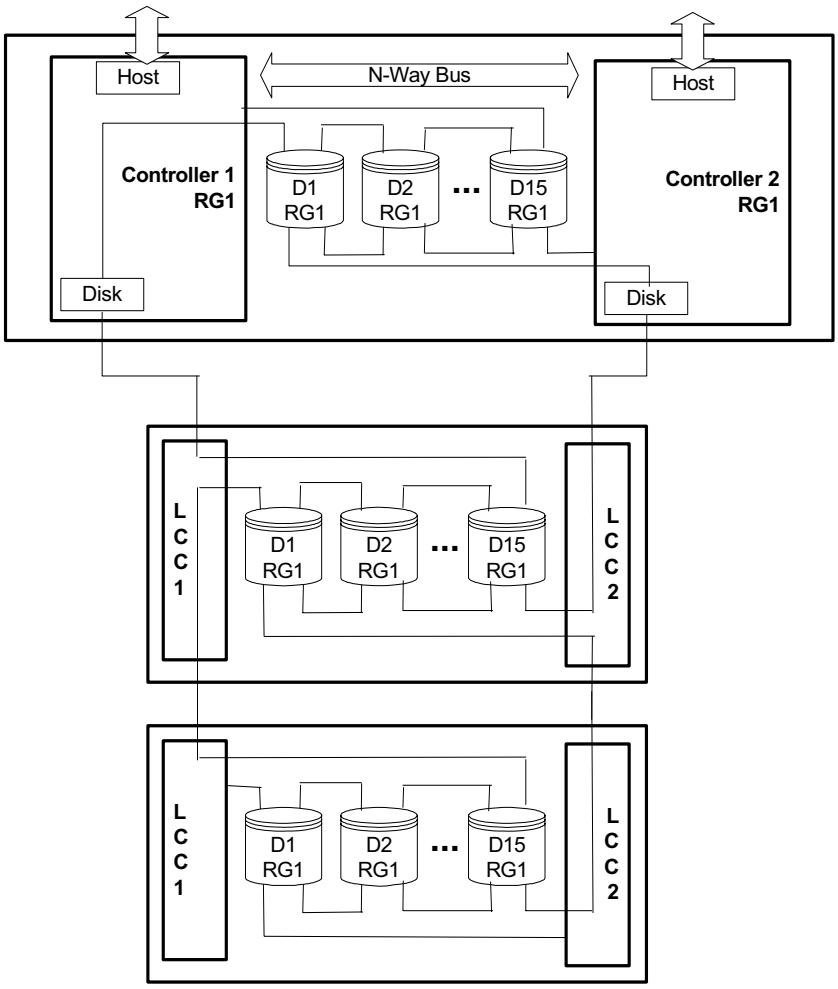

VA 7400/7410 Redundancy Groups

The VA 7400 and VA 7410 have two redundancy groups (RG1 and RG2). See Figure 15 and Figure 16.

- Controller 1 manages Redundancy Group 1 (RG1), which consists of all disks in odd numbered slots (D1, D3, D5, D7, D9, D11, D13, D15) in the controller enclosure, and in all disk enclosures (JA0-JA5).
- Controller 2 manages Redundancy Group 2 (RG2), which consists of all disks in even numbered slots (D2, D4, D6, D8, D10, D12, D14) in the controller enclosure, and in all disk enclosures (JA0-JA5).

On the VA 7410, Redundancy Group are independent of both back-end FC loops. Management of the redundancy group disks is independent of which disk enclosure LCC the array controller is connected to. For example, array controller 1 can be connected to LCC 1 or LCC 2 and it will still manage the disks in the odd numbered slots.

The array controllers are connected via an internal N-Way bus, which used for controller-to-controller communication and loop failover.

Figure 15 VA 7400 Redundancy Groups

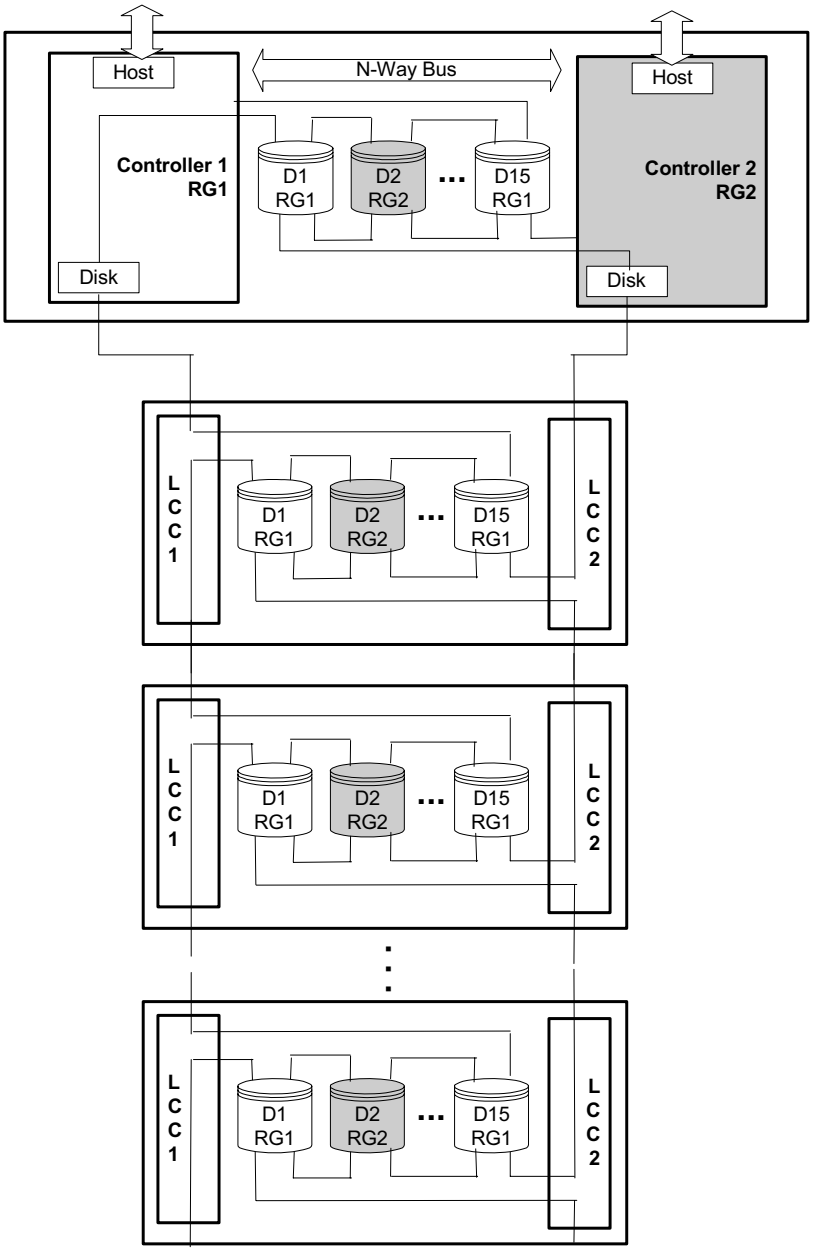

Figure 16 VA 7410 Redundancy Groups

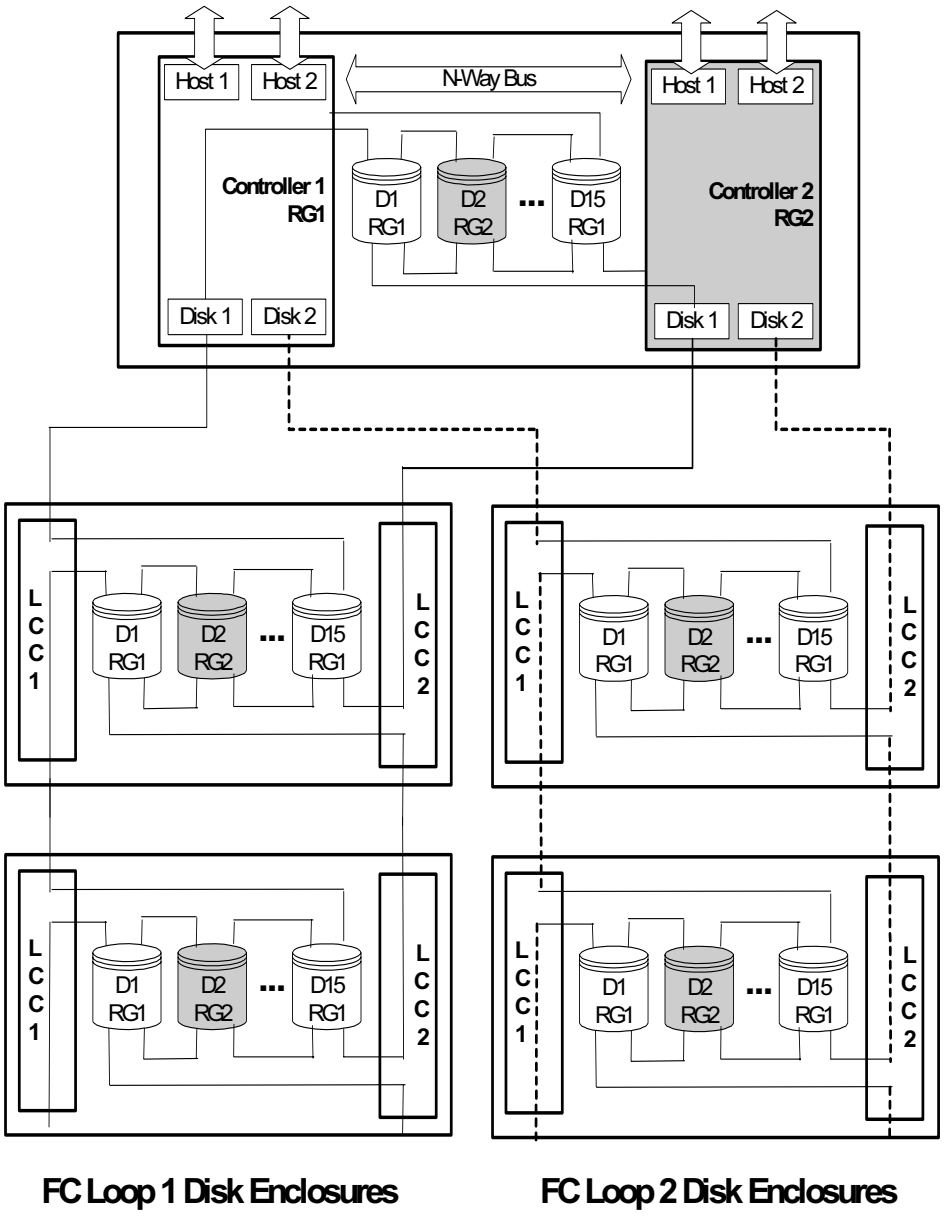

Performance Path

The performance path is the most direct path from the host to the data in the array. It is specified by two separate device files that direct the data either through Controller 1 or through Controller 2. The performance path is always the faster path in terms of data transfer rate.

Because the array has two active controllers, the host will typically have two paths to data, as shown in Figure 17.

- The primary path is through the controller which owns the LUN being accessed. That is, the controller that manages the RG the LUN belongs to. On the VA 7400 and 7410 each LUN is assigned to RG1 or RG2, managed by controller 1 and controller 2 respectively. When accessing data on a LUN, the host should send I/Os to the controller which owns the LUN.
- The secondary path is through the controller which does not own the LUN being accessed. In this situation, the non-owning controller must use the internal N-Way bus to send the I/O to the controller that owns the LUN. Whenever the secondary path is used, I/O performance is impacted due to the inter-controller communication required.

System and SAN configuration with the knowledge of the performance path is a technique to maximize the array performance. For normal workloads this provides very little performance improvements, but for benchmarking and highly utilized arrays, this can provide modest performance gains. The biggest gains can be found with the VA 7100/7400, improvements with the VA 7110/7410 have reduced the performance gained through performance path management.

The use of load balancing software in normal workloads, such as HP AutoPath, can, in many cases, offset any gains in performance by managing the configuration of the performance path.

VA 7100/7110 Performance Path

In the VA 7100, the performance path is always specified by the device file for Controller 1. Because the VA 7100 has only one redundancy group, and the secondary controller is recommended only for failover, the primary controller is always the most direct path to the data. If Controller 1 fails, the host should use the secondary path to Controller 2.

VA 7400/7410 Performance Path

The following example illustrates how the performance path is used in a VA 7400/7410:

Assume LUN 4 is part of Redundancy Group 2 under Controller 2. An HP-UX host has two device files that have two separate paths to LUN 4: The primary device file that addresses Controller 2, and the secondary device file that addresses Controller 1. The performance path uses the primary device file, because Controller 2 owns LUN 2. The non-performance path uses the secondary device file. If the secondary device file is used, data flows through Controller 1, across the N-way bus to Controller 2, and then to LUN 2 and its associated disk drives.

Figure 17 Data Paths on the VA 7400/7410