



CacheIO High Availability Deployment

Table of Contents

Solution Overview	3
Product Overview	4
Before CacheIO Deployment	5
HA Deployment with a Single CacheIO Appliance	6
HA Deployment with Redundant CacheIO Appliances	7
Configuration at Orange Silicon Valley	9

Table of Figures

Figure 1: CacheIO Product Architecture	4
Figure 2: Before CacheIO Physical Configuration	5
Figure 3: Before CacheIO Logical Configuration	5
Figure 4: Single CacheIO Physical Configuration	6
Figure 5: Single CacheIO Logical Configuration	6
Figure 6: Redundant CacheIO Physical Configuration.....	7
Figure 7: Server Area Network and Storage Area Network	7
Figure 8: Active Active Application Access.....	8
Figure 9: Active Active Storage.....	8
Figure 10: OSV Physical Configuration.....	9
Figure 11: OSV Logical Configuration.....	9

Solution Overview

Performance issues of business critical applications can cause revenue loss, and decline in customer satisfaction, employee productivity, or brand reputation. Increasingly data center operators find themselves struggling with an enterprise bottleneck not of their own making. Servers' hard disk access times have not kept up with the increasing speed of their CPUs, and the resulting lags have become the limiting factor in performance critical applications, particularly those involving software as a service and cloud computing. As these applications become even more popular, the bottleneck is about to get worse.

It's well understood in the IT industry that application performance is mostly limited by storage bottleneck and therefore cannot be addressed by faster servers or networks. It is also widely accepted that Flash, with its rapidly declining cost per GB, is emerging as the ideal storage media to address this performance bottleneck. Storage vendors vary widely in their architectural approaches to leveraging flash media:

Server-side Flash: Server-side Flash can be an effective point solution for single instance applications (not clustered), especially when it's implemented as a cache to the SAN storage. However, this approach does not scale to cache all active data, making its latency and throughput profile more similar to that of spinning disks than flash memory. Caching software also competes with the application for CPU and memory resources.

Array-side Flash: Most storage vendors offer SSDs in their arrays as a high performance tier. However performance gain with this approach is severely limited by the storage controllers in the arrays, which are also burdened with performing storage management features such as disk rebuild, snapshot, clone, and replication.

All-Flash Array: All-flash arrays require a complete replacement of existing disk arrays, incurring high capital and data migration costs, as well as lacking in mature storage management features that customers have depended on for years.

Network Cache Appliance: Network cache appliances leverage your current infrastructure and management practices. They scale flash capacity and performance beyond that of server-side and array-side flash without interfering with server and storage controller performance.

At CacheIO we believe that network cache appliance is the best storage architecture to decouple performance from capacity and achieve the best of both worlds. Once deployed as a "bump in the wire" performance accelerator, our network cache appliance is in a great position to deliver additional value added services such as synchronous and asynchronous remote replications for heterogeneous arrays and applications. Compared to server-side flash cache, our network cache appliance is a shared resource that is more scalable, more reliable, supports clustered applications, and most importantly allows customers, especially cloud service providers, to monetize performance by dynamically provisioning performance based on changing SLAs.

Product Overview

CacheIO removes performance as an issue by providing accelerated data paths between the servers and storage. It services nearly all of the active data from its 46TB flash cache, resulting in 100x lower latency and higher throughput. In another word, CacheIO makes applications screaming fast! And it does it in a way that is cost effective, operationally transparent, and highly available.

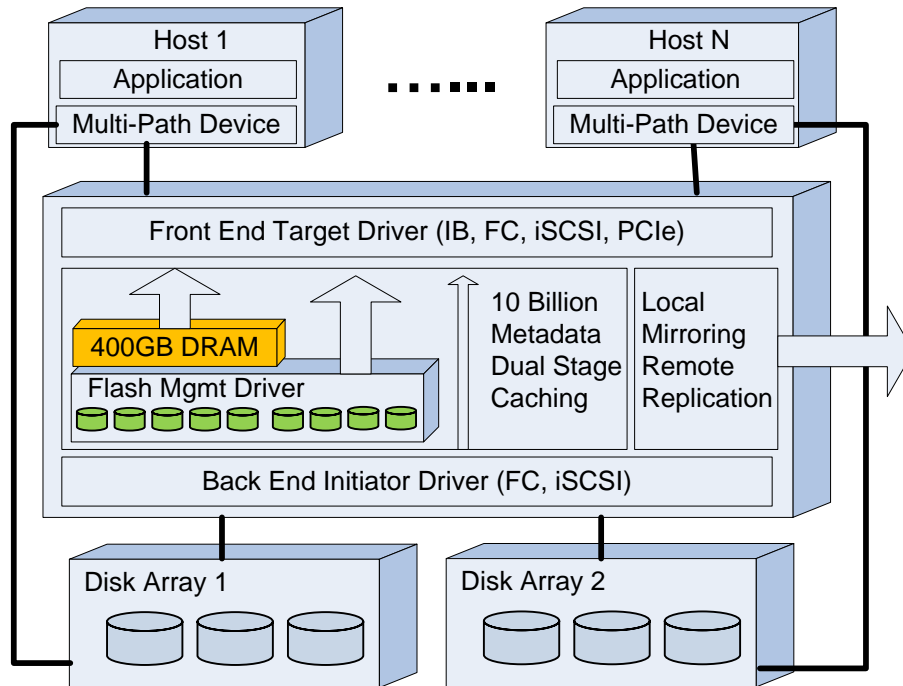


Figure 1: CacheIO Product Architecture

Cost Effective: Large and affordable flash cache capacity is the key to accelerating application performance since a mere 2% cache miss doubles application latency. CacheIO employs proprietary flash management software to eliminate copy on write garbage collection and increase the endurance of commodity flash by 40 times. Combining innovative software with commodity hardware, CacheIO provides the highest flash density, 46TB usable in 2RU, and by far the lowest cost in the industry at \$2/GB. With CacheIO removing performance as an issue, disk array utilization can increase from 30% to 80% and life cycle from 3 years to 6 years, drastically reducing both CAPEX and OPEX.

Operational Transparent: Implementing CacheIO network appliance requires no change to existing applications, servers, or storage. CacheIO can be slotted in, turned on to accelerate applications, and turned off if necessary, often without needing to stop the applications.

High Availability: CacheIO supports multi-path based high availability, You can either create an active-passive HA configuration with a single CacheIO appliance or an active-active configuration with redundant CacheIO appliances. The rest of the document focuses on these high availability deployment options.

Before CacheIO Deployment

Figure 2 shows a typical high availability configuration prior to CacheIO deployment. The host has two HBAs, each with two Fibre Channel ports. The Fibre Channel ports connect to redundant switches, through which the host connects to redundant controllers on the storage array.

Figure 3 illustrates the logical configuration after Fibre Channel zoning. The application is configured to use multi-path devices for transparent high availability and load balancing. There are four paths to each LUN, protecting against single HBA, switch, or controller failure. The four paths can be active-active, active-passive, or ALUA (active-active asymmetric), depending on the high availability capability of the storage array.

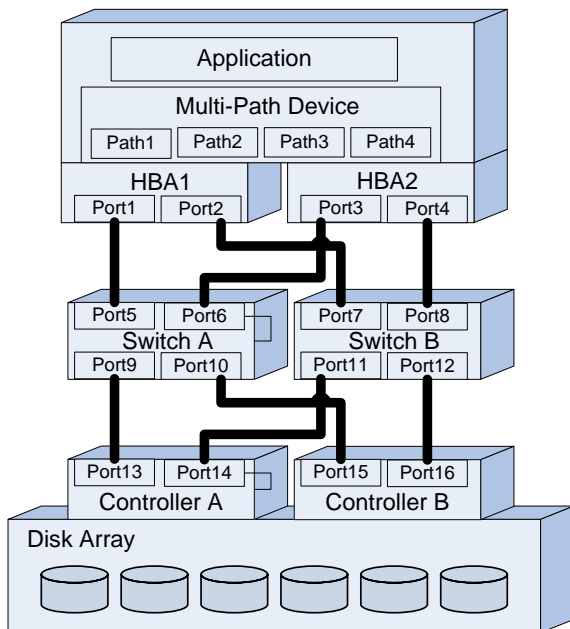


Figure 2: Before CacheIO Physical Configuration

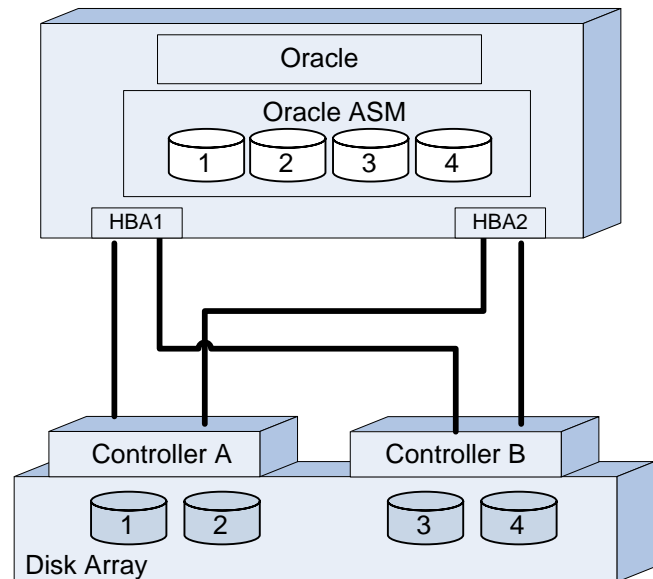


Figure 3: Before CacheIO Logical Configuration

HA Deployment with a Single CacheIO Appliance

Figure 4-5 show a high availability deployment with a single CacheIO appliance. The CacheIO appliance is inserted between the server and storage to accelerate two of the four available paths to each LUN. During the insertion process, the application accesses the storage through the two direct data paths. After the CacheIO appliance is inserted and ready to go, the multi-path device is switched over to the two CacheIO accelerated paths without interrupting application access. The CacheIO appliance can also be removed non-disruptively by switching the multi-path device over back to the direct paths.

Figure 4 illustrates the physical configuration. The host has two HBAs, each with two Fibre Channel ports. The Fibre Channel ports connect to redundant switches, through which the host connects to the CacheIO appliance and the storage array. The CacheIO appliance has two Fibre Channel ports connecting to the storage array via redundant switches.

Figure 5 illustrates the active-passive logical configuration during normal operation. On the host there are four paths to each LUN – two CacheIO accelerated paths (active) and two direct paths (passive). On the CacheIO appliance there are two paths to each LUN via the redundant controllers on the storage array. The two paths can be active-active, active-passive, or ALUA depending on the HA capability of the storage array.

This configuration protects application access to the storage against single HBA, switch, controller or CacheIO failure.

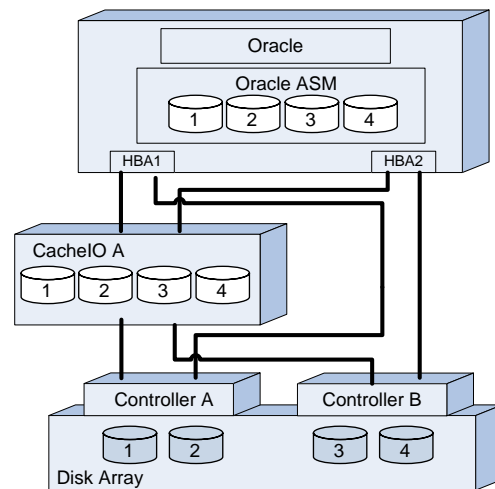



Figure 4: Single CacheIO Physical Configuration

Figure 5: Single CacheIO Logical Configuration

HA Deployment with Redundant CacheIO Appliances

Once a single CacheIO appliance is in place, a second CacheIO appliance can be inserted non-disruptively to establish an active-active configuration, as illustrated in Figure 6. The host has two HBAs, each with two Fibre Channel ports. The Fibre Channels ports connect to redundant switches, through which the host connects to both CacheIO appliances. Each CacheIO appliance has two Fibre Channel ports connecting to the storage array via redundant switches.




Figure 6: Redundant CacheIO Physical Configuration




Figure 7: Server Area Network and Storage Area Network

Alternatively CacheIO can be deployed in an environment where the Server Area Network is separate from the Storage Area Network, as illustrated in Figure 7. For example the Server Area Network can be InfiniBand or the emerging PCI Express Network, which provides lower latency and higher throughput than traditional Storage Area Networks such as Fibre Channel and 10G IP. In this configuration CacheIO serves both as a performance accelerator and as a gateway between the two networks, providing high throughput, low latency via the Server Area Network while offloading most of the traffic from the Storage Area Network.

There are two active-active configuration options for redundant CacheIO deployment. In the first option as shown in Figure 8 an even number of LUNs are split between two CacheIO appliances. The application, in this case Oracle ASM, load balances all I/Os across all the LUNs. 50% of traffic goes through CacheIO A and 50% CacheIO B. There is no data traffic and associated latency between CacheIO A and CacheIO B. This active-active configuration option is recommended for applications or LVMs (Logical Volume Manager) that can load balance across an even number of LUNs.

For applications that cannot load balance themselves, Figure 9 shows the second active-active configuration option where the multi-path device load balances I/Os across all four available paths. This load balancing is transparent to the application. Again 50% of traffic goes through CacheIO A and 50% CacheIO B. In this case there is data traffic and associated latency between the two CacheIO appliances via the InfiniBand interconnect.




Figure 8: Active Active Application Access




Figure 9: Active Active Storage

Configuration at Orange Silicon Valley

Figure 10 illustrates the physical configuration at Orange Silicon Valley (OSV). Five Oracle RAC nodes are connected to redundant CacheIO appliances via a single InfiniBand switch as the Server Area Network. Redundant CacheIO appliances serve as performance accelerators and gateways between the Server Area Network and the Storage Area Network iSCSI over 10G IP. Only one InfiniBand port is available on each RAC node so a virtual HBA port is created to connect to the second CacheIO. Each CacheIO appliance has two 10G ports connecting to the redundant controllers on the NetApp array.

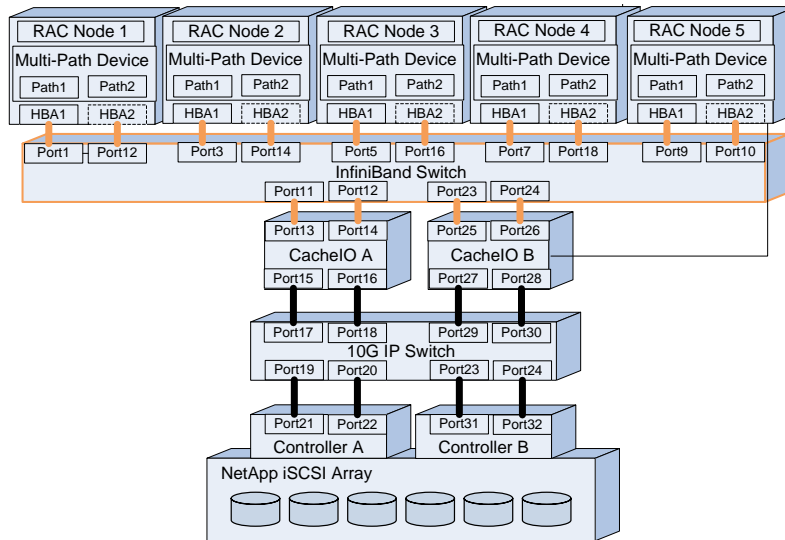


Figure 10: OSV Physical Configuration

Figure 11 illustrates the active-active logical configuration at OSV. Oracle best practice requires mirroring for redo logs, which reside on the SSD storage of the CacheIO appliances. Oracle ASM stripes all data I/Os across all data LUNs, which are evenly split between two CacheIO appliances.

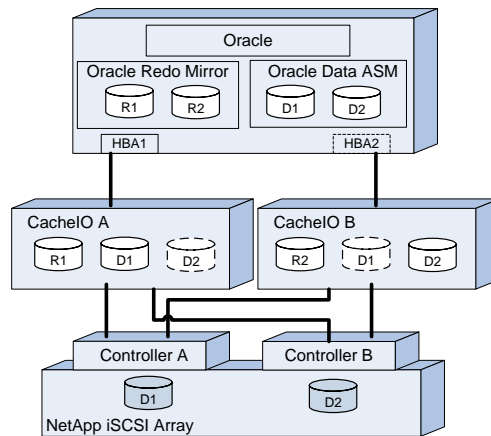


Figure 11: OSV Logical Configuration