



Tech Brief

VMware vSphere Flash Read Cache and OCZ VXL Software

Providing Your VMware Environment with the Full Power of Flash Caching
and Virtualization

Contents

	Page	
1	Introduction	2
2	ESXi 5.5 & vSphere Flash Read Cache	2
3	VXL Caching & Virtualization	3
4	Key Capabilities of Flash Caching	4
	4.1 Cache Allocation	4
	4.2 Distribution of Flash Resources	4
	4.3 vMotion Support	5
	4.4 Advanced Caching Algorithms	5
	4.5 High Availability (HA) & Fault Tolerance (FT)	5
	4.6 Remote Central Management	6
5	The Bottom Line	6

1 Introduction

VMware® vSphere® is one of the world's leading virtualization platforms for building cloud infrastructures and its architecture includes powerful computing capabilities that enable x86 server resources to be virtualized for allocating multiple application workloads simultaneously. The platform includes network services, storage, data protection and security, and uptime services (such as High Availability and Fault Tolerance), all designed specifically to enable an efficient, virtualized environment.

The latest 5.5 release of vSphere includes a generic read cache functionality as part of the platform bundle (available as an 'Enterprise Plus' license). This vSphere Flash Read Cache (vFRC) enables VMware users to utilize introductory server-side flash caching and/or experiment with flash-based storage especially if their experience has been hard drive-centric. VMware users that want robust enterprise functionality for their virtualized flash resources should evaluate OCZ's VXL Software.

The purpose of this Tech Brief is to present the new vFRC platform, identify key capabilities of host-based flash caching, and showcase the advanced features available in OCZ VXL Software. The result is a VMware environment enabled with optimal choices in flash caching and virtualization.

2 ESXi 5.5 & vSphere Flash Read Cache

ESXi 5.5 is the latest version of VMware's vSphere, and vFRC 5.5 is the new caching system integrated into vSphere to enable flash-based read caching. The vFRC architecture allows multiple flash-based devices to be pooled as a single vSphere entity (through Virtual Flash Resource). Once the flash pool has been created, each VM can be assigned a static portion of the cache resource pool (through vSphere Web Client). Figure 1 depicts the vSphere 5.5 vFRC configuration.

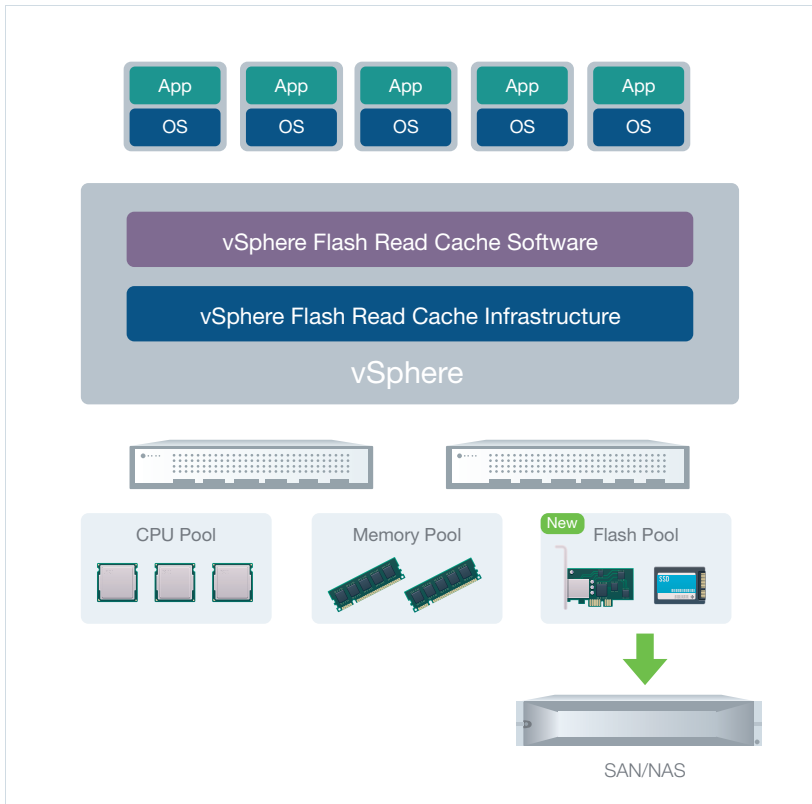


Figure 1: vSphere Flash Read Cache (vFRC) configuration

3 VXL Caching & Virtualization

When combined with an OCZ Z-Drive R4 PCIe SSD, VXL Software enables the delivery of a complete virtual performance system for enterprise customers looking to efficiently distribute flash resources across VMs as a means to maximize performance of key applications. It delivers data caching and flash virtualization into VMware ESXi hypervisor platforms and uses the SSD as an acceleration cache that resides in the host VMware ESXi server as outlined in Figure 2.

Flash volumes are automatically provisioned from host flash, with unallocated flash transparently distributed as a dynamic cache resource. No matter how many VMs are running simultaneously, VXL Software dynamically assigns Z-Drive SSD flash resources by need at all times.

Deployed as a virtual appliance, VXL Software works directly with the ESX or ESXi hypervisor layer to manage and distribute the on-host flash resources. It does not require guest agents or special drivers to communicate with each accelerated VM.

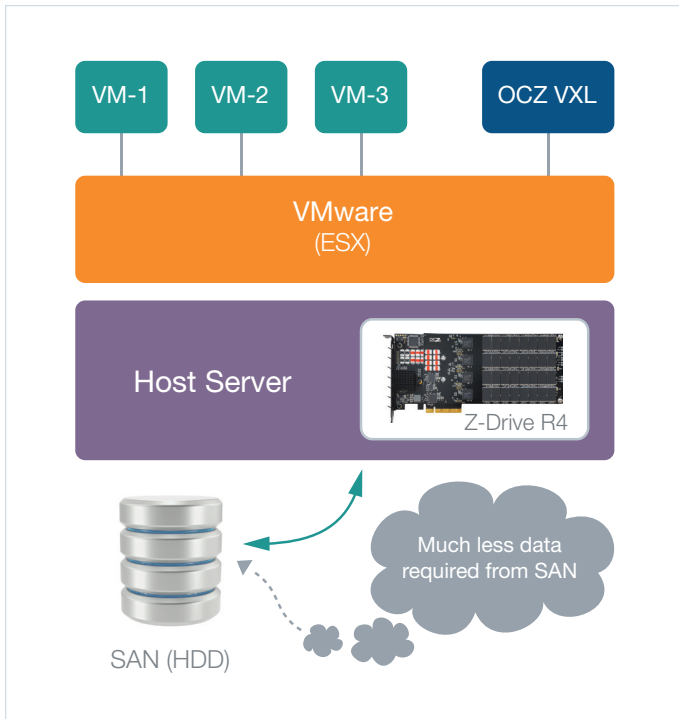


Figure 2: OCZ's Virtualization Platform

4 Key Capabilities of Flash Caching¹

vFRC 5.5 provides general read cache functionality as part of the platform bundle while VXL Software provides advanced enterprise caching and virtualization. Each provides key capabilities of flash caching as follows:

4.1 Cache Allocation

vFRC uses a static cache allocation model that requires a fixed amount of cache space to be designated for each guest VM. As workloads shift (depending on importance) the amount of cache available remains static and does not dynamically adjust to workload changes. Each VM must be configured individually to determine how much flash each VM deserves which will extend deployment times.

VXL Software uses a dynamic cache allocation model whereby the flash cache is pre-loaded with application data to accommodate specific business-critical workloads that need to run at scheduled times. It utilizes a unique 'business-rule' pre-warming cache engine that adapts the flash cache to the workload cycle assuring that relevant data resides in cache in time for use by the application.

4.2 Distribution of Flash Resources

All SSDs configured in the vFRC layer are allocated for use as a cache partition only (and not as a flash partition) limiting the flash resources available for storage capacity.

VXL Software provides on-demand distribution of flash resources between all connected VMs so that the Z-Drive R4 card that it supports can be virtualized as a highly-available network resource to be shared amongst any VM residing on the local ESXi server or on other servers in the virtualized cluster. With this design approach, no VM inefficiently occupies flash when it can be better used elsewhere in the environment and is not limited to accelerating only the application running on the same server that the cache resides.

1. This section includes vFRC 5.5 functionality as understood by OCZ from the VMware ESX and VMware ESXi Data Sheet and from available open industry resources. It does not cover VMware vSAN as this is a separate product introduced as a beta by VMware and is not a generally available VMware product at the time of this writing. vFRC and vSAN are not inter-compatible as the flash device can either be assigned to vFRC or to vSAN and cannot be used by both at the same.

4.3 vMotion Support

VMware vMotion is a capability that reduces application downtime of planned server maintenance by migrating live VMs between host servers. In the vFRC model, the cache needs to be migrated prior to enabling vMotion requiring a careful IT plan for these migrations to not lose cached data, experience performance drops or extend completion times.

In the VXL Software model, cached data is treated as a virtualized storage entity that can be continually accessed between ESX servers whenever VMs are migrated. As VMs migrate from the local ESX host to a remote host, VXL Software identifies the VMs that are being remotely serviced and transfers the local flash connectivity to remote connectivity eliminating drastic performance drops while enabling transparent and dynamic VM migration from one server to another without the loss of cache.

4.4 Advanced Caching Algorithms

VXL Software features advanced policy-based algorithms that enable IT administrators to select from a set of pre-optimized application-specific caching policies which help them make knowledgeable selections on what data to store in SSD cache. The VXL caching algorithms take into account specific needs of each VM and their priorities based on application policies as each policy utilizes data collected from hot zone detection, command size inspection and a host of other variables.

4.5 High Availability (HA) & Fault Tolerance (FT)

The ability to provide uninterrupted service for business-critical applications even during severe system failures is a requirement for any sized data center. To enable this 'no data loss/no VM downtime' environment, virtualized data must be synchronously mirrored and continuously available across the network, through HA and FT services, especially for mission-critical VMs.

VXL Software is designed to provide the critical storage services that maximize uptime across the virtualized infrastructure reducing both planned and unplanned downtime of server and storage maintenance. It enables Z-Drive R4 flash volumes to be virtualized and synchronously mirrored making them continuously available to support advanced HA and end-to-end FT services from within the virtualization host without the need for any back-end SAN or storage appliance. To enable this level of all-silicon 'SAN-less' virtualized services, VXL Software synchronously replicates flash data through mirroring (between Z-Drive R4 cards residing on redundant cluster servers), keeping two live identical copies of VM data (on two separate ESX hosts), down to the last command as outlined in Figure 3.

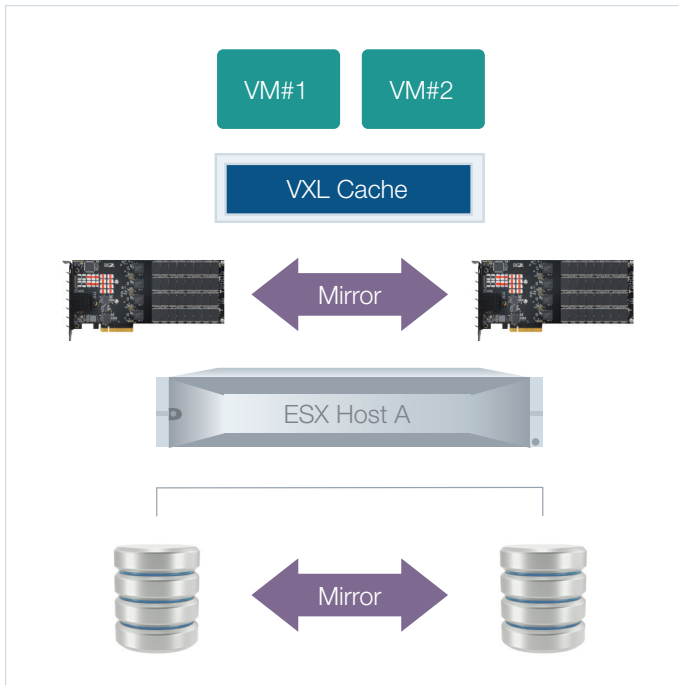


Figure 3: VXL Software mirroring in support of VMware HA and FT services

The net impact of this SAN-less, flash-only infrastructure makes the virtual environment significantly simpler to manage, and since all data is stored on Z-Drive cards inside of each server, flash-level performance and lower power and cooling is delivered.

vFRC does not support FT and VMs using this mechanism are automatically blocked from being enabled as fault tolerant VMs. To support HA, vFRC requires that the secondary servers include flash resources at the level allocated to the original VM. If there are not enough resources available at the time of the failover, the VM will not load on the alternate server.

4.6 Remote Central Management

Both vFRC and VXL Software support local device management. OCZ extends this capability through its StoragePro flash management system that centrally monitors and manages Z-Drive SSDs connected to

network servers, storage arrays or appliances. As a network-accessible management system, StoragePro provides a central view of its enterprise flash resources and simplifies flash administration by enabling connection to multiple host systems.

5 The Bottom Line

Utilizing flash technology in a VMware environment gives an organization the opportunity to increase its overall performance levels and quality of service (QoS) dramatically, be more flexible in day to day tasks, secure data center uptime while simultaneously reduce cost and maintenance resources. As caching and virtualization are added to the mix, the number of VMs that can run on a host infrastructure increases enabling growth in the data center without excessive increases in TCO.

With the availability of flash-based caching and virtualization solutions, such as VMware vFRC and OCZ VXL Software, more enterprises and cloud providers are transforming the way that data is stored and handled in virtualized environments. Reaping the benefits of increased server utilization, efficient flash distribution and the ability to use cost-efficient commodity hard drive storage, CAPEX and OPEX are reduced in the data center while boosting enterprise applications with flash performance. OCZ's VXL Software platform provides a complete set of flash virtualization and acceleration functionality designed to address the needs of enterprise workloads.