

# WHERE TO PUT YOUR CACHE

## Best Practices for Deploying Flash Cache Resources in IT Environments

Dr Allon Cohen, VP Marketing and Product Management, SANRAD

### Why the World Needs Flash

Nearly every IT professional has time and time again received similar distress calls from frustrated users: "My application keeps freezing on me; new page loads are taking ages", "The application isn't responding; I keep pressing the OK button but it takes so long for the screen to refresh..." More often than not, the culprit is well known: *a storage access bottleneck*.

Increased use of server virtualization and high IOPS applications such as Databases and E-mail servers have improved server efficiencies but at the same time have created a strain on SAN storage resources. In growing installations with more and more users being added, the IOPS and bandwidth demands on the same legacy storage controllers continually increase. The physical performance limitations of the storage disks and connectivity are quickly reached and end-users begin to suffer from delayed responses.

Today, a key tool for alleviating the IOPS limitations of traditional storage systems is the use of flash based storage as a caching tier. Flash based storage can produce up to thousands fold increases in IOPS when compared with HDD and hence is a perfect candidate for use when storage performance is becoming a bottleneck in the IT organization. However, flash based storage does not come without its own limitations. Prices per Gigabyte are still very high and life time endurance is severely limited. Both pose considerable barriers to the adaptation of flash only storage systems.

Ultimately the best value for your investment is by deploying "just enough" flash as a cache in front of larger amounts of traditional storage. Common deployments have shown that in real life operations storage access is not evenly accessed across large data stores. Most of the data stored in arrays remains relatively dormant while a small collection of frequently accessed "hot zones" see a very large percentage of storage access operations. Therefore, to receive an immediate and significant boost in performance it is enough to deploy a flash cache to protect only these "hot zones". It is easy to see that saving dormant data zones on costly flash storage produces no performance gains, why still inflating the storage costs. By merely deploying a relatively small flash cache protecting only the "hot zones", the organization gets a similar boost in performance, but at a fraction of the cost.

## Now that I Have a Cache, Where Should I Put It?

While there is widespread agreement in the storage industry on the benefit of using flash based storage as a cache, there are competing approaches as to how to deploy this cache in practice.

The varying approaches can be group in the three basic categories:

A. **Host based Caches** – Under this approach the cache is deployed inside the servers hosting the performance challenged applications. Flash is added to the host server either onboard a PCI based extension card or as a local SSD disk replacement. Data from the applications is cached locally before even leaving the host towards the storage arrays. The data cached under this approach is particular to the instance of the application running on this specific host.

B. **Storage Array based Caches** – Adherents of this approach prefer to place the flash cache inside the storage array, usually as SSDs replacing some of the HDD bays. The storage array will have built in logic to decide which data to place in the cache. When read requests are received the array will answer them from the flash cache when appropriate. The cache data servers to accelerate the particular SAN array it is residing on.

C. **Network Acceleration Appliance based Caches** – With this approach the cache is placed in a storage network appliance. The appliance resides between the serviced hosts and any underlying storage arrays. Hot data is cached as it is transferred through the SAN (storage area network) and is ready for retrieval when ever it is needed again. In this case, as the cache resides in the network it can service multiple hosts and multiple storage arrays.

Each of these approaches carries with it some advantages and some disadvantages, and ultimately the best approach to take depends on the use case IT requires the cache for.

Table I below summarizes the key benefits and challenges of each approach and shows some of the workloads that best fit each approach.

**Table I - Key Attributes of Different Flash Cache Locations**

	Host-Based	Storage-Based	Appliance-Based
Pros	<ul style="list-style-type: none"> <li>• High speed/High IOPS PCIe connectivity</li> <li>• Dedicated application acceleration</li> </ul>	<ul style="list-style-type: none"> <li>• Supports Server Virtualization</li> <li>• Centralized Maintenance and Replacement</li> </ul>	<ul style="list-style-type: none"> <li>• High speed/High IOPS PCIe connectivity</li> <li>• Supports Virtualization</li> <li>• No storage bottleneck/contention</li> <li>• Cluster friendly</li> </ul>
Cons	<ul style="list-style-type: none"> <li>• No support for server virtualization and vMotion</li> <li>• Not Cluster friendly</li> <li>• Difficult to replace/maintain</li> </ul>	<ul style="list-style-type: none"> <li>• Lower IOPS/bandwidth connectivity</li> <li>• Storage controller bottle neck/contention</li> <li>• Vendor lock-in/ High Cost</li> </ul>	<ul style="list-style-type: none"> <li>• 10 GE hop between storage and server.</li> </ul>
Best Fit For	<ul style="list-style-type: none"> <li>• Dedicated servers for specific tasks</li> </ul>	<ul style="list-style-type: none"> <li>• Storage vendor preference (e.g. “Single Vendor Shop”).</li> </ul>	<ul style="list-style-type: none"> <li>• High IOPS applications Exchange/Databases</li> <li>• VMware/XEN/Hyper-V</li> <li>• VDI</li> <li>• VoD/Media Streaming</li> </ul>


## When Should I Use My Networking?

On the one hand, the key advantage of placing a flash cache in the host servers is the fast access it provides to the flash. Because the flash cache resides in the host it can connect using the internal PCIe expansion interface. PCIe is an ideal connectivity option for flash because it is a much faster connection to the CPU than traditional storage based SAS or SATA. It can also carry a higher IOPS load due to its higher power feed. In the world of flash access the higher wattage translates into more IOPS and faster bandwidths.

On the other hand, placing the flash in the host does not come without its downsides. Because the flash is local to the host machine it can service only the particular host it resides on. This limitation is particularly troubling when the environment is either a virtual server environment or a server cluster. Powerful virtualization tools such as server vMotion are negated because the local cache can't be transferred with a virtual machine. Similarly server cluster failovers become more difficult because the local cache is lost with the loss of any server. These issues limit the usefulness of host based caches for general SAN usage and relegates them mainly to dedicated servers running very specific workloads.

The opposite holds for array based caches. In array based caches the cache can service any host connected to the specific array. Therefore, server vMotion and clustering can be supported between different physical hosts. However, the connectivity itself becomes a challenge. As the flash is placed within the storage array it uses the same controller as the other (HDD) disks in the array. This places the flash storage in contention with the other slower disks accessed with the same controller – the very same contention that was the initial cause of the performance issues in the first place. Furthermore, in many array solutions the flash cache will be placed in a storage bay initially designed for slower SAS connectivity, thus negating the performance boost achievable through flash. Placing flash inside an array is in many ways similar to buying a race car and trying to use it in morning rush hour traffic – you might get some points for show but you won't reach your office faster.

When your IT environment utilizes virtual servers or clustering for high IOPS applications such as Exchange or SQL Server, **appliance based caches** present the best caching solution. On the one hand, such an appliance is a central location for a cache that can service multiple servers and arrays, while on the other hand network appliances may deploy PCIe connectivity making the best use of the flash resources they contain.



With network based caching appliances, all cache hits are directly answered from the appliance without any storage controller contention. At the same time the physical location of a server workload becomes irrelevant to the cached data, so virtual server vMotion and server clustering is fully supported. Furthermore the appliance serves as a central location which provides easy access for servicing of the cache. Cache sizes can be easily managed to fit the operational requirements, providing just the right

amount of cache at the right time the enterprise requires it. For such IT environments network based *cache acceleration appliances* provide a solution that is both the best fit in terms of flash cost efficiencies and the best technological fit to the application usage requirements.

## **About the SANRAD V-Switch XL Series**

The V-Switch XL Series is SANRAD's series of network based SAN acceleration appliances. Utilizing appliance embedded FLASH as a high performance cache, the V-Switch XL appliances deliver an immediate performance increase to SANs. The V-Switch XL series can further serve as a connectivity bridge enabling the connection of 10-GE servers to an FC SAN by upgrading 2Gb or 4Gb FC SANs to 10-GE.

The V-Switch XL appliances work either in transparent acceleration mode (accelerating existing SAN volumes and VMware datastores) or in virtualize & accelerate mode (creating new accelerated virtual volumes using the SAN as a storage pool). Users deploy the V-Switch XL appliances to deliver significant acceleration of storage response time, deliver high IOPS for demanding applications, and maximize utilization of 10-GE networks.