



## Why use SSD Caching & Tiering for VDI

WHITEPAPER



In today's market, storage vendors have biased towards offering either caching-only or tiering-only solutions. However, detailed analysis shows that neither solution on its own provides a complete solution for a variety of workloads, including the growing popularity of VDI environments. In this context, this whitepaper will discuss the StorTrends approach to advancing the efficiency and performance for VDI deployments.

Before dissecting each approach, it is helpful to be mindful of the different I/O profiles of typical VDI users as shown below. These typical I/O loads will be referenced later in this document.

User Classification	IOPS requirement per user	R/W Ratio Steady State	R/W Ratio Boot Storm
Task User (Light)	3-7	80:20	80:20
Knowledge User (Medium)	8-16	60:40	80:20
Power User (Standard)	17-25	50:50	80:20
Power User Heavy	25+	20:80	80:20

*Chart 1: VDI User Profiles*

These user profiles will be discussed within two use cases: Boot Storm and Steady State. Note that Steady State is the more important of the two as the boot storm can be controlled as a staggered boot.

## The Caching Approach

In truth, most storage vendors adopt the Caching approach as it is the easiest. With Caching, SSDs can be used to augment the Read cache, while NVRAM is used for Write cache. The drawback is that majority of VDI environments have I/O profiles which are more write-biased. So, the Read cache is only used during boot time, but when it comes to Steady State, more Write cache is required to support the amount of data being written to the storage. The available NVRAM size is comparatively small compared to a size of an SSD, so in a typical 1000-user VDI environment it would be not be able to sustain the throughput over a period of time, due to the bottleneck of the traditional/spinning disks at the back.



Some vendors claim that they are able to alleviate this handicap by writing sequentially to the disk. While this method initially delivers high performance, eventually the sustained performance eventually drops off over time - as continued data writes to spinning disk must move data around as it becomes fragmented on the disk. In contrast, the Read cache is not of much use outside of boot time, as the overall I/O profile is mostly writes. This cache is repeatedly invalidated, so the expensive SSD space goes to waste. Also, due to the consistent thrashing, changes to the SSD Read Cache could quickly degrade the endurance of the SSD. Ultimately, upgrading these expensive NVRAM devices would require the controller to be shut down and a purge of the data in the NVRAM to the disk.

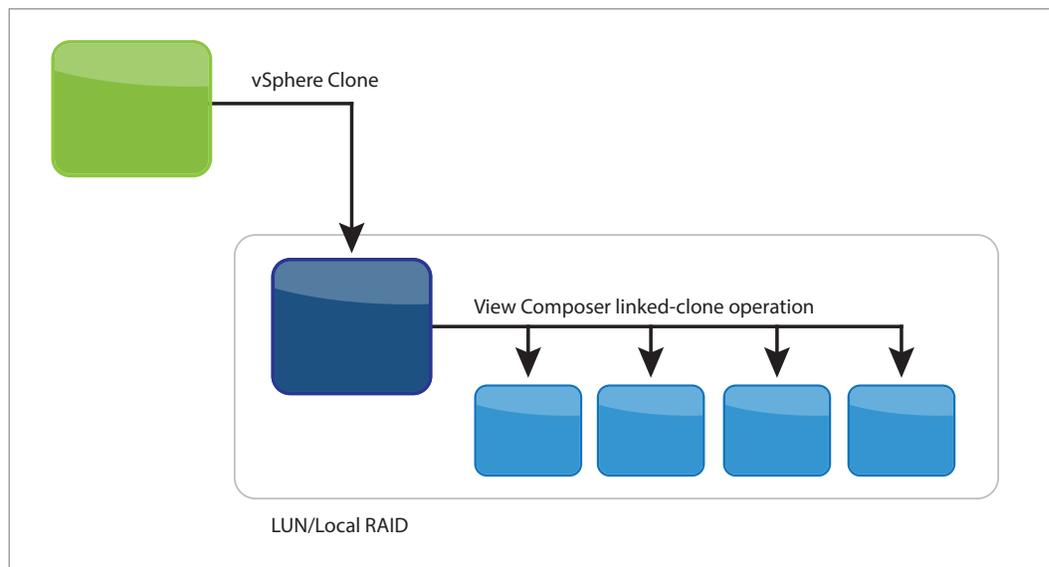
### The Tiering Approach

The Tiering approach is adopted by some storage vendors who have efficient tiering algorithms. The drawback of this method is that the SSD tier would need to be large enough to cater towards both new I/O and hot data at the same time. If the SSD tier is under provisioned, then it will lead to new I/O hitting the spindle disks. If the SSD tier is over provisioned, it would increase the cost of the solution as SSD drives are expensive for a customer looking to minimize cost by going for a hybrid solution. Certain vendors, who have done tiering for a long time, have had major changes in core architecture adding SSD tiers. For those vendors who do not allow for manipulation of the tiering promotion and demotion strategy, a full SSD tier results in spinning disk speeds for all applications.

### StorTrends Tiering & Caching Approach

After detailed analysis of the I/O profiles in a VDI environment, StorTrends has come up with a combined approach of both tiering and caching so that we cater to both the use cases in a VDI environment, Boot Storm and Steady State.

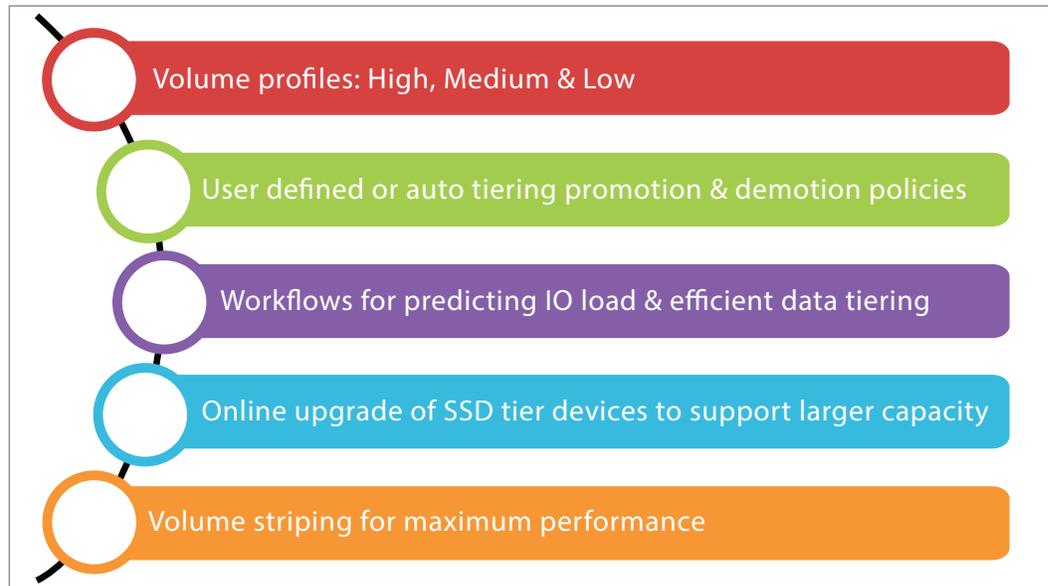
StorTrends can use the VMware linked clone approach to our customer's advantage.



*Figure 1: VMware vSphere Linked Clone*

A typical StorTrends VDI configuration would be the use of our 3500i Hybrid Array with 4 SSD drives, 2 for SSD caching and 2 for SSD tiering, and 12 high capacity SAS or NL-SAS drives.

There are only 3 types of IO patterns from StorTrends point of view, New Writes, Over Writes, and Reads. How effectively these are taken care of in a VDI solution to maximize the use of SSD space determines how effective the VDI solution would be. Below we go through both the tiering and the caching aspects of the StorTrends 3500i and how they work hand in hand to provide an effective VDI solution with hybrid storage



## StorTrends Tiering

StorTrends divides volumes into three different profiles: High, Medium, and Low. In a high profile volume, IO would be redirected to the SSD tier. IO for a medium profile volume would be redirected to a less expensive tier such as the 15k SAS tier. IO for a low profile volume would be redirected to the least expensive tier, such as the NLSAS tier. StorTrends stores the replica images in the High Volume Profile so this performance volume resides in the SSD tier. This SSD tier will handle all reads during a boot storm and also store frequently accessed data.

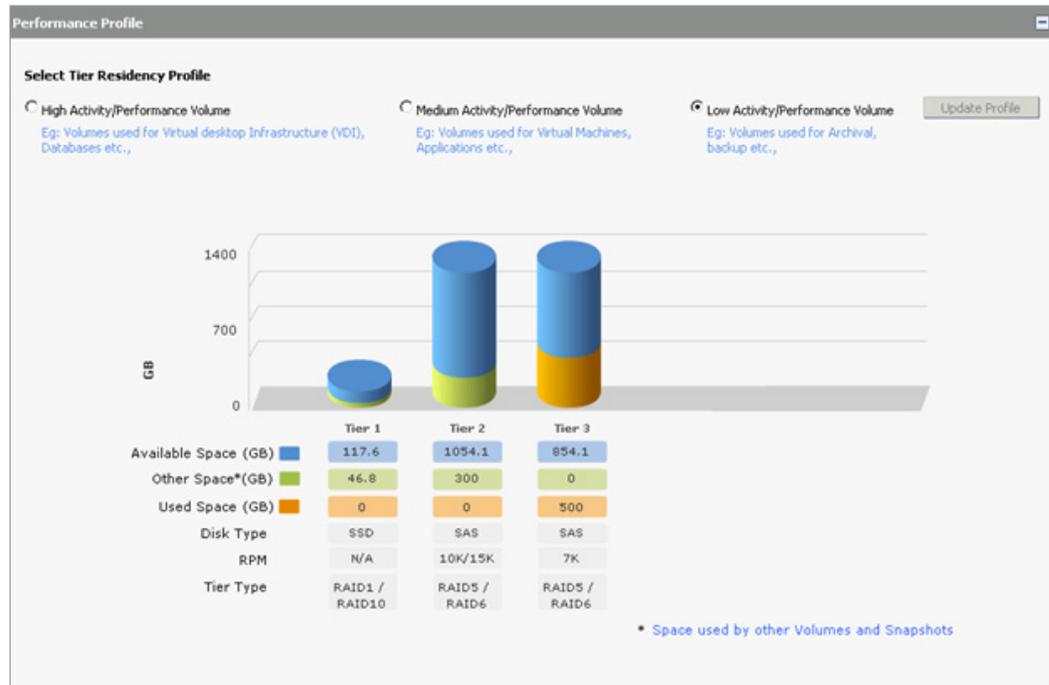


Figure 2: StorTrends Volume Profiles

Depending on the promotion policies, which StorTrends allows to be user defined, the hot data is promoted to the SSD tier and based on the demotion policies the cold data is demoted down to the lower tiers. This leads to a smaller foot print of the replica image on the SSD tier over a period of time, which is required to boot the virtual machines, thereby making best use of space and allowing for more hot data to be promoted and new IO to land on the SSD tier.

Overwrites and reads land on the SSD tier for frequently accessed data and new writes also land on the SSD tier. Data stored in the SSD Tier is not duplicated in the SSD cache; thereby in effect the SSD cache space augments the available SSD space because it also acts as a tier.

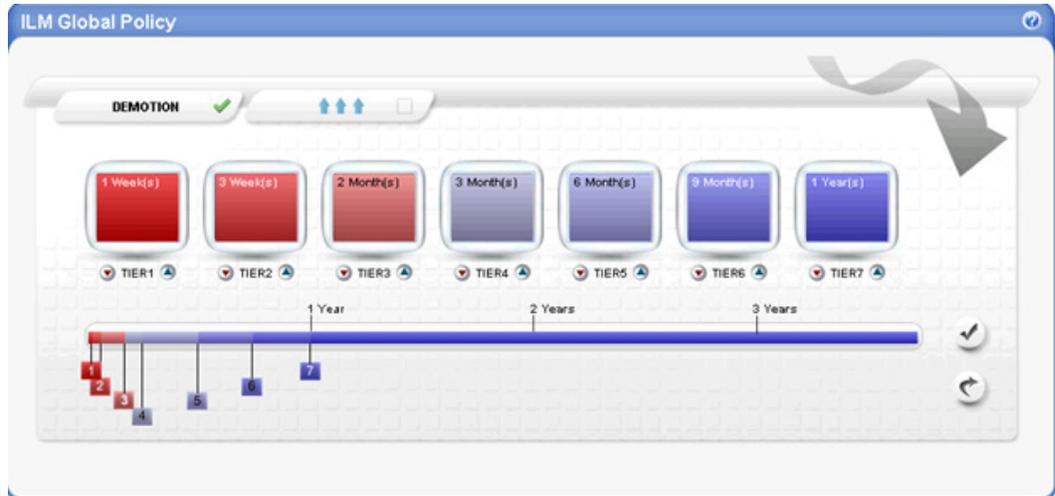


Figure 3: StorTrends Demotion Policies

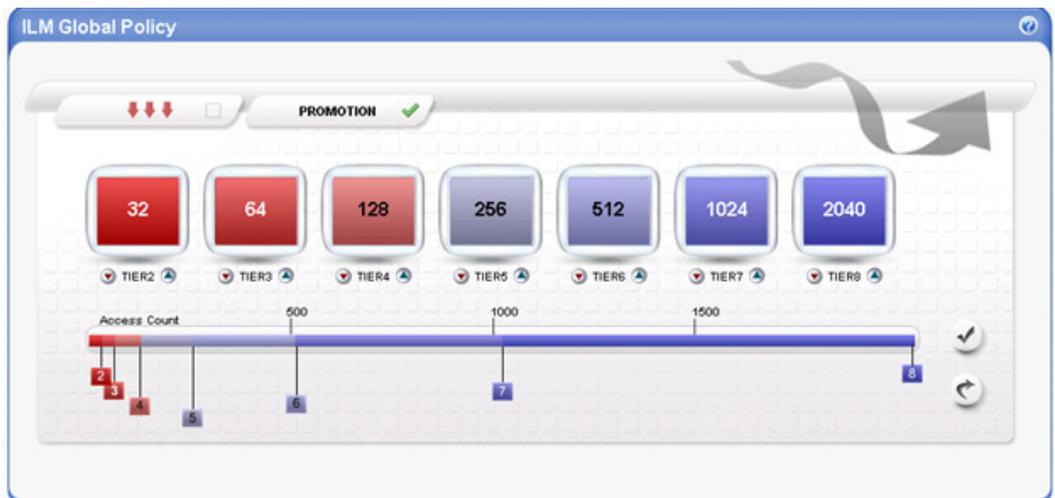
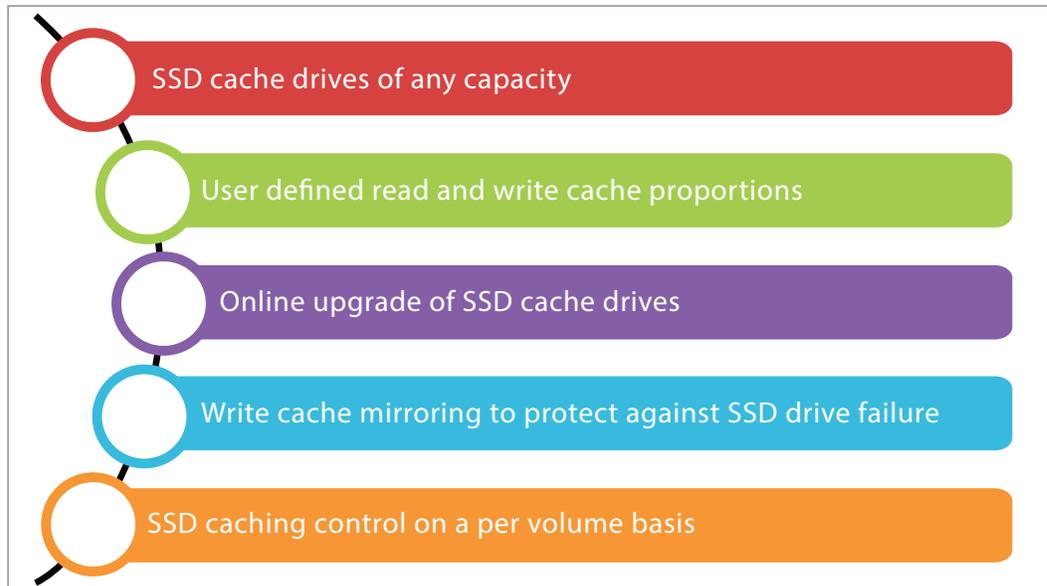


Figure 4: StorTrends Promotion Policies

StorTrends handles both boot storm and steady state IO by the SSD tier in an optimal fashion. As tiering is based on policies and workflow, it may take time to move inactive data down to the lower spindle disks. StorTrends' Workflow management and ILM need to predict the workload and demote data down during less active periods of the day. This is where caching comes into the picture.



The user has the flexibility to choose between drives which can be used for the SSD cache to fit the required solution. These drives can also be upgraded while the system is online in case of a need for higher capacity drives as business needs change, drive failures occur, or drives nearing their endurance limits need to be replaced.

The SSDs are partitioned into read cache and write cache. The proportion can be either default 60:40 or it could be specified by the user depending on the type of IO profile users that are going to exist in the VDI environment. The user has the ability configure or change the SSD cache drives as read-only cache or write-only cache or r/w cache at any point in time, thereby offering the flexibility to change the r/w proportion based on changing IO load.

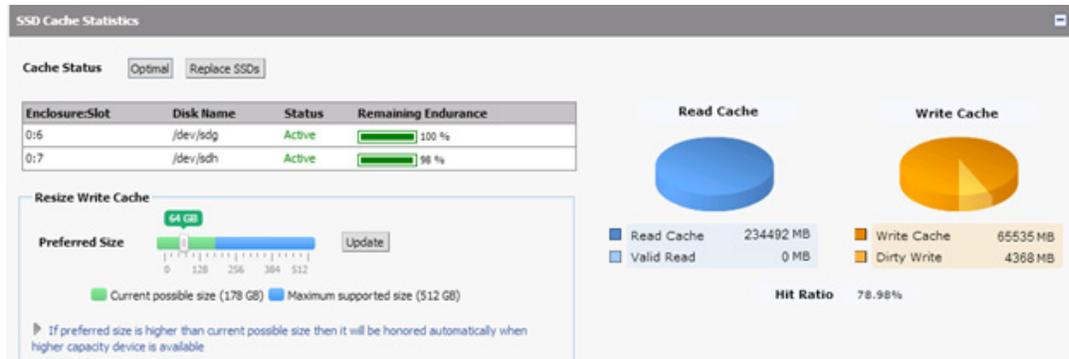


Figure 5: StorTrends SSD Caching

The write cache is mirrored for redundancy and the read cache is striped for maximizing available read space and performance.

Caching is enabled on a per volume basis which gives StorTrends customers more control in selecting what data needs to be cached. This is shown in the figure below:

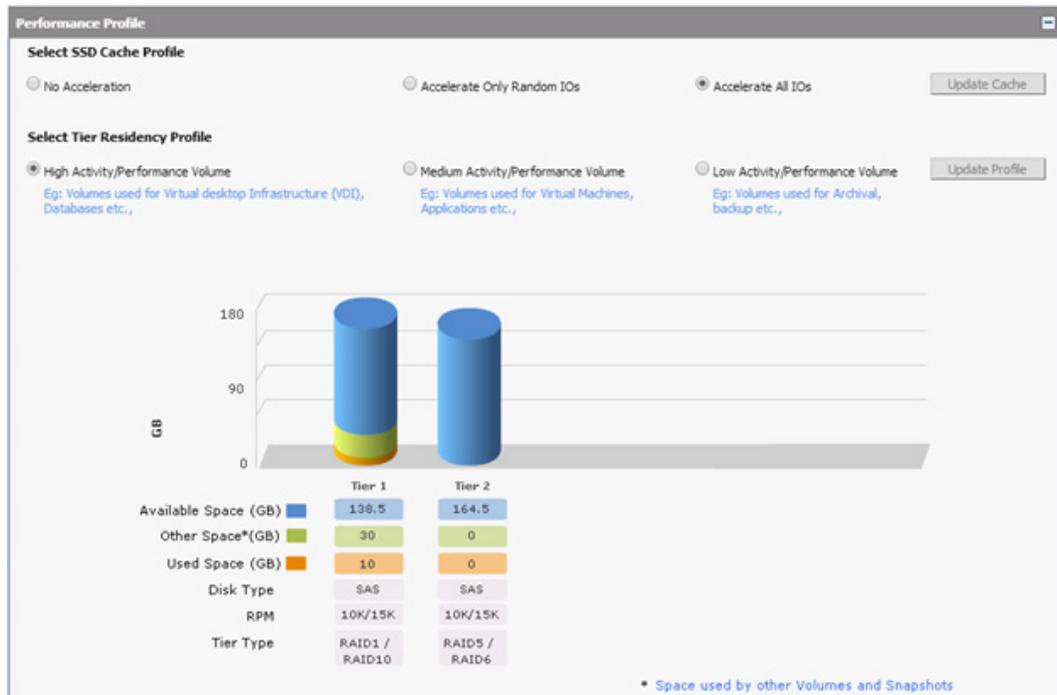


Figure 6: StorTrends Promotion Policies



## Caching & Tiering Working together to provide performance in VDI

StorTrends recommends the golden image, or the base replica, to be stored in a high performance volume without caching support. The linked clone volumes can be set as either high performance volumes or medium performance volumes, but on these volumes the SSD caching is enabled. With SSD caching enabled, any IO which does not land on the SSD tier would land on the SSD cache. Reads, new writes, and overwrites are handled by the SSD cache for any miss from the SSD tier. The purging of the write cache data is done in the background to the spindle disks thereby maintaining free space in the SSD cache drives at all times. StorTrends has engineered the purge to be done in a method that it is sequential to the disk, thereby taking advantage of the throughput of the disk, rather than the IOPS. This purge happens in the background by the StorTrends system sensing the IO load. StorTrends continuously monitors the IO load, picks the idle time in the IO flow, and varies the purge bandwidth dynamically based on the IO load. The IO is maintained while old data is moved to the spinning layer to make room for fresh data.

By using SSD tiering and SSD caching together, StorTrends manages the SSD endurance very efficiently. In the cache only approach, data may be evicted prematurely during the block replacement. The data which apparently would have got more cache hits in the future may have been evicted from cache in order to accommodate new data into the cache. Since StorTrends can keep the long-term hot data in the SSD tier, the hot data from SSD cache promotes to the SSD tier instead of cycling the data between SSD cache and spindle drives, thereby preventing thrashing and also allowing for new data population in SSD cache.

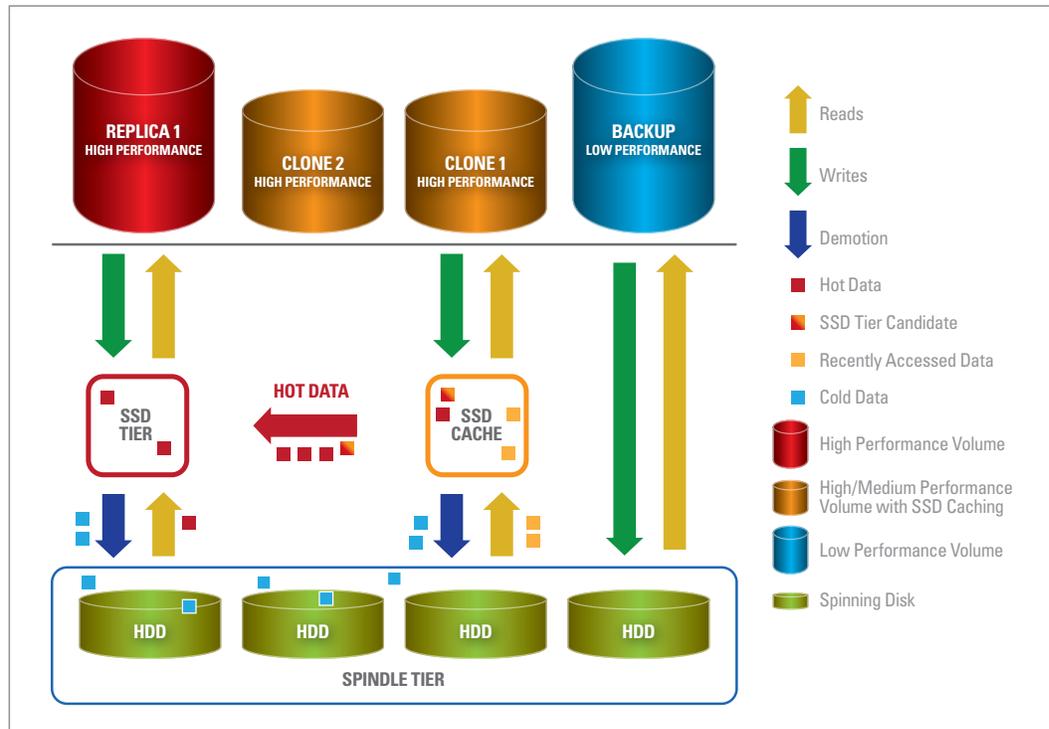


Figure 7: SSD Caching & Tiering working together harmoniously

Similarly in a SSD tier only approach, any access on cold data, which resides on the spindle drives, requires the policies to kick in for the promotion to happen to the higher flash based tier. Until this threshold is met, performance suffers as the application reads the same data over and over again from the spinning disk layer.



In StorTrends, cold data which is recently accessed is brought to the SSD cache first before it is truly identified as hot and gets a place holder in the SSD Tier. This provides a cushion by delivering the performance needed instantaneously if the same data were to be accessed again before the policy kicks in for promotion. It also offers a cushion, if the data accessed does not meet the policy requirements for promotion to the SSD tier. So in effect all the SSD drives are used in a balanced manner which also increases the endurance of the SSD drive.

StorTrends' goal is to have all IO land on the high performance SSD drives, utilizing SSD for caching and tiering gives high performance to the VDI user at any point in time. StorTrends engineers have meticulously tuned the StorTrends features to provide the highest performance available to our customers. StorTrends 3500i Hybrid Array is a product designed to meet the demanding nature of VDI implementations big and small.