

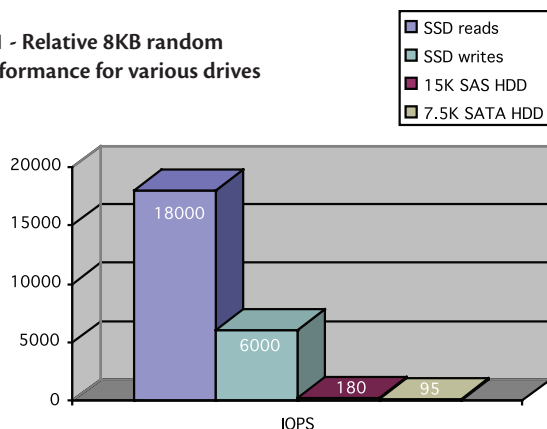
Many applications today are demanding very high levels of I/O performance from storage systems. These performance requirements often become the determining factor in the cost of the storage array as the number of drives needed to meet the I/O requirements exceeds what's needed for capacity. To address these performance requirements in a different way, IceWEB leverages high performance Solid State Drives (SSDs) and other memory-based drive technologies to provide storage systems capable of ultra high performance for I/O intensive applications. SSDs are still expensive relative to SAS/SATA drives from a capacity standpoint, but they have now reached the point where they are a cost-effective solution where ever I/Os per second (IOPS) are the critical requirement.

IceWEB offers storage system configurations utilizing SSDs which provide incredible application performance benefits, while at the same time lowering power consumption and cooling requirements. Let's take a look at how these configurations can help you solve your performance problems while lowering your product, administrative and environmental costs.

Solid State Drives Deliver Amazing I/O Performance

Solid State Drives (SSDs) utilize flash memory instead of spinning disks to store data. Since there are no moving parts, you eliminate much of the latency compared to a normal hard disk drive (HDD), thus allowing faster reads and writes of data. In particular, an SSD is very effective for small block random I/Os; about 30 times as fast for writes, and an astounding 100 times as fast for reads (as shown in Figure 1 below). One SSD is as fast as 50 or so HDDs for random I/O, the dominant I/O pattern for databases, most unstructured data, and web applications.

Figure 1 - Relative 8KB random I/O performance for various drives



With data center storage capacity requirements increasing 50-60% per year without an end in sight, it is more important than ever to keep volatile power and cooling requirements in check. Without motors to spin or parts to move, SSDs consume about a third of the power of a typical data center disk drive and require far less cooling. As an example, for 100 8K read operations, an SSD uses about 30 Watts, whereas a HDD will use around 6,000 Watts for the same work. Obviously there is a capacity difference, but if performance is the driving requirement, the power savings can be significant.

Typically, SSDs today support from 200GB to 600GB of capacity, so with their outstanding performance characteristics, the number of SSDs required are likely to be driven by capacity needs rather than performance. Due to their high cost relative to HDDs, it makes sense to only use SSDs where absolutely needed for performance. So, let's explore the various ways in which SSDs can be configured to optimize their benefits.

High Performance Primary Storage

The most obvious way to leverage SSDs is as a set of primary storage devices. IceWEB supports the creation of storage pools that consist entirely of SSDs. These are standard pools in that you can create volumes (LUNs) and or file shares, and the IceWEB OS will stripe data across all the drives in the pool. For applications where all the data is "performance critical"; that is, all the data needs to be accessed at the best possible speed, you create the needed volumes / file shares in SSD pools and assign those back to the servers and applications that require them. This is the simplest method to assign SSD resources to applications, but it may not be optimal, as with many applications not all the data is performance critical.

The IceWEB storage system also supports the mixing and matching drive types, so you may elect to configure a mix of HDDs and SSDs into the same array. This allows you to then create separate storage pools for each drive type. You may create as many pools as needed, so a typical strategy is to create a "high performance" pool consisting of SSDs and a "high capacity" pool consisting of SATA drives. The high performance pool is then used for only that data where I/O performance is critical (e.g. database log volumes).

Performance Boost for Normal Storage

SSDs can also be used to boost the performance of SAS and/or SATA drives by using the SSDs to extend cache and contain a "write index". The IceWEB Management Console allows SSDs to be assigned as an "L2ARC" which is a level 2 cache which augments main memory cache. If so configured, the IceWEB OS will use those SSDs both as a Level 2 cache for reads and writes, as well as extra space for deduplication and virtualization tables. SSDs can also be assigned as the space for the ZFS Intent Log or "ZIL". The ZIL is a non-volatile index that records all writes until they are committed to disk. All writes to the array will be written first to the ZIL and committed to the application, and then subsequently to disk using a "lazy" update process. This means that writes get the benefit of SSD speeds without the cost of having all SSDs for storage. Ultimately you need enough drives to handle the sustained I/O rate from the applications, but tactically this can make a huge positive impact to application performance.

A BIG ENOUGH PIPE...

A key factor to getting the most out of Solid State Drives (SSDs) is ensuring sufficient network bandwidth. IceWEB Storage Systems support six (6) 1GbE connections for each head unit which provides sufficient bandwidth for up to ten (10) SSDs running a "typical" I/O load. For larger systems, you can add up to four (4) 10GbE ports or up to four (4) 8Gb Fibre Channel ports. This is enough bandwidth for over 50 SSDs, and provides an amazing amount of performance for your applications.

Best of Both Worlds - Hybrid Systems

The most cost effective method of utilizing SSDs is a combination of the two methods above in conjunction with HDDs to provide the bulk of the capacity. Putting together a hybrid system requires a more specific understanding of the application environment, but can deliver the most dramatic results.

For example, let's say we have a database application where overall performance is important, but the order entry portion is critical to the business and needs to be very fast. The overall application databases are very large (12TB), with the order entry portion consisting of around 200GB. A possible configuration would be to configure six SSDs along with 22 1TB SATA drives. You would configure the drives as follows:

- 1 200GB SSD (L2ARC)
- 1 200GB SSD (ZIL)
- 4 400GB SSDs (High-Perf-Pool)
- 14 1TB SATA drives (DB-Capacity-Pool)
- 8 1TB SATA drives (Gen-Pool)

With this configuration, all database logs plus the order entry database are put into volumes within the High-Perf-Pool while all other databases are put into the DB-Capacity-Pool. Since there are dedicated SSDs for both a level 2 cache and for the write indexes, overall performance will be 2 to 3 times as good for the general database, and the critical order entry database gets the full benefit of running all on SSDs. The Gen-Pool is used for all other applications and general file sharing, so as to not interfere with the database I/Os.



Energize Your Storage

IceWEB Storage Systems based on Solid State Drives breathe new life into existing storage infrastructures, increase server utilization, and reduce the total cost of ownership. SSDs eliminate the wait time for processors associated with disk drive latency. SSDs can replace multiple hard drives, eliminate the need for additional volatile memory, and assist overall performance by augmenting memory as a non-volatile cache and write index. Hybrid storage systems from IceWEB are revolutionizing the way businesses access critical data by overcoming the performance bottleneck inherent to traditional rotating media HDD storage.