

StarWind LSFS™: Turning TBs into IOPS

Introduction

The quantity of data increases each day and the traditional spinning disks evolve accordingly by increasing the capacity as the result of Moore's Law. Despite the flash-based storage is on the market, spindles and tapes will never die – they cost much less.



Problem

Typical spindle based storage system is running out IOPS long before it runs out of capacity. While the disk's density and capacity is increased, the RPM (revolutions per minute) value is not changing, the spinning disks IOPS remains the same. The table below represents the approximate number of IOPS that traditional spindle disks produce.

DEVICE	TYPE	IOPS
7,200 rpm	SATA	~75-100 IOPS
10,000 rpm	SATA	~125-150 IOPS
10,000 rpm	SAS	~140 IOPS
15,000 rpm	SAS	~175-210 IOPS

Spindle disks are physically limited in terms of performance

The average 4 TB 7200 rpm HDD is capable of storing up to 80 VDI VMs, 50 GBs. Since single VDI VMs require ~10 IOPS, 10 VMs is the theoretical maximum for this particular disk, which is by order of magnitude lower in comparison to expectations.



Random I/O reduces performance of HDD

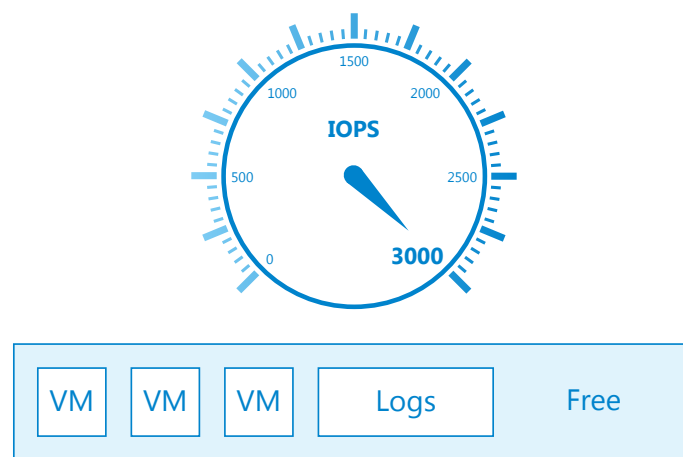
A typical VM workload is dominated by small reads and writes thus processed by one spindle only. That is why the common performance increase technologies like RAID's or I/O distribution across the nodes (i.e. wide striping) do not help much, thus keeping the performance low.

Situation gets even worse if Parity RAID's or wide striping is used on the storage array.

Solution

StarWind's log structuring is boosting performance coalescing multiple small random writes into a stream of big sequential writes. Therefore the IOPS that a single disk or an entire storage cluster can handle is increased exponentially.

The TBs are turned into IOPS by the side effect of using StarWind LSFS. The log structuring uses Redirect-on-Write for snapshotting, writes, etc. That means every new data block is written to the next available place on the disk organizing the data blocks sequentially. Therefore Log-Structuring consumes the space on the cheap spindle storage that would still be just wasted.



StarWind LSFS turns small random writes into big sequential ones boosting performance of HDD

Moreover, Read-Modify-Write sequence with parity RAIDs is no longer an issue since the small random I/O requests are combined into one big I/O request. Since the size of the I/O request is always bigger than one stripe, it is processed by multiple spindles, which greatly increases the performance of entire storage system tenfold.

Conclusion

StarWind's Log-structuring converts wasted inexpensive disk space into the costly IOPS.

The performance of parity RAID systems increases by involving multiple spindles into the processing of every request.

In 2016, Gartner named StarWind "Cool Vendor for Compute Platforms".

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.