

VERITAS™

Veritas SmartIO

Content

Introduction3

What is SmartIO3

Enabling SmartIO4

SmartIO Benefit: Increased Application Performance4

SmartIO Benefit: Tier-1 Performance at Tier-2 Costs5

SmartIO Benefit: Cascaded Performance Increase5

Summary6

Appendix7

Introduction

When you look closely, you'll see that some key aspects of Storage Area Networks (SAN) continue to get bigger and faster: Servers have more computing power and RAM to process bigger data sets under more demanding applications; Network infrastructures are increasing throughput capabilities and taking advantage of the increased CPU through advanced APIs like RDMA. The 15,000 RPM disk, however, is over 10 years old and the laws of physics dictate that the performance characteristics of hard drive (HDD) based storage have peaked.

Yet, despite the continuing increase in data generation, information management and storage budgets are flat or shrinking¹. This means that storage administrators are stretched to meet the needs of the latest applications with fewer resources. The disconnect between application needs and storage capabilities is leading to the increasing adoption of Solid State technology in various forms in the Data Center. Solid State devices provide unmatched IOPs and latency when compared to spinning disk. This allows applications to process more data at faster rates. To take full advantage of the performance characteristics, while still ensuring a low \$/IOP and data availability, enterprises must turn to advanced software solutions to ensure they maximize investment in both new and legacy technologies.

Veritas InfoScale Storage allows customers to utilize the latest in hardware technology with the industry leader in storage management software. SmartIO in Veritas InfoScale Storage for Linux brings the data closer to the application to drive up to 4x performance gains while enabling 60% cost reductions. SmartIO allows the enterprise to decouple capacity from performance and better align architectures to business needs.

What is SmartIO

Integrated into Veritas InfoScale Storage, SmartIO brings granular, online caching to system and application owners to move reads and writes inside the server. SmartIO can be enabled at the Block (VxVM) or File (VxFS) level to meet the SLAs and performance characteristics of the application, without the need to manipulate the architecture or, in a worse case scenario, require downtime.

SmartIO offers two modes of caching, read and writeback, in both single node and multi-node configurations to ensure high performance and high availability of mission critical data. Cache reflection in Cluster File System ensures data coherency and application performance in the event of hardware outage. SmartIO cached data is also persistent across a server reboot, which reduces the cache "warm" period and brings applications back to optimal states in shorter time.

Deep, transparent application integration, combined with advanced caching heuristics, enables SmartIO to target the most relevant data based on different workload patterns.

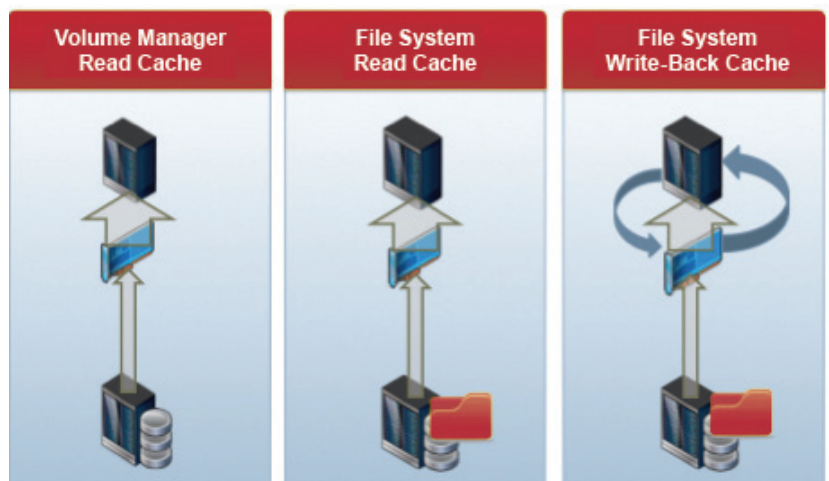


Figure 1. SmartIO Modes

1. Forrester: "Software-Defined Storage Will Sound The Death Knell For Traditional Storage Provisioning"

Enabling SmartIO

SmartIO introduces a new command line interface (CLI), called sfcache, for full management and visibility into cached volumes, file systems, and files. With sfcache system administrators, application owners, and database administrators (DBAs) can control, tune, and report on SmartIO cache areas to increase performance and utilization of solid state devices, while keeping applications running. No volume, file system, or application tuning is needed to realize the performance benefits of SmartIO. Designed to be simple yet flexible, sfcache can be used to increase performance for both a wide range of data and the most granular targets with minimal interaction. With a single command, users can cache all existing and future volumes and file systems, while simultaneously choosing individual files to permanently reside, or be “pinned”, within the SmartIO cache area.

sfcache control
Create / Delete
Grow / Shrink
Enable / Disable
Pin / Unpin
Flush / Purge
Online / Offline
Stats

SmartIO also brings deep application insight around database workloads from Oracle or Sybase. With the app parameter within sfcache, DBA's of all levels of expertise can enable SmartIO caching control at varying levels within the application. Different databases and varying workloads will each benefit from different levels of data being cached, or in some cases, ignored as cache candidates. SmartIO can take insights from Oracle Disk Manager (ODM) and Automated Workload Repository (AWR) output to proactively move portions of Oracle data in and out of the cache, with no end-user interaction, to increase performance of targeted workloads. SmartIO works with ODM and caches transactional data, while ignoring RMAN activities (back-up, copy, restore) and internal reads, such as Archive reads.

SmartIO Benefit: Increased Application Performance

Traditional Online Transaction Protocol (OLTP) databases have different read/write patterns than Data Warehouse (DW) based applications. OLTP and DW also have different I/O characteristics on the back-end disk. Whereas OLTP is more random, DW is more sequential. Understanding those patterns is not only critical to placement of cached data, but smarter caching can increase the lifespan of Solid State devices through reducing the number of write operations.

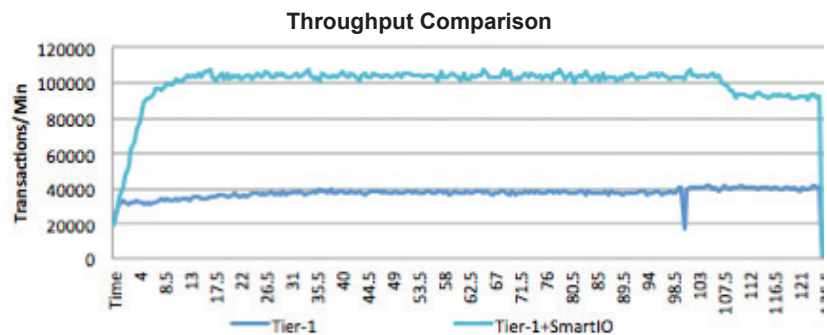


Figure 2. OLTP Comparison

For a mixed OLTP workload, SmartIO provides more than 2.5 times the transactions per minute when compared to a traditional “tier-1” array. Simply by adding any Solid State device (PCIe/SATA/SAS) from any Enterprise vendor into existing servers can yield immediate gains in application performance and productivity without re-architecting the storage or server environment.

SmartIO Benefit: Tier-1 Performance at Tier-2 Costs

While those gains are significant, the ancillary benefits may provide the largest value in Enterprise Data Center. To see where those secondary gains occur, you need to examine the typical mixed workload and HDD arrays. HDD Arrays have been architected to squeeze as many IOPs as possible out of a limited medium. The physical movement of the disks and read/write heads can only handle so many requests in a period of time.

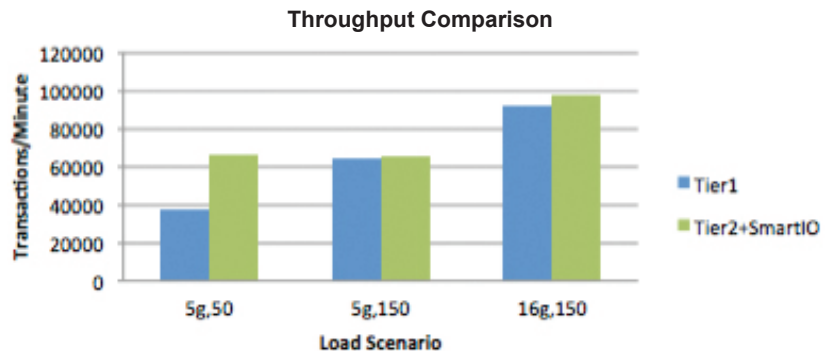


Figure 3. Tier 1 vs Tier 2+SmartIO OLTP Performance

With SmartIO, the majority of requests can be handled within the server. In cases where the workload size is less than the available cache size, a 100% cache hit ratio is a possibility. Handling reads within the server with SmartIO reduces the performance needs of the back-end array. Those same platters need to less seeking and less spinning to service I/O, especially in applications, like OLTP databases, that exhibit smaller, more random I/O patterns. Our scenario uses an EMC VMAX 20k array with 15k drives as the standard “tier-1” array. An EMC Clariion Cx4, also with 15k drives, serves as our “tier-2” array. The combination of SmartIO and cheaper SAN brings better performance across multiple OLTP workloads than our tier-1 array at 39% of the cost. SmartIO provides a key tool as data centers migrate and organizations continue to look for ways to cut cost, while minimizing the disruption to revenue generating applications.

SmartIO Benefit: Increasing Storage Utilization

As SANs have been more widely adopted over the past 20 years, storage administrators and architects have implemented new techniques to get more IOPs out of hardware that alone maxed out around 20,000. By putting bigger front-end cache inside the controllers, using “wide striping” across multiple arrays to minimize seek times, and keeping storage utilization rates low, the performance of high-end arrays were able to keep up with the increasing capabilities of the server and application. These practices, however, came at a cost of: un-used storage; increased floor space and cooling for large storage farms; and complex infrastructures to gain a few hundred or few thousand IOPs. As the performance capabilities of servers and applications continue to increase exponentially, the benefits seen from these methods are showing decreasing returns.

SmartIO reduces the I/O needs of the back-end storage infrastructure. Techniques to get more IOPS are no longer required. SmartIO drives up the utilization of existing arrays and decreases footprint and cooling needs. Users of SmartIO can make better use out of existing storage arrays by more than doubling the number of applications within the same storage infrastructure, without impacting the performance of the application.

Running a “tier-1” array with 4 attached servers, each running an OLTP workload, we see an average of 125,000 transactions per minute per server. When enabling SmartIO on each of those instances, the number more than doubles to close to 300,000 transactions.

The area of bigger interest is the second set of bars in figure 4. Running with the same storage array and OLTP workload, but doubling the number of servers, we see a near linear regression in performance as the average transactions per server drops to around 6000. By enabling SmartIO on those servers, it brings us back higher than the initial benchmark with just four servers. We have doubled our workload on the array, without impacting our overall performance improvement.

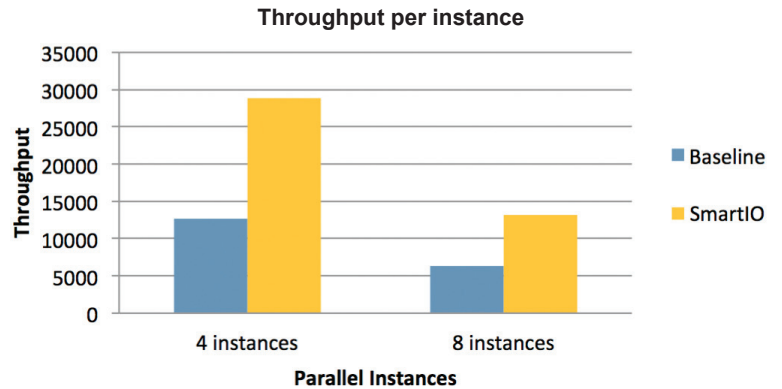


Figure 4. Instance Stacking

With SmartIO on the server, data center architects and storage administrators can consolidate applications on fewer arrays, driving up storage utilization and reducing data center costs associated with floor space, power, and cooling of running high end arrays while meeting stringent Service Level Agreements.

SmartIO Benefit: Cascaded Performance Increase

By servicing I/O requests inside the host, SmartIO not only increases performance of the front-end application and the back-end storage, but can also improve the performance of applications that do not run SmartIO, or have any internal SSD installed. SmartIO frees up the performance capacities of shared storage. This has a cascading performance impact on other servers that rely on that storage as those applications get more access to front-end cache and can take advantage of the reduced workload on the spinning disk.

Figure 5 above outlines a test workload of two servers, connected to a single array, across a shared set of spindles. Each node ran simultaneous OLTP workloads to create a bottleneck across the shared storage. “Node 1” shows a slightly higher increase in transactions as it contains more RAM and CPU than “Node 2”.

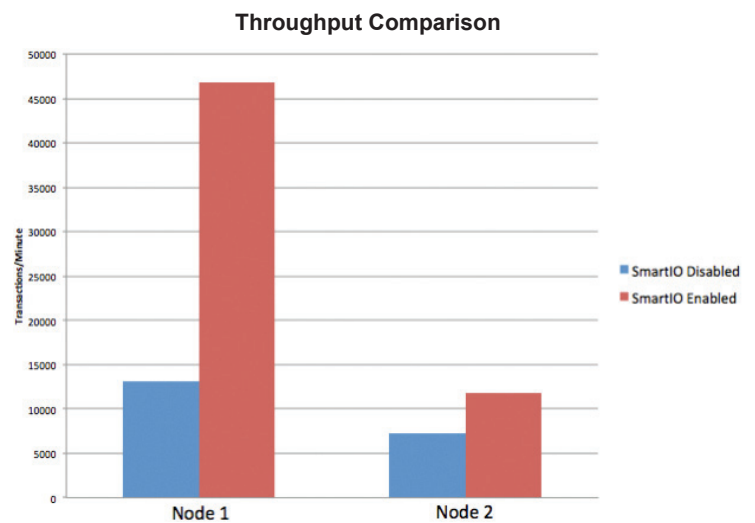


Figure 5. Cascaded Performance

After completing the OLTP workload, the same tests were run again, this time with SmartIO enabled on “Node 1”. The increased performance with SmartIO enabled, over 4x in this test, is expected with SmartIO’s intelligent caching and application insight capabilities. The performance increase for “Node 2”, however is not due to any changes on that node, but due completely to enabling SmartIO on “Node 1”. SmartIO has removed a large percentage of reads from the shared storage. Those performance benefits cascade through all nodes running on that storage. SmartIO and internal SSD can have impacts on systems that are never touched.

Summary

SmartIO in Veritas InfoScale Storage brings customers complete, online control of performance for their most mission critical applications. Caching in single and multi-node environments reduces the needs of expensive tier-1 storage, while increasing performance and utilization of legacy storage. SmartIO is integrated directly into existing Veritas File System and Veritas Volume manager technologies to allow customers to ensure availability of their applications and data with no downtime.

For more information and details on Veritas InfoScale Storage, see <http://www.veritas.com/infoScale>

Appendix

Figure 2 Details:

Hardware	
Server	Dell 820 (32 Cores, 256 GB, Hyperthreading on)
Storage	IBM DS8100 (16*Raid10 luns, 16GB cache, 64 spindles)
SSD	FusionIO IODrive2
Connectivity	2*4Gbps (8 path)
Software	
OS	Redhat Enterprise Linux 6.3
Veritas	InfoScale Storage (version 6.1, known as Storage Foundation)
Oracle	11.2.0.3
Setup	
Data file system	1.5TB (blocksize=8k,15luns)
Log file system	50GB(blocksize=1k,1lun)
OLTP	Users: 150; SGA: 7g

Figure 3 Details:

Tier-1 Configuration		
Host	Storage	
Dell 820 (Intel Xeon E5-4620 4*8 core with hyperthreading enabled @ 2.2Ghz) 256 GB RAM	EMC VMAX	Disks: 128* 15k rpm Luns: 12*Raid10 Cache: 16GB Controllers: 2 (1 engine)
Tier-1 Configuration		
Host	Storage	
Dell 820 (Intel Xeon E5-4620 4*8 core with hyperthreading enabled @ 2.2Ghz) 128 GB RAM	EMC Cx4	Disks: 80* 15k rpm Luns: 20*Raid0 Cache: 16GB
	Fusion-IO IODrive 2	1.2 TB
Setup		
OLTP	Users: 50; SGA 5g Users: 150; SGA 5g Users: 150; SGA 16g	

Figure 4 Details:

Hardware	
Server (x4)	Dell 820 (32 Cores, 256 GB , Hyperthreading on)
Storage	IBM DS8100 (16*Raid10 luns, 16GB cache, 64 spindles)
SSD	FusionIO IODrive2
Connectivity	2*4Gbps (8 path)
Software	
OS	Redhat Enterprise Linux 6.3
Veritas	InfoScale Storage (version 6.1, known as Storage Foundation)
Oracle	11.2.0.3
Setup	
OLTP	5000 warehouse (1.5TB)

Figure 5 Details:

Server	
Node 1	Dell 820 (Intel Zeon 4*8 core with hyperthreading enabled 2.2Ghz, 256GB, 4*4Gbps FC HBA)
Node 2	Dell 720 (Inter Zeon 2*8 core with hyperthreading enabled 2.2Ghz, 96GB, 4*4Gbps FC HBA)
Storage	
IBM DS8100 12 Raid10 luns (4 spindles, 4 path each)	
8*550 GB luns	
Software	
OS	Redhat Enterprise Linux 6.3
Oracle	11.2.0.3
Veritas	InfoScale Storage (version 6.1, known as Storage Foundation)
OLTP	5000 warehouse (1.3TB)

About Veritas Technologies Corporation

Veritas Technologies Corporation enables organizations to harness the power of their information, with solutions designed to serve the world's largest and most complex heterogeneous environments. Veritas works with 86 percent of Fortune 500 companies today, improving data availability and revealing insights to drive competitive advantage.

For specific country offices and contact numbers, please visit our website.

Veritas World Headquarters
500 East Middlefield Road
Mountain View, CA 94043
+1 (650) 933 1000
www.veritas.com

© 2015 Veritas Technologies Corporation. All rights reserved. Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

21357612 08/2015