DataCore SOFTWARE

The Storage Hypervisor Company

Virtualizing Business Critical Apps

Clearing the Performance Hurdle

ORACLE

SAP

Microsoft Exchange Server

Microsoft SharePoint

Microsoft SQL Server

TIPS FOR:

- Speeding up vitualized applications
- Removing I/O bottlenecks
- Eliminating storage-related downtime

INTRODUCTION

Stories of unmet performance expectations continue to deter IT organizations from virtualizing their business critical apps. The two most common complaints:

1. Programs run unacceptably slow after virtualizing
2. The response is too erratic to be useful

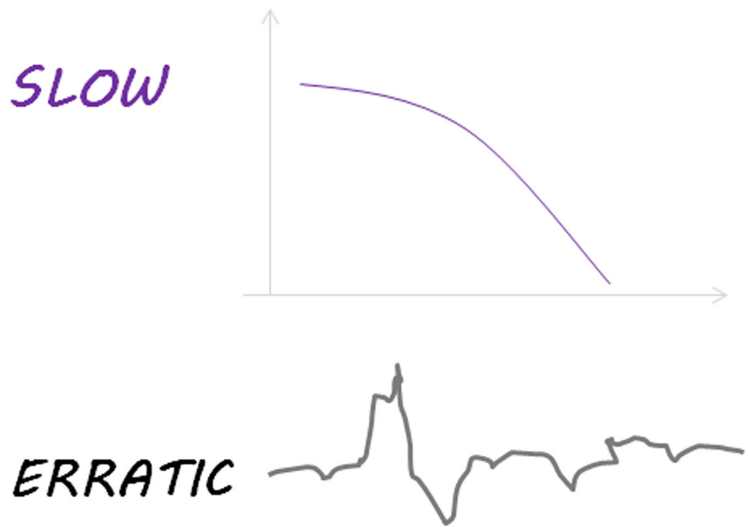

Countless cases with similar experiences cast doubt on the benefits of moving SQL Server, Oracle, SAP, Exchange and SharePoint from segregated servers into fewer, consolidated virtual machines (VMs).

In this paper we consider the root causes behind the unacceptable behavior and offer a practical, proven solution.

ROOT CAUSE:
CONTENTION FOR SHARED RESOURCES

While apps may display several symptoms indicative of slow or erratic response after being virtualized, the problem boils down to contention for shared storage resources; contention that did not occur when the apps had the storage all to themselves.

These so called "bottlenecks" occur in spurts as application requests collide randomly, resulting in spikes of sluggish, unpredictable latency. The more frequent, the greater the users' dissatisfaction. You may recall that one of the primary reasons these business critical apps were originally sequestered on separate physical machines was to avoid such collisions.

OLD WORKAROUNDS PROVE EXPENSIVE

In scenarios where each physical server hosts a single app, one can readily isolate bottlenecks and troubleshoot them. Traditionally, the remedy involved spreading the load across more disk drives using additional disk channels. The technique, known as "short stroking", places data only on the fastest edges of the disk platters, thereby limiting the distance (or stroke) the disk actuator travels radially. Reducing the stroke reduces the mechanical delay in getting to the data.

What's the tradeoff? The rest of the disk platter goes unused, wasting much of the drive's capacity. Often more than 40% of the disk space is left vacant in favor of quicker response.

The premium priced caches on the storage controllers that front-end the short stroked drives are also maxxed out to improve speed.

Depending on how the storage system is packaged, the limits on disk drive slots, channels and cache memory force one to buy additional underutilized disk arrays in an effort to attain adequate response.

Sure sounds like the same problem we had with poorly used servers before virtualization.

More recently, very fast solid state disks (SSDs) and flash memory cards have been incorporated inside servers in order to speed up I/O performance. The approach significantly reduces the number of Hard Disk Drives (HDDs) by factors of 10x or more. In general, SSD/Flash technology turns around read requests very quickly, but is not optimized for write requests.

SSDs work nicely when dedicated to a physical server. However, two new challenges arise:

- How to share the premium resource across multiple physical servers to amortize their cost?
- How to replicate their contents to prevent them from becoming single points of failure?

COST IMPLICATIONS

Although technically feasible, the money spent customizing a specialized storage infrastructure using the methods described above often exceeds the projected value of virtualizing and consolidating tier 1 apps.

Furthermore, once you arrive at a stable configuration, any major changes or expansion requires additional, expensive customization. The poor economics force companies to leave the apps alone.

Interestingly, new storage virtualization technology has been developed to significantly shave costs and complexity from the solution. These software innovations make the transition to virtual machines viable. Let's see how that's done.

SOLUTION:
HIGH-PERFORMANCE STORAGE HYPERVISOR

The solution described below employs DataCore's SANsymphony™-V storage hypervisor software. This new member of the virtualization stack (complementing server hypervisors) is responsible for circumventing the physical constraints of storage devices, while making the best use of their capacity and connectivity. The solution has been shown time and time again to overcome the storage-related bottlenecks encountered when virtualizing Tier 1 apps, and to do so affordably.

One measure of its effectiveness comes from independent lab tests (see: **ESG Lab Validation Report**) comparing the maximum number of Exchange users hosted before and after DataCore virtualized the storage in controlled environments. These standardized benchmarks reveal an impressive 5X performance boost, which grows to 6X for file sharing workloads.

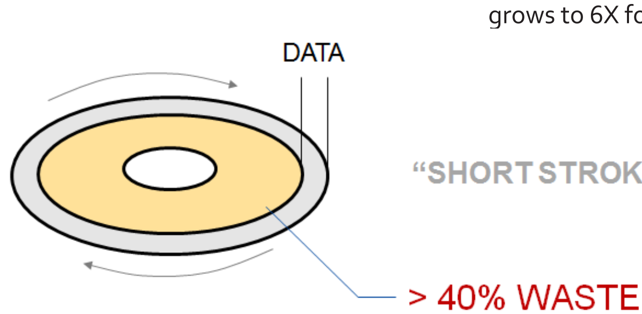

Perhaps a more representative number comes from the large number of customers reporting 3x to 5x faster speeds across their production systems spanning Oracle, SAP, SQL Server, Exchange and SharePoint.

How's that possible, especially when layering software between apps and storage? It has a lot to do with how those disk resources are leveraged, and where server resources are redirected.

HARNESS SERVER RESOURCES DIFFERENTLY

Essentially, the breakthrough comes from employing CPUs, memories backplanes and channels differently from the standard way that the base operating system and server hypervisor use them. In effect, these server components are assigned I/O tasks, rather than the processing activities usually associated with them. In mid to larger IT environments, servers (nodes) are dedicated to the storage virtualization role, whereas in smaller configurations, the services co-reside with the virtualized apps on the same physical hardware.

The innovative strategy produces outstanding results on several aspects:

- Substantially speeds up response when reading and writing to disk
- Consistently accelerates throughput to realize the desired Transactions per Minute (TPM) and IOPS (I/O Operations per Second) expected from the virtualized apps.
- Significantly reduces the cost to achieve both the predictable behavior and continuous availability that Tier 1 apps demand.

Next we consider those technical factors having the most influence on performance and resiliency.

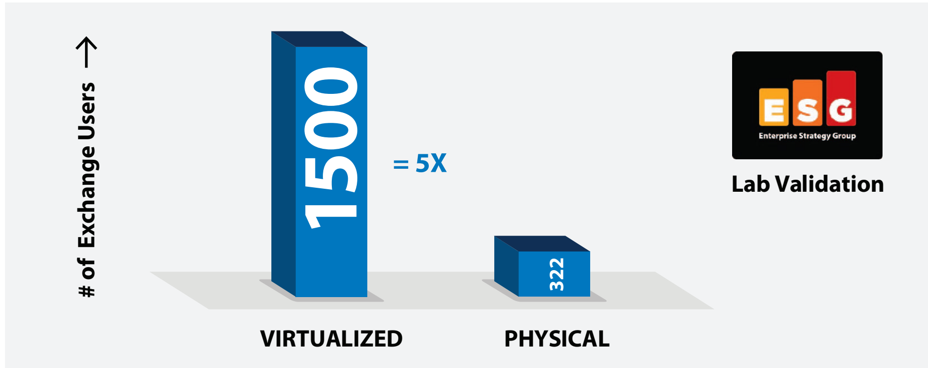

INFRASTRUCTURE-WIDE CACHING

The advances we discuss here are not so much in new caching algorithms. Instead, they are in how those well-known techniques, including write coalescing and read ahead, speed up read and write requests to disk when applied over a wider scope of resources.

Rather than limiting caches to the narrow confines of a disk array, DataCore software caches across pools of storage devices consisting of multiple, independent storage systems. The caching hierarchy works much like how web caches accelerate browsing across numerous Internet servers- essentially extending local caches closer to the user in a cost-effective manner.

The pools typically consist of various storage classes covering a range of price points. They often encompass a combination of the fastest SSD/Flash technology, very high performance HDDs and low cost, high capacity drives, possibly from different manufacturers.

Clients find that SANsymphony™-V can manage very large working sets of several virtualized Tier 1 apps in its large distributed caches without incurring costly delays for accessing the back-end disks.

This is true with OLTP loads as well as in Business Intelligence (BI) and messaging / collaboration use cases. Incidentally, VDI (virtual desktop) scenarios display similar characteristics.

Adaptive caching across storage devices results in:

- Faster response from electronic memory (DRAM)
- Better economies of scale without the need to overprovision (short stroke) spindles
- Best use of SSD/Flash where this technology is appropriate

AUTO-TIERING ACROSS STORAGE SYSTEMS

DataCore software enhances performance in other ways. It automatically migrates disk blocks to the most appropriate class of storage systems to ensure the best results for the money. Auto-tiering optimizes the use of premium space without depending on system administrators to make absurd allocation guesses.

At the same time, less demanding workloads are automatically redirected to the lower cost drives.

Heat maps in conjunction with auto-tuning within tiers further load balance disk blocks among the available resources, thereby avoiding back-end bottlenecks.

MIRRORING BETWEEN FAULT DOMAINS

As mentioned earlier, it's not just about being fast. Keeping the apps running smoothly, and uninterrupted is equally important.

Like caching, mirroring disk updates between different storage systems for redundancy is not new. However, shop around for products capable of synchronously mirroring between separate SSDs or high-end disk arrays located in different locations and you will be in for big sticker shock.

Nevertheless, these precautions are necessary to eliminate single points of failure in the shared storage infrastructure where the occasional facility glitches and technician errors can quickly cripple business critical operations.

Here again, the SANsymphony™-V storage hypervisor takes responsibility-economically replicating disk changes in real time from location A to location B; each in a separate fault domain. Synchronous mirroring occurs between infrastructure-wide caches to:

- Achieve multi-cast stable storage between locations on separate ends of a metropolitan area
- Minimize the round-trip latency before I/Os are acknowledged to the app
- Allow later de-staging to disks without impacting application performance

SUMMARY

The combined benefits of DataCore's adaptive caching, auto-tiering and synchronous mirroring software implemented outside the box, so to speak, provide the desired high performance response and reliable behavior from virtualized tier 1 apps on which the business can depend. Though not discussed here, the collective capabilities also help you surmount the difficulties introduced when transitioning from physical servers to virtual servers, since the solution can simultaneously provide robust shared storage for both.

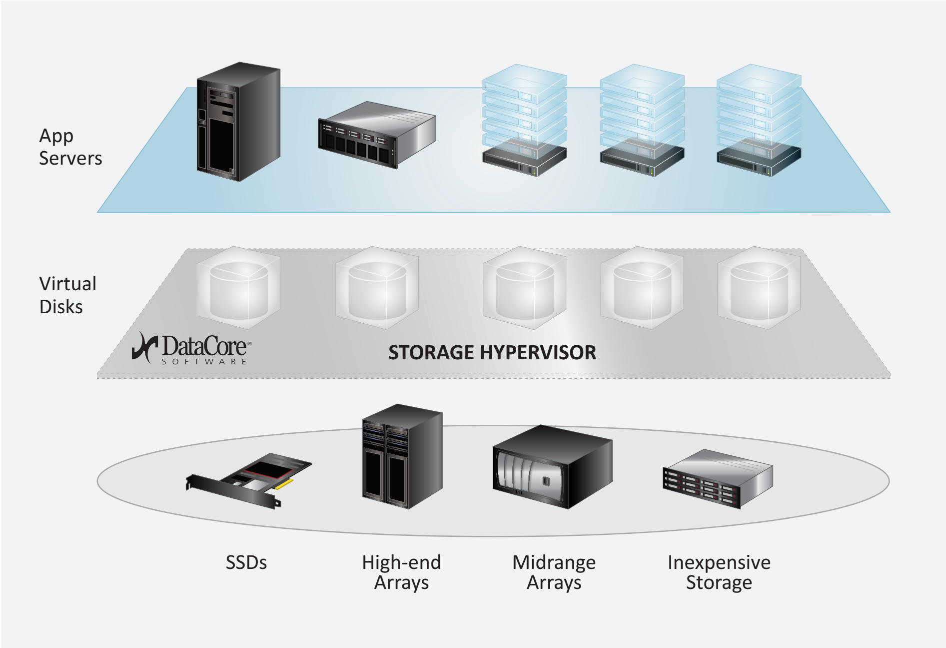

0912

For additional information, please visit **www.datacore.com** or email **info@datacore.com**