



Case Study

***How to Improve Productivity
by Reducing VMware IOPS 31%***

The Client

Savills PLC is a global real estate services provider listed on the London Stock Exchange and a constituent of the Financial Times Stock Exchange (FTSE) 250. The company has an international network of more than 600 offices throughout the Americas, the UK, continental Europe, Asia Pacific, Africa and the Middle East.



Savills offers a broad range of specialist advice in the areas of building and environmental consultancy, valuation, landlord/tenant planning and strategic projects and research in conjunction with commercial, residential and agricultural properties. Their activities include property-related financial management and investment management for institutional and professional investors. Their staff of 27,000 is well-versed to handle most any kind of real estate management or funding transaction.

The HW/SW Environment

The hardware/software environment includes 20 ESXi hosts with 340 VMware virtual machines on VMware vSphere 5.5 Enterprise Edition. The virtual machines support SQL, SharePoint, Exchange and application and file servers. The operating systems cover everything from Windows 2003 to Windows 2012 R2 Server.

The Problem

In the spring of 2015, the Savills IT organization had a major issue that needed immediate attention:

VMware virtual machines were experiencing performance issues: Higher-than-normal disk latency that was impacting overall performance and extending backup windows into production hours.

Two of the classic virtualization performance problems are HBA-LUN queue contention and disk latency and Savills was being affected by both. Queue contention is caused by excess SCSI commands clogging the queue while disk latency problems occur when the controller issues an overwhelming number of physical I/O to the SAN and the disks can no longer respond in acceptable times. The two problems go hand in hand since the number of SCSI commands directly affects the disk I/O count. VMware contends latency of 15ms should be monitored and over 30ms it is a cause for concern. VMware's Storage I/O Control (SIOC) actually throttles the LUN queue when latency hits 30ms thereby trading performance for better latency¹.

¹ Scott Drummonds (VMware) presentation at 2010 EMCWorld

Savills noted their disk latency was affecting SQL, SharePoint and Exchange performance. Average latency was 15ms but frequent peaks of more than 100ms were affecting productivity and causing long overruns on backups.

The IT staff realized that the situation was likely to worsen unless the Windows guest systems were properly maintained.

As with physical servers, virtual servers are adversely impacted by file and free space fragmentation in the guests. Windows accesses files by referencing “chunks” (fragments) of logical address space (LCN’s) reserved for the file.

Each fragment becomes a SCSI command from the guest to the hypervisor and controller as shown in Figure 1 below. A file in 100 fragments needs 100 SCSI commands to move it to the controller.

The controller breaks down each SCSI command into smaller disk I/O on the SAN. As Windows file fragmentation increases the SCSI traffic increases and there is a corresponding increase in disk I/O. It is the high volume of disk I/O that increases latency as the disks can no longer respond at their rated times.

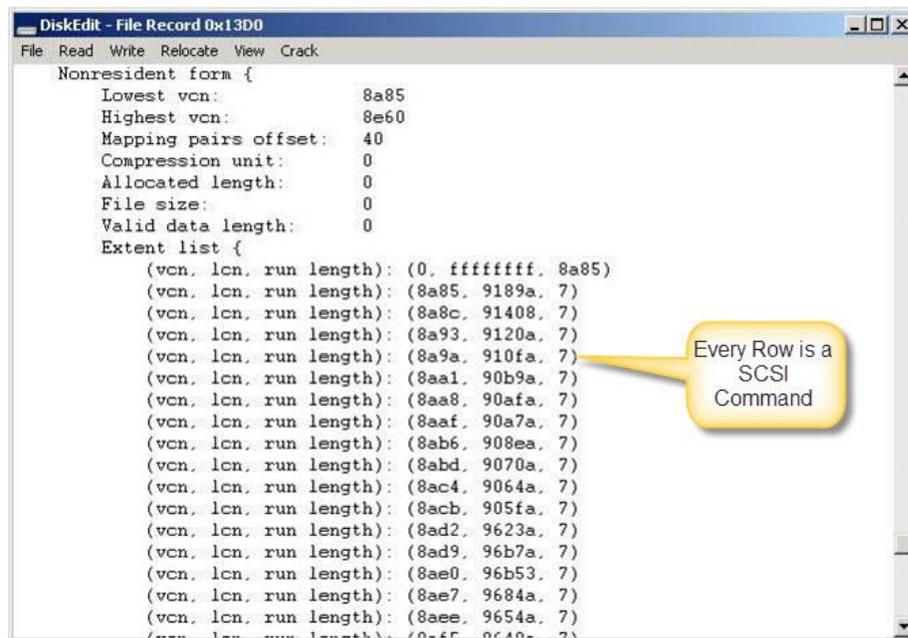


Figure 1. List of Fragments Inside File Record in the MFT

In order to take control of the latency issue Savills needed a comprehensive virtual disk defragmentation program that would reduce the total SCSI traffic and correspondingly reduce the physical I/O demands on the array.

The Solution

VMware Performance Issues and Expanding Backup Windows

VMware offers nothing to optimize virtual guest systems suffering from file and free space fragmentation. In terms of defragmentation, the Windows defragmentation tool was the only non-commercial option and it did not meet Savills requirements for a robust disk management solution.

Savills wanted a solution that was easy to implement, non-intrusive, with the ability to be scheduled to run at specific times, central management and the ability to produce both alerts and reports. In a VMware environment there are unique characteristics that need to be considered. **They opted for Raxco's PerfectDisk vSphere which is specifically tailored for VMware implementations.**

Daran Clarke, the project manager for both the reclamation and the defragmentation efforts, was looking to automate as much of the process as possible: **"PerfectDisk had the features we needed to implement a continuously scheduled defragmentation plan specifically for VMware."**

Solution Rollout Plan

Due to the large number of virtual machines involved and the nature of the workloads they supported Savills designed an implementation plan that started with the least critical systems and worked its way to the most critical systems.

This approach allowed them to create proven procedures that would be well-practiced by the time they got to the critical systems and they would be thoroughly familiar with PerfectDisk by the end of the project in six months' time.

PerfectDisk was deployed and run on this basis for the term of the project. At the project's end PerfectDisk was fully operational.

Results

PerfectDisk vSphere was run on the 340 VMs with impressive results.

- Prior to any defragmentation HP monitored the array for several days with Performance Analyzer and it averaged 50,202 IOPS, after the PerfectDisk defragmentation HP monitored the same LUNs and the average had dropped to 34,505 IOPS.
- Latency dropped from an average of 15ms to 10ms and peaks dropped from around 100ms to 20ms, ending complaints about SQL, SharePoint and Exchange performance.
- VMDK backup windows were reduced 43 hours.*

This means Savills was doing the same work with only 2/3 of the I/O. This 31% reduction in SCSI commands meant less resource demand on the guests, the hypervisor, controller and the disks. Latency and throughput improved.

*Backup windows were reduced 43 hours, originally completing Tuesdays at 10AM and now completing Sundays at 3PM but there were a number of factors at play in addition to running PerfectDisk vSphere because Savills was backing up at the VMDK level, including the reconfiguration of a storage management tool which had just about every property set for data collection. However, PerfectDisk has a history of significantly reducing backup windows for virtual and physical servers.

Return on Investment

PerfectDisk ROI: 31% Reduction in I/O Across Virtual Guests, Noticeably Improving Performance and Helping Reduce VMDK Backup Window

PerfectDisk vSphere effectively eliminated almost 1-in-3 of the total SCSI commands coming through the system, improved SQL, SharePoint and Exchange performance and helped reduce the VMDK backup window by 43 hours, now completing outside of production hours. In terms of resources this is the equivalent of getting another 110 VMs. The 31% performance improvement across all the guest systems affords Savills a number of options in terms of VM growth, number of VMs per host and load distribution across all the hosts. The current hosts will be able to support more virtual machines in their current configurations for some time to come.

How Savills Uses PerfectDisk Today

PerfectDisk vSphere is installed on all VMs including build templates, scheduled to run each weekday at 12am and runs for a maximum duration of 3.5 hours as well.

Making PerfectDisk a part of routine disk maintenance will prevent the kind of poor latency issues Savills experienced at the beginning of the project.

