

TIPS AND TRICKS FOR OPTIMIZING STORAGE MANAGEMENT

**Learn eight best practices to
optimize your storage and
become a great IT pro.**

Quest



Table of Contents

Introduction	3
Tip 1: Control VM sprawl	5
Tip 2: Consider solid state storage	7
Tip 3: Proactively manage capacity	8
Tip 4: Leverage thin provisioning and deduplication	9
Tip 5: Obtain storage statistics	12
Tip 6: Ensure Optimal Storage Performance	13
Tip 7: Track the right stats for your storage environment	14
Tip 8: Analyze performance and capacity needs	15
Conclusion	16



Introduction

The first and last consideration for any virtual environment is storage management. Like George Carlin so entertainingly observed, you need a place for your stuff. In this case of a managing a virtual data center, all that stuff is servers, operating systems, applications and file storage. Ensuring you have just the right amount of storage, both for the overall infrastructure and for file storage as well, can be a complex undertaking on the best day. You have to balance function and performance against having so much capacity that it impacts your department's profitability. A good IT pro can attempt to find this balance without tools that tell you minute by minute where storage should be allocated, but a great IT pro knows exactly what's happening in their virtual environment from cluster to LUNS and uses this knowledge to deliver an extraordinary, high-performance, low-cost experience.

“Maybe put some of your stuff in storage. Imagine that – there’s a whole industry based on keeping an eye on your stuff!”

- George Carlin, A Place for Your Stuff

Want to be a great IT pro? Since you're reading this e-book, it's likely you already are. Well, embrace your greatness because it's about to go to the next level. Read on to add eight tips to your bag of tricks. What you're about to discover in this eBook will help you optimize your storage and make even the most complex virtual data center sing with efficiency.



Storage Resource Planning



Tip 1: Control VM sprawl

THE IMPACT OF UNUSED AND DISUSED RESOURCES


Any idea how much unused resources you have in your virtual data centers? Waste in virtual data centers is particularly insidious because it's challenging to even locate. Virtual machines (VMs) that are no longer being used, powered-down VMs that haven't been touched in months, templates that are no longer relevant and snapshots that just keep aging — these are all examples of virtual data center waste. All consume expensive software, hardware and storage resources, and all will continue to accumulate unless you have a plan of action for orchestrated waste removal. Eventually your virtual environments will become overpopulated with VMs of dubious value, and VM sprawl will impact both performance and operating expenditures.

Spend some time implementing policy controls, as well as doing an audit of resources and eliminating underutilized and disused VMs, CPUs, memory and especially storage capacity. Then take your storage optimization to the next level to realize reduced expenses by controlling VM sprawl using VM lifecycle management technology. Putting automated policy controls and governance in place delivers visibility into the lifecycle of all VMs. Plus, it enables businesses to be in strict compliance with IT policies and standards.

KEEP TRACK OF RESOURCES

The right solution keeps track of resources and flags those that are prime for retirement, which in turn facilitates better planning. The right technology will alert you to retirement requirements across your virtual environment, helping you reclaim productivity and capacity from:

- **Abandoned VM images.** An abandoned VM is one that has been removed from inventory but that still remains in the virtual environment, consuming resources.
- **Powered-off VMs.** VMs that have been powered off for extended periods but remain in inventory are also excellent candidates for retirement.
- **Unused template images.** Today's VMs are almost always created from a template, and most virtual environments require more than one template. That said, templates that haven't been used for extended periods are good candidates for retirement.
- **Snapshots.** VM snapshots offer another opportunity for savings — particularly snapshots that grow over extended periods of time, getting larger and slower. An effective toolset will monitor for changes to snapshot files and alert you when a snapshot has been running for an excessive time.
- **Zombie VMs.** A private cloud's self-service VM-creation functionality offers incredible automation potential. But this can come with a price: resources allocated to disused VMs. By analyzing resource consumption over time "zombie" VMs can be easily identified.
- **VM sprawl.** Because VMs are so easily created and replicated, many companies experience an increase in VMs that are unapproved, unlicensed or unused. A company can recoup tens or hundreds of thousands of dollars in wasted resources by simply cleaning up their virtual environment.



RIGHT-SIZING YOUR VIRTUAL INFRASTRUCTURE

Virtual infrastructure optimization is important because VM resource requirements can be difficult to determine and generally vary over time. For example, if VMs are overprovisioned, resources – such as storage – that could be used by other VMs are wasted. Overprovisioning also wastes your investment in server capacity. If VMs are under provisioned, the applications running inside will perform slowly. When a VM is at its optimal resource configuration, it is considered “right-sized.”

Virtualization optimization can save enough in recouped resources to, in effect, pay for the virtualization optimization tool. New VMs can be added without needing to purchase additional servers, or VMs can be consolidated onto fewer servers, allowing other servers to be returned on lease or shut down.

Virtualization optimization can save enough in recouped resources to, in effect, pay for the virtualization optimization tool.

Tip 2: Consider solid state storage

WHAT IS SOLID STATE STORAGE AND WHY IS IT IMPORTANT?

Solid state disks are transforming the storage market more than just about any other factor today. For decades, hard drive technology has been the primary data center storage option, despite its lack of performance gains over the years. Virtualization administrators had no other options but to build their virtual environments around those inherent performance limitations.

Now solid state storage has become a viable option for many, if not most, data centers. While their price point has dropped, solid state disks remain the most expensive type of storage when it comes to capacity. But when performance is the primary metric, solid state disks can be surprisingly affordable. Solid state disks have low latency and very good random access performance because they have no moving parts.

To get comparable performance from rotating disks, you would need to buy dozens or hundreds of disks. When it comes to overall reliability, solid state disks beat rotating storage, hands down. Also, with no moving parts, there is no physical wear and tear or mechanical breakdowns that can take place. That's why solid state storage can withstand harsher environments than rotating storage,

too. In addition, without the need to continuously spin a disk and move a drive head, solid state disks are also more power efficient than rotating disks.

SOLID STATE STORAGE INNER WORKINGS

It is common knowledge that rotating disks work by spinning magnetic platters at high RPMs and then suspending read/write heads on a cushion of air above those platters. Solid state disks, on the other hand, use a cell-based structure that uses electricity to continually erase and reprogram individual data cells to store information. There are two primary types of solid state storage and each carries with it distinct characteristics and use cases (see table below).

So evaluate your business needs and consider whether the performance benefits of solid state storage outweigh the overall cost, size and lifespan limitations. In some cases, it definitely will.

	Description	Performance	Capacity	Cost	Use cases
Single-level cell (SLC)	SLC solid state disks store a single yes/no value per single layered cell (one bit of data).	Because of this binary, single value system, performance is exceptional.	The single value per cell nature limits the capacity potential for SLC solid state disks.	SLC solid state disks are the most expensive type of solid state storage available.	For situations where speed rules; big data, large VDI installations, analysis
Multi-level cell (MLC)	MLC solid state disks can store multiple states per cell — in general, up to two bits per cell.	Because the system spends more time analyzing each cell to determine its value, MLCs are slower than SLCs.	MLC storage boasts much greater capacity than SLC because each cell can hold more data.	MLC solid state disks are much less expensive than SLC and are commonly used in storage arrays.	General purpose line-of-business applications, mainstream applications

Tip 3: Proactively manage capacity

Clusters are the foundation of VM performance and high availability. Clusters ensure that resources get used efficiently. But they can create problems when they're not planned and implemented correctly.

To protect each of your VMs, you should set aside one host's worth of resources in reserve. These resources are to reduce disruption if there's a cluster host failure and standby ready to process the VMs that existed on the cluster's lost node. Some clusters are built without the necessary reserve. Many more lose it in the face of unexpected VM growth or due to strapped hardware budgets.

This is where planning plays a critical role. That way you've mapped out cluster resources to be reserved as failover capacity. Just keep in mind that this will change as your cluster grows (or shrinks): Add a new cluster node and you might reduce your reserve percentage. Add more VMs and that percentage might increase. Modify resources on existing VMs and your percentage could go in either direction. Maintaining the best reserve balance requires vigilance in a dynamic environment.

The right solution provides performance data and user-defined thresholds at all levels of the virtual infrastructure to help you see and predict how efficiently capacity is used. Also look for a solution that can help you determine additional VM capacity and monitor thresholds with cluster-based capacity utilization. Most important, make sure that the operations management solution does more than simply alert you to a problem — you also want single-click remediation of issues to save you time and keep your environment running smoothly.



Tip 4: Leverage thin provisioning and deduplication

THIN PROVISIONING

Thin provisioning in virtual environments can take two forms – at a hardware storage array or at the virtual disk – with the most common being the array. When thin provisioning is done at the array level, administrators can make better use of storage. For instance, when storage is allocated “thick,” the array sets aside all of the storage defined at the time (for example, 100 GB) and allows the file system or database to grow into it. However, any storage that goes unused, or takes time to use, the “thick provision” is used. This also eats up storage that could be used by other applications. For example, if only 5GB will be used initially, 95GB goes unused, and is wasted.

However, a “thin provision” would allow 100GB to be created, only using this physical space as needed while freeing up storage on the array to be used by other applications. Thin provisioning also enables the administrator to allocate a much larger provision, such as 1000 GB, in order to prevent the file from running out of space. That said, care must be taken not to over provision the array, otherwise, you could run out of storage.

The second, form of thin provisioning is at the virtual machine disk (VMDK) level, and is basically the same concept as array level thin disk provisioning, but is administered within the hypervisor. However, many argue it is best to use thin provisioning at the array level vs. the virtual

file system/disk to make better use of physical storage and possibly improve performance

The following are considerations for thin provisioning:

- **Physical disk space constraints.** Thin provisioning can effectively allow an administrator to allocate more storage capacity to VMs than is physically available in the storage array. Using thin provisioning means that administrators need to be vigilant in their monitoring virtual disk space.
- **Potential performance penalty.** As a VM grows and needs to request more physical space, it takes a little bit of time for the array to locate and prepare that space, meaning that, in some instances, thin provisioning may introduce a small amount of additional storage latency.

DEDUPLICATION

For organizations that want to maximize the return on investment for their storage environments, deduplication is a must-have feature that can reap major benefits. Here’s how it works: When the deduplication process is running, the deduplication engine keeps watch over every block and every file that comes in and attempts to match patterns with existing data. When a match is found, rather than simply saving a second copy of the information, the deduplication engine simply writes a pointer to the original copy of the data. Every time this happens, storage capacity need is reduced and, over time, the savings can be substantial. Take, for example, a virtual desktop environment in which the storage houses multiple copies of hundreds of desktop computer images. In these kinds of situations, deduplication can save organizations large sums of money in capacity savings.



There are two kinds of deduplication:

- Inline. From a capacity-savings perspective, inline deduplication is the most efficient. Although this method can also introduce some latency in storage operations. As the data passes through the storage array, the array's deduplication engine immediately jumps into action and deduplicates the data on the fly.
- Post-process. Some systems – such as Microsoft's new deduplication engine – use post-process operation. In this kind of deduplication, the array simply operates as usual, saving multiple copies of data in routine operations. Then, on a predetermined schedule, a process kicks off during which the array performs the deduplication. This prevents any ongoing performance issues, although it also means that there needs to be enough capacity to initially store all copies of the data.

For virtualization administrators, deduplication is important to understand. Its use can effect overall storage design and performance, depending on the application.

Storage Resource Performance

Tip 5: Obtain storage statistics

Because virtual infrastructures are used to consolidate servers more and more densely, the majority of performance issues stem from storage I/O contention. Any tool that analyzes virtual infrastructure performance and capacity would be incomplete without some kind of storage I/O and capacity analysis piece. The source of raw storage statistics also is an important consideration. You want your virtualization management tool to obtain storage statistics from your virtualization management platform; however, it's also valuable to obtain storage statistics directly from your physical storage systems as well.

More advanced virtualization management tools can directly access your storage arrays and storage fabrics using industry standard protocols such as CIM and SIM-S. Many virtualization management tools that manage infrastructures to report storage properties, latencies and availabilities. If your virtualization management tool can go directly to your storage, you can easily identify configuration errors, resource starvation issues and performance bottlenecks from the firsthand perspective of the storage system.

If your virtualization management tool can go directly to your storage, you can easily identify configuration errors, resource starvation issues and performance bottlenecks from the firsthand perspective of the storage system.




Tip 6: Ensure Optimal Storage Performance

Physical server networking is usually in good supply — physical servers rarely exceed a fraction of their available bandwidth. The same isn't true for virtual hosts. Consolidate dozens of VMs onto a single host, and those individual fractions start adding up. The right storage configuration can make or break infrastructure performance and availability. Visualizing what's going on is arguably the biggest challenge with storage connections. Even today, it remains a complex task to glean relevant information from storage performance statistics.

Add networking or disk spindles and rebalance VM processing — or do all three — to free up storage connections. In addition, keep track of how storage is impacting the virtual environment. You'll either have to become a storage IOPS master or rely on technology to get visibility from the VM down to the physical disk spindle. A unified and logical end-to-end view of the virtualization infrastructure provides the information you need to understand how separate systems impact overall VM performance.

Look for a storage monitoring tool that proactively highlights storage bottlenecks and delivers one-click troubleshooting to quickly identify and remediate issues. This clear and linear visualization from the VM to the physical disk will put an end to cross-team finger-pointing when there's a performance issue and ensure optimal storage performance.



Tip 7: Track the right stats for your storage environment

In addition to measuring physical storage capacity, it's critical to track various performance-related metrics. When storage performance becomes problematic, the workloads running on that storage area network (SAN) may behave erratically and affect the business. As a virtualization administrator, there are a couple important performance statistics to track vigilantly. They are listed below. Beyond these statistics, consider what will serve you best in your unique environment.

LATENCY

Latency is the time it takes for an entire storage operation to take place. High latency values – above 20 to 30 ms – can create problems for workloads running on the storage. There are many opportunities for latency to be introduced into the storage equation, such as the time it takes:

- For the hypervisor to process a storage command
- To transfer the command over the storage link
- Disks to spin around to the location on the disk
- For the system to read or write the data from or to the disk
- To transfer the data over the storage link back to the host

IOPS

Perhaps the most well-known storage metric is the number of Input/Output Operations per Second (IOPS) that can take place. This metric describes how many read and write operations the storage can handle before running out of resources. When it comes to virtualization, IOPS is important because it dictates how much storage performance capacity the storage environment can withstand. There are a number of factors that have an impact on how many IOPS a storage system can sustain:

- **RAID level.** If you're using RAID, make sure you understand that some RAID levels impose major IOPS penalties to support the protection afforded. You will learn about the impact of various RAID levels later in this paper.
- **Media type.** If you use spinning disks, you'll get anywhere from about 70 IOPS per disk up to about 200 IOPS per disk. As you add more disks, you get more IOPS. If you decide to move to solid state disks, your storage will be able to support thousands or even tens of thousands of IOPS

When storage performance becomes problematic, the workloads running on that storage area network (SAN) may behave erratically and affect the business.

Tip 8: Analyze performance and capacity needs

To help you understand the performance of your VMs, as well as your overall hardware, storage and virtualization capacity utilization — historically and in real time — but only when you have the right management tools in place to deliver deep insight into how these infrastructures work. VMs are dynamically sized containers that move from one server to the next and from one storage container to the next. Their resource demands fluctuate on a daily, weekly, monthly or even yearly basis. Virtualization management tools must take all of this into account, know what a normal baseline is and use that knowledge to help predict and prevent potential performance problems and capacity shortages before they happen.





Conclusion

You and the users of your virtualized environment need a place to put all your stuff. With the proper tools, an optimized, cost-effective virtualized environment with optimized storage is within your grasp. By heeding the eight tips detailed in this e-book and choosing the right storage optimization tool, you can become a great IT pro.

For more information about Foglight for Virtualization, please visit:
www.quest.com/virt

ABOUT QUEST

Quest helps our customers reduce tedious administration tasks so they can focus on the innovation necessary for their businesses to grow. Quest® solutions are scalable, affordable and simple-to-use, and they deliver unmatched efficiency and productivity. Combined with Quest's invitation to the global community to be a part of its innovation, as well as our firm commitment to ensuring customer satisfaction, Quest will continue to accelerate the delivery of the most comprehensive solutions for Azure cloud management, SaaS, security, workforce mobility and data-driven insight.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our Web site (www.quest.com) for regional and international office information.

© 2016 Quest Software Inc. ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at www.quest.com/legal.

Trademarks

Quest, and the Quest logo are trademarks and registered trademarks of Quest Software Inc. in the U.S.A. and other countries. For a complete list of Quest Software trademarks, please visit our website at www.quest.com/legal. All other trademarks, servicemarks, registered trademarks, and registered servicemarks are the property of their respective owners.