

Good to Great: The Top Seven Optimization Tips and Tricks



Introduction

Virtualization is one of IT's most exciting technologies in a generation. The ability to combine software and hardware in a way that allows a single server to act like several servers has provided IT with vast benefits, from saving on capital expenses and energy, to faster application deployment and workload agility.

With rapidly advancing technology, power and flexibility comes a multitude of process changes and complexity. Mastering your virtual data center requires comprehending a dizzying array of software and hardware configurations, settings and workflows.

With so many combinations, dialing in a virtualized data center's performance has become something of an art form. A good IT pro can generally make it run, but a great IT pro can deliver an extraordinary performance and availability, while minimizing operations expenditures.

The difference is proactive visibility and analytics. Good virtual administrators know which configurations will allow an application to run, but the great ones have a holistic understanding of the entire infrastructure and the optimal configuration for performance assurance, while minimizing cost and maintenance.

Managing a virtual data center—with all its servers, operating systems, storage, applications and more—is a complex undertaking. A good IT pro can make it work, but a great IT pro can deliver an extraordinary, high-performance, low-cost experience. This whitepaper provides the information you need to go from good to great as quickly as possible, with the top seven tips for optimizing even the most complex virtual data center.

You can fast-track your path to virtual greatness by avoiding a few key mistakes and implementing smart technologies in just the right places

How do you get from good to great? One way involves experience—and a lot of it—adjusting configurations, making mistakes and analyzing completely the operational impact of each decision. Problem is, most of us don't have that kind of time; we need to be experts today. The best way to achieve that is to cheat.

Or, in better terms, to leverage best practices and the experiences of experts through technology. Virtualization has gained ridiculous maturity in just a few short years. As a result, you can fast-track your path to virtual greatness by avoiding a few key mistakes and implementing smart technologies in just the right places.

The top seven optimization tips and surprisingly common corresponding mistakes are described below. See how your approach—and your optimization techniques—stack up.

Optimization tip #1: Right-size VMs for optimal performance and maximum VM density.

Common mistake:

Creating “big” VMs (because you can)

Virtualization's early years started out with a bang—an explosion of VMs that were all too often a copy-and-paste from a physical computer. Physical-to-virtual (P2V) technologies fork-lifted many data centers into the virtualization experience, converting physical machines bit by bit into virtual ones.

Yet all of virtualization's benefits arrive with a singular caveat: virtualization works because resources are shared. Wasted resources on overpowered physical servers were an unfortunate cost of doing business; wasted resources on overpowered virtual servers create contention. Even vSphere's much-ballyhooed memory-sharing features lose their effectiveness when VMs get configured with resources they aren't using.

A good IT pro pays careful attention to CPU and memory usage and will tune the number of CPUs and quantity of vRAM to meet demands, based on their assumptions.

Becoming great

Right-sizing should not be a guessing game. A great IT pro strikes the right balance between over- and under-provisioning resources. Luckily the right virtualization operations management technology can provide sophisticated algorithms that use a combination of best practice resource allocation advice and actual workload demand to right-size virtual machine configurations. Capacity management tools for virtualization help you discover which virtual machines are experiencing performance issues and need to be moved, which virtual machines need more resources (undersized), and which virtual machines have too many resources (oversized) and are wasting resources.

Look for technology that will help you maximize VM density by clearly outlining workload placement and reporting on how many more virtual machines can be placed on a particular host. By right-sizing virtual machines, you can save your company thousands of dollars and save yourself some fire-fighting of performance issues.

Optimization Tip #2: Keep it clean and control VM sprawl.

Common mistake:

Ignoring the impact of waste

Waste in virtual data centers can be particularly insidious because it's challenging to even locate. Zombie VMs that are no longer 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 overcrowded with VMs of dubious value and VM sprawl will impact both performance and operating expenditures.

While virtual administrators may hold all the privileges, they don't possess the right to indiscriminately delete. A good IT pro will spend some time implementing policy controls, as well as identifying wasted resources and orchestrating their elimination.

Becoming great

A great IT pro will help realize the cost savings of virtualization technology by preventing VM sprawl using VM lifecycle management technology. Putting automated policy controls and governance in place delivers visibility into the lifecycle of all VMs. Automation and control of the entire virtual machine lifecycle also eliminates manual and repetitive tasks that often introduce

errors, and streamlines the process to save admins time. Plus, it enables businesses to be in strict compliance with IT policies and standards.

The right solution keeps track of resources with the goal of facilitating better planning, while at the same time flagging resources that are ready for retirement. The right technology will alert you to retirement requirements across your virtual environment, helping you reclaim waste from:

- **Abandoned VM images**—An abandoned VM is one that has been removed from inventory but that still remains in the virtual environment. Abandoned VMs are no longer directly visible, but they still consume resources. A good toolset will help you locate abandoned VMs so you can de-provision them and return their hardware resources back to the pool.
- **Powered-off VMs**—VMs that have been powered off for extended periods but remain in inventory are also excellent candidates for retirement. Technologies that monitor for extended periods are necessary to identify VMs no longer being used.
- **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 go without modification for extended periods. An effective toolset will monitor for changes to snapshot files and alert you when a file has not been changed in a long time.
- **Zombie VMs**—A private cloud's self-service functionality offers incredible automation potential for non-virtual administrators. But this automation can come with a price: resource use by VMs no longer being used. These "zombie" VMs can be identified by analyzing their resource consumption over time. The right solution enables IT to harvest the resources of these unused VMs once they've been deemed abandoned.

Not every VM needs to be powered on all the time, because not every VM is providing services all the time.

The right storage configuration can make or break infrastructure performance and availability.

Optimization tip #3: Ensure optimal storage performance.

Common mistake:

Squeezing storage connections

Physical server networking is usually in copious 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. It remains a complex task even today to glean relevant information out of vSphere storage performance statistics. A good IT pro will unsqueeze storage connections by adding networking, adding disk spindles, rebalancing VM processing, or any combination thereof.

Becoming great

Great IT pros know exactly how storage is impacting the virtual environment by ensuring they have complete visibility into critical storage subsystems. They have either taken the time to become storage IOPS masters, or—more likely—they rely on technology to get visibility from 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 can help identify storage bottlenecks and that also delivers one-click troubleshooting to quickly identify the source of each issue and remediate the problem. 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.

Optimization tip #4: Take advantage of vSphere's DRS load balancing.

Common mistake:

Overdoing reservations, limits, and shares

Planning is critical for any vSphere configuration activity, but over-using certain configurations can actually create problems. One area where overuse creates unexpected results is with resource reservations, limits and shares.

Resource reservations, limits and shares can be applied to individual VMs as well as the resource pools that contain those VMs. Well-meaning virtual administrators sometimes use reservations and limits to constrain or guarantee a VM's or pool's resource use. But applying them in multiple locations at once can significantly complicate resource calculations and the effectiveness of vSphere DRS load balancing.

Resource pool constraints can create significant performance problems when shares aren't configured correctly. Resources in constrained pools are divided at the resource pool first, and constraints are enabled only once contention occurs. These rules mean that the innocuous "test resource pool" (with four VMs and 1000 of the available 3000 shares) won't be a problem until the "production resource pool" (with 50 VMs and the remaining 2000 shares) experiences contention. When this happens, the four test VMs will share one-third of your vSphere cluster's resources, while the other 50 production VMs must share the other two-thirds.

In most cases, effective performance management requires a far lighter touch than vSphere's hard constraints. It also requires a long-term view to understand the relationship between today's activities and those in the past. Good IT pros avoid reservations, limits and shares until they understand completely—and can continuously monitor—their effect on everything else.

Becoming great

A great IT pro is an expert in high availability best practices and focuses on cost-effectively increasing the baseline level of availability provided for all applications. Understanding vSphere vMotion capabilities, as well as High Availability (HA) and the Distributed Resource Scheduler (DRS) helps eliminate downtime caused by hardware failures. Getting the right performance monitoring solution in place can provide a significant shortcut to success with these complex virtualization technologies. The right tool provides visibility across the IT stack and simplifies VM provisioning, resource allocation and load balancing.

Optimization tip #5: Maximize energy savings.

Common mistake:

Leaving VMs powered on

The prototypical data center is crowded with operating machinery. Traditional wisdom argues that servers must be powered on if they're to provide services. That needn't necessarily be the case with today's smartly planned virtual data centers. Not every VM needs to be powered on all the time, because not every VM is providing services all the time.

Virtual automation now allows additional servers to be spun up on demand to address client demands. Developer and test environments need constant rebuilding to verify code and complete testing activities. But desktop VMs needn't be operational when their users aren't awake.

Powering down unneeded VMs isn't a task taken lightly. Determining which VMs can be powered down, and when, is a primary challenge. Automatically powering them up at just the right time is another. You need to be able to justifying that effort by quantifying its benefit to the organization.

Becoming great

With power costs near the top of a data center's budget, figuring out ways to

reduce unnecessary power is good for the business. Moreover, it offers the opportunity to squeeze greater utility out of expensive hardware. The great virtual administrator thinks like a cloud administrator, and leverages intelligent tools that facilitate cloud-like automation.

The right operations management solution can help you confidently deploy power minimization functionality. With insight into the minimum number of host servers needed over time to safely run workloads, and estimates on the potential cost savings by powering down unneeded servers, you can easily go green.

Optimization tip #6: Proactively manage capacity.

Common mistake: Building vSphere clusters without reserves and forgetting about cluster growth

vSphere clusters are the foundation of VM performance and high availability. Clusters facilitate vMotion, which, in cooperation with VMware high availability (HA) and Distributed Resource Scheduler (DRS), ensures that resources get used efficiently. Clusters create problems, however, when they're not planned and implemented correctly.

Protecting every VM requires setting aside one host's worth of resources in reserve. These resources are unused. They lie in wait for a cluster host failure, standing ready to process the VMs that existed on the cluster's lost node. Some clusters get built without the necessary reserve. Many more lose it in the face of unexpected VM growth and unexpectedly non-existent hardware budgets.

Disabling a vSphere cluster's Admission Control might earn back resources for a last-minute request. Doing so can also exacerbate downtime when hosts fail. Instead, consider planning your Admission Control Policy to prioritize high-value workloads. When not every VM needs vSphere HA protection, use a percentage policy to balance spare capacity with production needs.

Resource management in shared environments involves careful calculus and a watchful eye.

Today's virtual environments are comprised of so many moving parts that they're becoming impossible to predict with an unaided eye.

The percentage policy makes sense when some VMs can experience downtime in emergency situations. Yet vSphere's percentage policy requires extra planning and a regular checkup. That planning must happen because your percentage of cluster resources to be reserved as failover capacity 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.

A good IT pro will take reserves and cluster growth into consideration and make changes as performance problems occur.

Becoming great

A great IT pro knows that resource management in shared environments involves careful calculus and a watchful eye. You can either get experienced with those equations yourself, or automate the process with intelligent monitoring that delivers actionable alerts and proactive capacity management tools.

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 remaining virtual machine capacity and monitor capacity thresholds with cluster-based capacity utilization. Most importantly, ensure 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.

Optimization tip #7: Spend time modeling capacity scenarios and planning for the future.

Common mistake: Failing to look ahead

This guide's final mistake is the easiest to miss, and the hardest to avoid without help. Forecasting in IT has long involved the "gut feeling" approach. In the past, when environments were simple, good IT pros could feel out future resource needs.

But virtualization and particularly cloud computing has dramatically increased environment complexity. In fact, today's virtual environments are comprised of so many moving parts that they're becoming impossible to predict with an unaided eye. Therefore, gut feelings have begun to lose their effectiveness as a forecasting strategy.

Becoming great

You're human. Ask for help. The right assistance involves a combination of monitoring and mathematics, keeping a watchful eye on virtual behaviors and reporting those that are amiss. Such data feeds directly into planning and forecasting activities to ensure your virtual data center is always prepared to meet business demands.

Look for an operations management solution that provides capacity trending, forecasting and alerting that will project time and resource consumption limits based on historical growth rates. You also need a tool that can help you analyze "what if" scenarios by modeling both virtual and physical workload placement on available hosts for improving efficiency in virtual infrastructure or moving workloads from physical to virtual infrastructure.

Conclusion

Managing a virtual data center today is undeniably a complex task. You need to not only optimize resource usage today, but keep an eye on growth so you can plan effectively for tomorrow. Fortunately, it doesn't have to take years to learn. With the seven key tips in this paper, along with the right tools, you can avoid common mistakes and fast-track your way from being a good data center administrator to becoming a great one.

For More Information

© 2013 Dell, Inc. ALL RIGHTS RESERVED. This document contains proprietary information protected by copyright. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose without the written permission of Dell, Inc. ("Dell").

Dell, Dell Software, the Dell Software logo and products—as identified in this document—are registered trademarks of Dell, Inc. in the U.S.A. and/or other countries. All other trademarks and registered trademarks are property of their respective owners.

The information in this document is provided in connection with Dell 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 Dell products. EXCEPT AS SET FORTH IN DELL'S TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT,

DELL 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 DELL 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 DELL HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Dell 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. Dell does not make any commitment to update the information contained in this document.

About Dell Software

Dell Software helps customers unlock greater potential through the power of technology—delivering scalable, affordable and simple-to-use solutions that simplify IT and mitigate risk. The Dell Software portfolio addresses five key areas of customer needs: data center and cloud management, information management, mobile workforce management, security and data protection. This software, when combined with Dell hardware and services, drives unmatched efficiency and productivity to accelerate business results. www.dellsoftware.com.

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

Dell Software

5 Polaris Way
Aliso Viejo, CA 92656
www.dellsoftware.com

Refer to our Web site for regional and international office information.

Share:

