



Storage Switzerland

Is It Time to Upgrade Your Storage Evaluation Process?

Author: George Crump, Lead Analyst

August 2013

Introduction

In the modern data center, storage system upgrades are rarely caused by a storage system running out of capacity. Rather, they are caused by an unanticipated lack of performance or exorbitant maintenance prices. In fact, performance related upgrades can occur at any time, often well before the current storage system was due to be refreshed and more importantly before the evaluation process had predicted a refresh would be needed.

An unexpected growth in applications, a significant new product launch, an acquisition, or rapid increases in server virtualization are just four examples of how storage platforms can be stressed beyond their initial performance specifications. In addition, application workloads are constantly changing or “drifting”. The primary culprits causing these pre-mature refreshes are the inability to properly understand storage infrastructure performance requirements on an ongoing basis and the inability to predict the performance limits of the newly acquired storage system. This can be easily resolved by simulating production workloads prior to implementing a new storage solution and by maintaining a consistent testing and validation process for constantly changing workloads. The key is to accurately model the current demand and the potential growth of the production workloads.

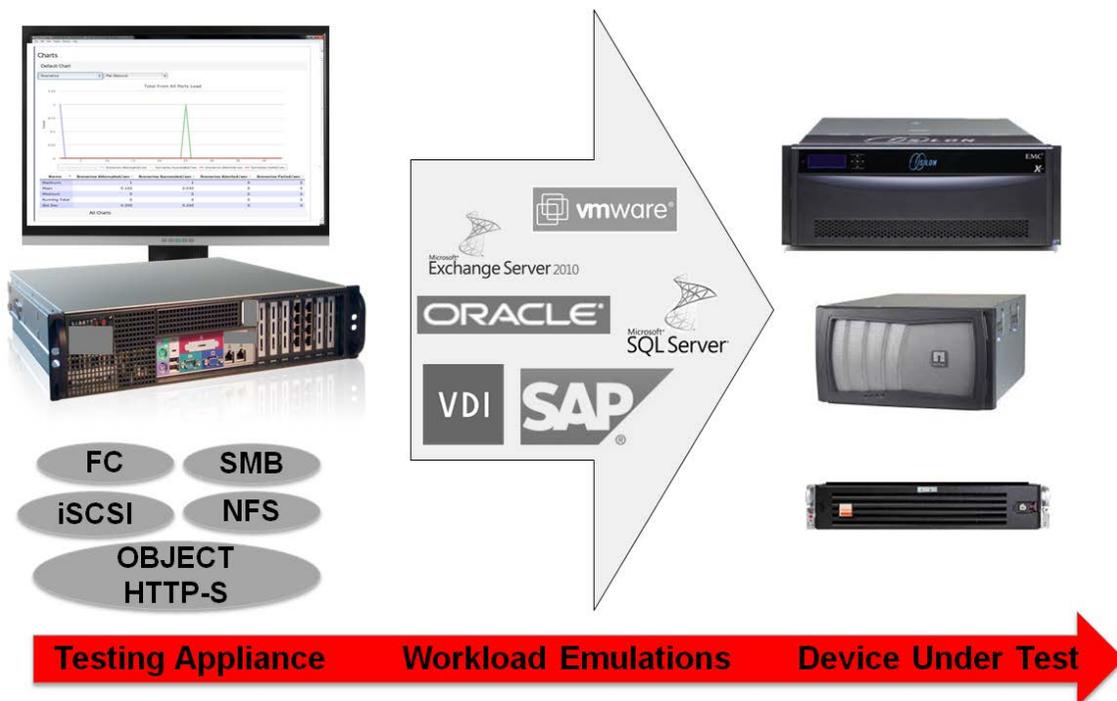


Figure 1. Illustration showing how to test real-world production workloads before they are deployed into production.

The State of the Storage Acquisition Process

The end user storage evaluation process has become much more challenging over the past 5 years. The first problem is that there are not enough people to adequately perform accurate testing and evaluation of new products and new technologies. IT staffing levels have at best been flat over this time. The "Do More With Less" data center environment is essentially a shell game where time is stolen from other less visible tasks and the evaluation process typically ends up on the short end of that reality. With technologies like solid state, tiering, deduplication, caching and new protocol versions all having the potential to dramatically affect application performance and the cost of your infrastructure; informed decision-making is essential.

The second problem with conducting accurate storage system evaluations is the nature of production systems themselves. The organization is so dependent on these systems that they can't ever be taken off-line. Likewise, gradual cut overs from System A to System B is a luxury of the past. Virtualization has helped evaluators spoof "worker" systems; however, most testing tools can't simulate real users accessing real data at a meaningful scale.

While virtualization helps to stand up a test lab, creating an accurate simulation of production storage I/O can be difficult. In fact, the very flexibility that makes virtualization so valuable, ends up being an impediment to recreating real world production workloads in the evaluation lab. At best, only a small subset of the application workload can be tested, which generally results in IT managers having to err on the side of caution and over-buy product.

Most importantly, the ability to model and validate what performance will be like under future load requirements requires special purpose devices. Without such a solution, the storage engineer or architect can't provide an accurate picture of how much performance growth ("headroom") the potential new storage system will be able to support. The impact of this is that when the new system hits its own performance limits, it may be unanticipated and could result in a mad frenzy to find a replacement system.

Another substantial obstacle to adequate testing today is the lack of mature tools that can model realistic workloads and the substantial hardware that must be acquired to generate the load. Historically, testing at production scale requires dozens of servers and substantial IT staff resources to create scripts and run tests using inadequate open source software. One example is the inability to model metadata operations, which is essential for file-based storage systems.

The Goals of Storage System Evaluation Testing

Storage evaluation testing should have a relatively simple set of goals. The primary goals are to reduce the risk of business interruption (outages) and to ensure the right product and amount of product is acquired. More specifically, the evaluation should try to model the production workload environment as accurately as possible. It should also enable storage engineers and

architects to pre-determine the performance limits of the storage system (“knees of the curve”) so that the next storage upgrade can be a carefully planned event instead of a disruptive fire drill. Lastly, it should allow IT planners to provide reliable information to application owners about the potential impact that new projects will have on current storage resources so that they can budget ahead of time for future systems.

The Cost of Getting "It" Wrong

When a product's current or future capabilities are misdiagnosed, there is a direct financial impact to the organization and the IT department that approved it. With roughly 40% of the typical IT hardware budget spent on storage, small miscalculations on storage requirements can have a multi-million dollar impact on the company's budget.

First, there is the obvious cost of procuring a replacement system prior to the current one being fully amortized. Second there is the cost of overbuying. A small 10% increase in savings due to buying only what is needed can lead to a significant ROI. Third, already undermanned IT staff has to dedicate time to perform due diligence and evaluate new storage system(s) in addition to their usual duties of managing production systems. Consequently, not only might the evaluation suffer from a lack of focus due to real-world time constraints, but so too may the critical business applications that these IT professionals manage.

The Storage Evaluation Challenge

Due to all these systemic business risks, proper evaluation testing that can validate performance is not a nice to have, it is a must have. The cost of getting it wrong carries too many negative consequences to the business. The problem is that for most organizations, storage evaluation testing is something that is spun up occasionally as new projects arise or as performance limits of the current system are breached.

Dedicating personnel for this task is hard to justify, especially considering the range of expertise required. The reality is that understanding infrastructure performance is an ongoing activity. New applications and users are constantly being added, firmware or operating system updates need to be implemented all too often and new storage technologies are always in need of evaluation. Knowledge is needed not only in storage infrastructure design but also in how the applications, updates and upgrades will utilize and affect the storage resources on a continual basis.

The Answer is In the Machines

Evaluating and assessing storage infrastructure needs to be a consistent ongoing process. The ideal solution is to acquire a workload modeling and simulation product that is purpose built for this task – such as the Load DynamiX storage performance validation appliances. Systems like

these are beneficial because they allow users active applications that are currently connected to networked storage systems to be fully profiled. This profile can be then turned into a workload model that represents real-world I/O traffic that can be automatically applied to any potential new storage system.

In addition, the profile can be scaled to simulate the expansion of the environment and its impact on the storage system. In other words, the limits of the system can be identified and planned for, eliminating the scramble otherwise caused by an unexpected performance wall. Most importantly, use of such workload modeling solutions ensures that no IT manager will over invest or under invest in their storage systems.

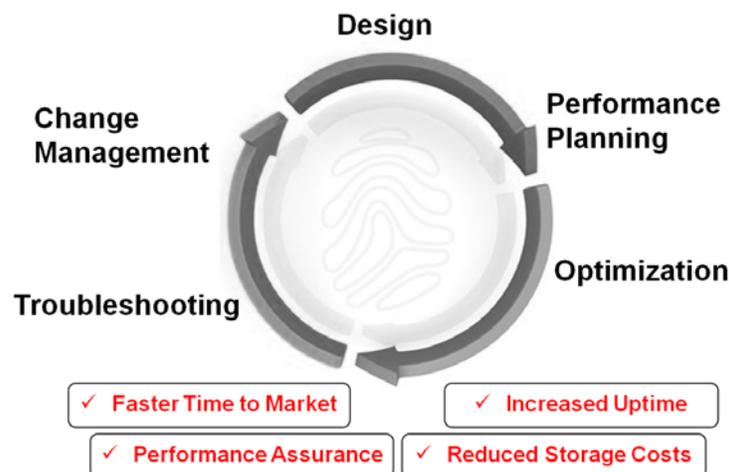


Figure 2. Illustration showing when to use a storage performance validation process and its benefits.

Conclusion

The difficulty in modeling the production workload and in scaling that workload to worst case scenarios will continue to be a challenge for the over-worked storage engineer. The chance for error is high, as is the business impact resulting from those errors. A more viable solution is to leverage technology, like Load DynamiX's, that can simulate production workloads without the expense of dedicated lab infrastructure or siphoning away limited cycles from already time constrained IT personnel.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage market places. For more information, visit: www.storage-switzerland.com.



Storage Switzerland, LLC