

NetApp: Lessons Learned In A Virtual Datacenter

Date: February, 2008

Author: Steve Duplessie, Senior Analyst

Abstract: NetApp is a major IT vendor—and a major IT user. As an internal service bureau, NetApp IT has found ways to creatively leverage next-generation datacenter design techniques that benefit its mission—proving that intelligent IT is a very real possibility not just for NetApp, but for everyone.

The Problem

Like any vendor of IT technology products, the ability to test new software and hardware, simulate user environments in a support crisis and maintain an endless array of compatibility tests is a brutal challenge. Largely, IT is an area with either loosely interpreted standards or none at all, which means that in order to become a tier one vendor trusted by the biggest shops in the world, huge investments have to be made. Three inch thick compatibility matrices are often seen as a badge of honor; distinguishing the few who can afford to maintain such an investment from those that can't. The fact that it is necessary in this day and age is a different issue—like it or not, it's reality. Technology developers have to spend huge amounts of money on infrastructure, tools and people in order to ensure that their products work in as many customer environments as possible.

Any IT shop knows the issues that exist attempting to become a legitimate service bureau to their respective business units. These departments are plagued with the problems of legacy applications and infrastructure, bureaucracies, regulations and a never ending tide of data growth. All of those issues inevitably contribute to the chasm that exists between IT and the business—the business needs things faster and IT isn't in a position to deliver.

Whether it's the IT department of a bank or a technology vendor, the issues are the same. The cost and complexity of infrastructure has been a huge and worsening problem that impacts IT's ability to quickly deliver services in the face of ever changing requirements. Power, cooling, footprint, people, time and money are universal issues—regardless of the type of business.

NetApp's Test and Development Quandary

Most people know NetApp as one of the rare storage players to hit a long-term home run. The company effectively created and then dominated in the era of distributed computing (centralized) file serving and has run almost unopposed in that space for about 20 years. Along the way, it has added dozens, if not hundreds, of features and models of products—each needing rapid and thorough regression, compatibility, performance and operational testing.

As company growth outpaced 99% of the industry, it is no surprise that internal IT, Quality Assurance, Support, Test and Development, and other support systems were tested time and time again. "It just works" has always been the top comment from NetApp customers and a top reason for buying more and more because NetApp has never had the luxury of putting products that were anything but perfect into the market. Like other IT shops, that meant spending and reacting. "Just buy more" is not a phrase relegated only to other IT shops—it was necessary at NetApp as well (though NetApp did have the advantage of providing its own storage, that didn't help with servers, networking, people and operations outside of its control).

By 2005, the company looked a lot like other huge IT shops—too much gear, too hard to manage and too much risk. The test and development organization alone had over 1000 servers, hundreds of switches, PBs of storage and was consuming ever growing amounts of data center space, power, cooling and man hours. Early attempts to consolidate were thwarted by the almost endless permutations of configuration support required. It was becoming almost impossible for Test/Dev to provide the type of quick turnaround times that the business units

required, so instead of “just add more,” the group was given the (rare) luxury of taking a breath and looking at the problem as it existed—and then making assumptions as to what that would mean as growth of gear, configuration permutations and footprint would mean in the face of shrinking delivery windows. The group came to the obvious conclusion that staying the course would lead to eventual collapse.

The Fix – Kilo Client

While the internal code name, “Kilo Client,” is bad by any measure (named due to the 1000 nodes each some form of client configuration), what NetApp actually did to solve its test/dev issues could be a model of datacenter efficiency, service delivery ability, flexibility and operational brilliance. It not only solved internal issues, but became a proven methodology that can work in any IT operation. The best part is that it did not require wholesale changes and has an effectively unlimited scale capability.

What the Kilo Client Lab Looks Like Today

- 1500 servers, now 90% blade servers with no per server physical system drives
 - 1GbE (multiple) to each blade
 - Fibre Channel (2, 4, 8 Gb) on several hundred blades
 - 10GbE backbone, smaller numbers of 10GbE connected blades
 - VMware
 - All types of NetApp systems as storage (file, iSCSI, FC running on each)
1. VMware was recently introduced across the server farm, enabling far simpler provisioning of a request for a server profile—i.e., a set of virtual resources capable of emulating and generating specific workloads in order to complete specific tests. By leveraging the liquidity of VMware, server resources can be set up and broken down in a few mouse clicks. Tens of thousands of virtual clients can be created or dismantled in minutes.
 2. By consolidating all the required images that could be needed in the test process onto its own arrays, NetApp has been able to boot a server off of any virtual disk it required—without server virtualization. Normally, a shop with 1,500 blade servers is also going to have 1,500-3,000 disks just for the boot and application images. Of those 3,000 images, only a microscopic percentage of the data is unique to a specific node or image. Netapp created single “gold” standard images—for each variant of Windows, Linux, UNIX, etc. that would be required—along with the same for each application type and each workload generator. The company then leveraged its own writeable snapshot technology to apply any individual differences or unique elements so that every node and every physical or virtual machine continued to boot and use its own unique image—without the additional disk space or physical disk requirements. Thin provisioning technology within the NetApp architecture has always meant that only the actual disk spaced used is truly allocated, so utilization is always optimized. FlexClone technology lets NetApp people create new images in seconds. This enabled NetApp to consolidate 3,000 boot disks down to 252—a 90% reduction—and eliminate the need for even more disks to house all the application and workload variations. By enabling remote boot, NetApp’s ability to create huge physical infrastructure stacks (presented as any number of virtual client configurations) went from days to minutes. When they are done, the entire stack is released back to the pool of resources.
 3. The entire operation is managed by 5 people. 1500 physical machines, countless configurations—all created, utilized and torn down by less people than typically manage tape cartridges in a big shop. That’s not just impressive; it’s amazing.
 4. As NetApp added or replaced equipment, new gear was migrated into far more efficient “Hot/Cold” isle configurations. Equipment was racked into the tallest cabinets available and stacked as densely as possible. While that certainly generates tons of heat, Netapp was able to design data center enclosed rows so that all the heat was exhausted in the same direction and sucked out through air handlers in the hot rows. Cool rows use air conditioning through the floor, pulling cool air through the front of each cabinet/rack. Hot air is exhausted out the rear of the cabinets into the contained hot rows, where it is sucked up and out. This system allowed NetApp to optimize density by jamming each rack as full as it could be, while consuming only 50% of the power and cooling that was previously required.

The Bottom Line

What started as an isolated Test and Development issue became a microcosm of the problems plagued by IT globally. Inside NetApp, the Kilo Client experiment has been so successful that the company has adopted its

principals across all data centers. Customers are using the Kilo Client lab (Research Triangle Park, NC) to perform proof-of-concept testing—bringing their own gold images in their backpacks to emulate their exact configurations and exact workloads, which eliminates any variables and ensures that testing will demonstrate how a change is going to affect their current operations and reduce huge risk factors. More customers are touring the facility to better understand just how they can add to the benefits that NetApp has been able to deliver to their businesses.

What NetApp has done is not revolutionary, nor entirely unique. Plenty of people use VMware and Netapp gear, and have figured out that there has to be a better way to control facility sprawl and costs. What they have done, however, is to prove that it is feasible and possible without spending huge amounts of money, hiring a number of skilled people and adding huge risks to the operation. “Cheaper, faster, simpler” has long been the NetApp product mantra, and clearly still is a corporate one. Its IT, like all IT folks, simply wanted to be able to support the business, react quickly to what was asked of them and stop being considered a tactical problem. NetApp shows us how the application of current technologies like server virtualization and storage virtualization can be married together to provide the ultimate IT goal; the ability to simply say “YES” when the business needs something. Kilo Client started as a way to solve a tactical set of problems and as such, was considered a cost center (like 99% of IT)—a.k.a., a necessary evil to the business. After being given the time and opportunity to find a better way, NetApp did. Along the way, the company made the quiet leap from being tactical to become entirely strategic—and if IT can do that, then all the problems of the past twenty years can go away. When used as a strategic competitive weapon, as it was originally intended, IT can rightly regain its seat at boardroom.