executive summary

virtualization, simplification, and storage

In 2001, HP announced two new disk arrays—the HP Surestore Virtual Array 7400 and the HP Surestore Virtual Array 7100. This paper will focus entirely on the virtualization technology that makes these the easiest arrays to manage and the most intelligently simple arrays on the market. This paper answers the question, "What is the value of HP's array virtualization?"

For those familiar with data center environments, virtualization is not a new concept. Virtualization already exists to some extent at every point in the solution stack. For example, servers, operating systems, databases, file systems, volume managers, drivers, switches, and storage devices all require virtualization to achieve their purposes. **At the array level, HP's virtual architecture simply expands on already familiar ground, and it cleanly fulfills the promise of virtualization—it hides complexity from the administrator and can have a dramatic positive effect on real-world performance.** This reduction in complexity greatly simplifies and streamlines the data center environment.

HP's array virtualization:

• reduces the time spent managing individual arrays

• allows storage administrators to manage more storage with less effort

• reduces the opportunities for human error

• frees up precious IT resources to work on revenue-generating projects

• self-manages the RAID configuration for optimum performance

These are not trivial benefits. Data centers are already complex, and with the continued explosion in storage capacity they will only become more so. HP's array virtualization is evolutionary, not revolutionary. It is a logical progression in array technology. It is proven. It is necessary. It is the perfect storage area network (SAN) technology. And it is offered **only** by **Hewlett-Packard!**

array virtualization defined

The purpose of virtualization in any technology is to hide complexity from the user, or in the case of disk arrays, to hide complexity from the storage administrator and provide a standard environment for application development and increased price/performance.

Disk arrays are complex devices designed for complex tasks. A disk array with 50 disk drives is more complex to manage than a disk array with one drive. Armed with virtualization, an array could potentially allow the 50 drives to be perceived and managed as one big drive or as one big pool of storage. The power of virtualization is the power of simplification.

Now for the definition: **Virtualization in arrays is about creating and managing virtual storage devices.** It is about taking blocks of storage on the disk drives and presenting them as LUNs (logical units of storage). What system administrators see, then, are not the actual physical disk drives but rather a created, simplified "virtual" view of the actual physical storage, i.e., the LUNs.

HP's virtual array works with LUNs just as traditional arrays do. However, the virtual array doesn't stop there. The virtual array actually manages the disks down to the level of the smallest available "cluster" or "chunk." Further, while traditional arrays utilize static address-translation algorithms for managing the chunks, the virtual array uses a dynamic mapping system. This allows the array to dynamically allocate and de-allocate clusters of any RAID type without affecting the logical view as seen from the server. In other words, clusters can be moved and new RAID 5 and RAID 1+0 stripes can be created or extended dynamically without the host view being changed and without system administrator intervention.

Just as a file system on the host presents a virtualized view of the storage to the application or user, the virtual array presents a virtualized view of the storage to the file system. For example, just as you can create and delete files in a file system, you can create and delete LUNs within the virtual array.

This degree of virtualization is what gives the virtual array its unique ability to relieve the system administrator of many of his/her mundane storage management responsibilities.

the reasons you <u>need</u> more virtualization in the array

- Capacity requirements will always grow.
- IT departments will always have a limited budget and limited human resources for managing storage.
- Virtualization improves the efficiency of storage administrators.
- Virtualization reduces the training required for people to manage storage.
- Data always grows, but human headcount doesn't. Virtualization reduces the number of people required to manage storage.

key features of hp's patented virtual technology

- Automatically stripes every LUN across all disks in a very large redundancy group.
  Benefit: simplifies overall management; reduces the number of LUNs required to achieve a balanced workload and eliminates "hot spots."
- Automatically adds new disk drives into existing RAID groups.
  Benefit: capacity upgrades require no IT training, are faster, and hot spots are virtually eliminated.
- Dynamically, and without any human intervention, optimizes the RAID level to the application workload.
  Benefit: performance tunes the array 24 hours per day and eliminates downtime for reconfigurations.
- Allows LUNs to be created in seconds.
  Benefit: reduces setup times.

- Allows immediate use of the array after LUN configuration.
  Benefit: speeds implementation time.

- Allows hundreds of LUNs to be created of virtually any size—from a single megabyte to over 2TB.
  Benefit: simplifies configurations, improves application performance, and greatly decreases management complexity and potential for security errors.

- Allows any combination of disk drive capacities and speeds within the array.
  Benefit: simplifies disk drive inventory control and capacity management.

- Allows disk drives and controllers to be moved to any slot in the array.
  Benefit: reduces the chance for human error.

why now is the time for array virtualization

Data centers have become enormously complex. Interestingly, new technologies that initially held the promise of simplifying the data center have, in fact, added to the complexity—at least for now. Fibre Channel, SANs, and the low cost and ease of adding storage have enabled data centers to expand both in size and complexity.

The world is moving from SCSI to Fibre Channel. Fibre Channel and SANs allow for more devices to be connected together, with greater performance, and at far greater distances. All of these are wonderful benefits, but along with those benefits comes an ability to now create configurations that are larger, more complex, further apart, and harder to manage than anything envisioned even a few years ago. Couple this new reality with the exponential increase in storage, fueled by its low cost, and you have a recipe for losing control of the data center.

Consider this: while human resources remain static, environments grow larger and more complex. Simplification is the only realistic answer. Adding newer and slicker management software oftentimes can help, but in the long term, adding software tends to have the opposite effect of reducing complexity. Software may give you a central place from which to manage your hardware and it may simplify processes, but it doesn't necessarily eliminate those processes. Armed with better management software, environments are encouraged to add more hardware and eventually the environment becomes as complex as it was before, but for different reasons. And when the human resources are already stretched to the breaking point, this is a recipe for complexity, stress, long hours, and human error. Virtualization is the answer. It solves the basic underlying problem. It permanently simplifies the environment for the system administrator.

managing traditional storage

"The mistakes are all there waiting to be made."
—chessmaster Savielly Grigorievitch Tartakower
  on the game's opening position

Configuring and managing traditional storage is time-consuming and challenging even for experienced system administrators. When configuring storage, a wide variety of factors must be accounted for, judgments must be made, and steps repeated over and over again. The potential for error is high. The cost of configuring an array improperly is also high. The administrator must consider the following factors:

- capacity, cost, performance, and availability requirements

- requirements for future additional capacity and performance

- number of disk drives and their capacity

- performance characteristics of the disks

- which RAID level will meet desired capacity, cost, and availability needs

- number and type of RAID groups

- number of LUNs based on application, performance, and array configuration needs

- size of LUNs

- configuration of the server volume manager

- cache configuration options

- stripe depth configuration

- implementation plan: who, what, and when

- time to bind LUNs

In addition, the same processes must be followed whenever disk drives are added or the environment changes. Plus, these additional factors must be considered:

- current configuration

- desired additional capacity, performance, and availability

- whether the new disks will be stripe extensions of existing disks or be independent groups

Here is a typical process for setting up a traditional array:

1. determine number of disks, number of RAID groups, disks and disk type per RAID group, RAID level of each group, total LUNs, LUNs per RAID group, stripe depth

2. determine volume manager configuration, stripe size and depth, LUNs per logical volume

3. using the command station, set up the LUNs and their RAID levels and assigns them to particular disks

4. set up the cache page size depending on the size of the I/Os coming in from the host

5. finally, before the new LUNs can be used, disks must be formatted, which can take many hours per array

configuring an array for a database

Properly configuring an array for a database typically involves a large problem set with many variables. Many database administrators have been taught to isolate different pieces of the database in an attempt to optimize performance, availability, and recovery. This process, although based on sound objectives, is far too error-prone. This typically involves a large problem set with many, many variables. Unfortunately, database administrators oftentimes don't have all the critical information. They don't know the precise database performance requirements for each of the pieces, and they don't know the performance behavior of the array in its multitude of configurations.

In these real-world environments, it is typically far too time-consuming to try a number of different storage configurations, so database administrators typically apply rules from previous installations. The changing characteristics of newer versions of the database typically result in an unbalanced configuration that has "hot spots" that limit the performance of the system.

This entire process can take from a few hours to several days, depending on the skill of the administrator and the number and size of the LUNs. During much of this time the array is either unusable or must operate in a degraded performance mode. In other words, LUNs cannot be utilized until they have been formatted. This formatting takes up a lot of the array internal resources and bandwidth. After a LUN has been formatted, it can be used; but as long as other LUNs in the array are also going through their format process, the entire array will suffer from degraded performance.

Now, a short word about human error. Every step of this process has the potential for human error. Except in the grossest cases, errors would probably not result in data loss, but every miscalculation in this process would easily result in a decline in performance. Some of these declines could be huge. For example, miscalculating the RAID levels or the cache page size could severely degrade the array's performance.

the system administrator's dilemma

Change is the issue. Many environments change over time and this makes their initial configurations progressively less and less optimal. The dilemma is that an optimum configuration today eventually becomes out-of-date, and typically over time, performance degrades in traditional arrays, but the administrator usually has neither the time nor the window of opportunity to bring the system down and perform the necessary reconfiguration.

Access patterns change. Different areas of the database become more highly used. Sometimes at the end of the month certain data that normally lies unused now becomes highly used and requires higher performance. All of this poses problems for the conscientious system administrator.

At first, the administrator notices a change in performance, but it is not enough to justify bringing the system down. However, over time, the performance continues to degrade until eventually the problem is so severe that the administrator is forced to bring the system down over the protests of the users. He must then go through all 13 configuration steps, including binding all the LUNs and waiting for the reformat to complete. After that the cycle begins again: Performance initially is great, but over time it degrades until finally the pain increases to the point that a reconfiguration again becomes justifiable.

managing the hp virtual array

Configuring an HP virtual array is much simpler than the process for configuring a traditional array. Remember the complex steps involved in configuring a traditional array? The steps involved in configuring the HP virtual array consist of:

• Determine the total capacity and performance requirements.

• For each application, determine the number and size of LUNs.

• Determine the number of disks required for the necessary performance.

• Create the LUNs. Note: creating LUNs takes only a few seconds per LUN.

Step away from the array; the configuration is now complete. Every other step is automatic. RAID levels are automatic. The different capacity, number, and speeds of disk drives are automatically accounted for. The cache page size is automatically set. Even the disk formatting is performed automatically. Moreover, after the LUNs are set, the array is immediately available to accept data. The array does the work, not the administrator.

Not only is this the initial configuration process, but the process is just as simple for any subsequent reconfigurations. In other words, if LUNs have to be deleted and new ones created, the process is just as simple. Note: as with any array, if you wanted to delete the LUNs but save the data, you would have to do a backup and restore.

adding capacity with hp virtual arrays

HP's Virtual Array Architecture also simplifies the process of adding capacity to an array. Today, many traditional arrays allow the administrator to add disks on-the-fly—in other words, to add a disk drive when the array is up and running and accepting I/Os.

However, when the disk drive is added to a traditional array, it is not part of any LUN. It is not formatted, and it is not able to accept data. An administrator must go in and manually perform those functions. If the disk drive is to be added to an already existing RAID group, then the data in that group must first be backed-up and later restored to the newly created LUN that now contains the new disk drive. And except for the backup, the steps are the same as when an array is first configured.

The HP virtual array accepts new disks while the array is up and running and accepting I/Os as with some higher-end traditional arrays. However, the HP virtual array takes it one step further. Once the disk is inserted, the array automatically includes that disk into the existing disk space and stripes all LUNs across that disk. This means that even without the creation of any additional LUNs, the array performance will improve because of the additional available spindle. Only the HP virtual array automatically adds the new disks to existing LUNs. Further, any newly created LUNs are also automatically spread across all the disks in the array, including the additional disk.

time to implementation: formatting the array

As mentioned earlier, after new disks are added to a traditional array, it then takes several hours to complete the formatting of the RAID group. During this format phase, no data can be written to the new LUNs. With some implementations, the array is offline until all the LUNs have been formatted. In other implementations, I/Os can be written to already formatted LUNs even while other LUNs are going through the format process, although performance is very slow.

Because executing the disk format command uses up so much of the array's internal bandwidth, array performance is greatly reduced until all of the disk formatting has been completed.

With HP's virtual array technology, the array is immediately available as soon as the LUNs have been configured. The disk formatting is done as the writes are done. In other words, as writes are sent to disk, the formatting is accomplished for only those blocks being written to. This means that while there is a small hit to performance for that individual write, there is very, very little impact on overall array performance.

automating the cache parameters

Configuring a traditional array typically requires setting the cache parameters such as the percentage of read and write cache, the size of the cache pages, and, in some cases, the allocation of cache to specific LUNs. In making these determinations, there is ample opportunity for error.

With HP's virtual arrays, all of this is preset and automatic. And this means that all the parameters within the array are tuned to work in unison with the stripe size and the array hardware. First, the cache is set at 80% read and 20% write, is shared between controllers, and is treated as a "pool." Second, the cache page size is set at 64K and is set to automatically destage to disk every 4 seconds whether the page is full or not. The 64K size minimizes the number of I/Os to the back-end in sequential environments and provides a carefully calculated balance within the array between the number of cache pages and the speed of the back-end in random environments.

performance

Traditional arrays are susceptible to "hot spots" and to changes in the environment that make the initial configuration obsolete. The HP virtual array virtually eliminates these critical performance issues.

First, the HP virtual array is far less likely to experience a hot spot—in other words, it will almost never experience a condition where a few disk drives become a performance bottleneck in the array. Here's why: the virtual array always (and automatically) stripes all of the LUNs across all of the disks in the RAID group. For example: assume a virtual array loaded with a total of 60 disks had 30 disks in each of its two RAID/redundancy groups. Every LUN in that group would be spread across all 30 disks. Every LUN would have 30 spindles at its disposal. And don't forget, in the virtual array the spare disk capacity is also spread across all of the spindles, i.e., there are no unused spindles in an HP virtual array.

Second, the virtual array automatically performance tunes the array 24 hours per day, 365 days per year. The RAID level is matched with the workload. Data with access patterns that would benefit from RAID 1+0 storage are automatically directed to a RAID 1+0 section of the array. Infrequently used data, or data whose access patterns match RAID 5DP performance characteristics, are directed to the RAID 5DP section of the array. These are the same rules that a storage expert would use to optimize an application's performance. In addition, like an expert storage administrator, any changes to the configuration are made only during low-usage periods or when disks are added. Note: when disks are added, the existing LUNs are automatically extended across the new disks. This normally requires data movement within the array. However, administrators have the option to postpone this data movement by simply turning off the array's "Auto-Include" feature.

FAQs

Question: *Isn't this virtualization technology new and doesn't that make it risky and untried?*

Answer: It would be new for HP's competitors, but HP has been virtualizing arrays since 1995, and since then HP has sold over 20,000 virtualized arrays. The technology is neither new nor risky and it offers a simple and compelling value proposition, i.e., great ease of management and great data protection. For arrays, management costs are far more costly than the initial purchase price. The HP virtual array is the easiest array in the world to configure, add capacity, and manage—the administration savings are significant. It also has the best availability of any mid-range array on the market. See the white paper titled "VA7100 Hardware High Availability Features" for more details.

Question: *Shouldn't a good system administrator know which RAID level the data is in? With the virtual array, I have no idea if it is in RAID 5 or RAID 1+0.*

Short answer: Both the HP Surestore Virtual Array 7100 and the HP Surestore Virtual Array 7400 can be configured to run in either fixed RAID 1+0 or AutoRAID modes. And to help with performance analysis, the controller can provide data on the actual usage of each RAID level.

Long answer: With the virtual array, if your data is frequently used, it WILL be in RAID 1+0 and will have the best performance. The virtual array RAID level policies were developed after researching the decision patterns of experienced system and database administrators. Almost always, technology progresses from highly manual to highly automated operations. In almost every case, the developers of automation simply replicate the best of the already developed manual processes. This is exactly what HP did with the HP virtual array. Also please note: the virtual array policies are improved over those of the Model 12H. The virtual array really does strive in almost every instance to do any background data movement during periods of low array activity.

Question: *On which disks is my data kept? I suppose the virtual array can be trusted to handle the RAID level decisions, but I also need to know on which disks my indexes and redo logs are kept. The way the virtual array moves data and stripes across all the disks means I have no idea where my data is kept.*

Answer: First, if you are worried about the integrity of your data, the HP virtual array does a combination of things to protect your data that no other array does. First, the HP virtual array offers end-to-end checksum, ECC protected memory, parity coherence, disk scrubbing, and RAID 5DP which, in a typical configuration, gives 100X the data protection of traditional RAID 5 and 10X the data protection of RAID 1. Second: the HP virtual array stripes across all the disks in the redundancy group. This involves more disks but is not so different from any other RAID 1+0 implementation. Third: array striping is the way of the future. It's faster and safer.

Question: *Doesn't all this "behind the scenes" movement of data require a huge performance hit? Surely, the trade-off for virtual technology is slow performance.*

Answer: Not true. For years, HP's original virtual array, the AutoRAID Model 12H, was used for HP's V-class TPC-C benchmark tests. And today, the HP virtual array products have been used for the new rp8400 benchmark testing and are scheduled to be used for the Superdome TPC-C tests. Remember, the HP virtual array mimics the policies of experienced system administrators. Would an experienced system administrator do a reconfiguration of the array during a period of high workloads? No! And neither would the HP virtual array. The HP virtual array policy is to NOT perform the background tuning operations when the array is under a greater than 60% workload.

Question: *Isn't manual always better than automatic? In cars, manual transmissions give better performance than automatic transmissions because they give the driver more control over performance; likewise wouldn't manual RAID configurations be better than HP's Virtual Array Technology and its automatic RAID configurations?*

Answer: Manual transmissions in cars would not give better performance if drivers were only allowed to shift gears once and could never change them after that. In essence, this is what you have with traditional disk arrays. You are stuck with the initial configuration unless you bring the array down and go through a time-consuming and complex reconfiguration every time the environment changes. HP's Virtual Array Architecture tunes the array automatically, 24 hours per day. It is the hands-down winner in real-world performance. A better analogy would be to compare the multiple manual processes required to set the type, load the paper, and actually print books on the old-style printing presses with those of the automated printing systems of our generation. Obviously, automation in printing presses adds to greater performance. It works the same for arrays.

Question: *Wouldn't striping the data across such large RAID groups make the disk rebuild times very long and take up a high percentage of the array resources?*

Answer: Because of HP's RAID 5DP this is practically a non-issue for the HP virtual array products. The HP virtual array's RAID 5DP gives each redundancy group 10X the protection of RAID 1. RAID 5DP requires that 3 drives would have to fail before there would be data loss. Thus, even AFTER a drive fails, the data is STILL protected with the same degree of protection as standard RAID 5. Finally, this means that a single drive failure does not put the data at risk and therefore does not require an emergency rebuild. The HP virtual array can take the time and do the rebuild in the background without impacting incoming I/Os. Also, the HP virtual array will rebuild the RAID 1+0 data first since that is the most vulnerable after a failure. In all cases, data integrity is ensured and performance is preserved. No other array can make this claim.

Question: *If your environment is totally stable, wouldn't a manual configuration by an experienced system administrator result in better performance than one derived from the virtual array's policies?*

Answer: Let's first admit it: totally stable environments are rare. But the answer is that if the environment were totally stable, and if the administrator configured the array absolutely accurately, the traditional array and the HP virtual array would have similar performance. However, the HP virtual array would still have the performance advantage because of its ability to efficiently load balance across all the disks in the array both at the initial configuration and after capacity growth. In both stable and changing environments, the HP virtual array is the right choice.

Now, if the environment is not totally stable, then the HP virtual array is the hands-down winner. When both arrays are initially configured, the traditional array should be faster at least for a day or so, but after that the HP virtual array will catch up and continue to operate at peak performance for as long as the array is plugged in, while the traditional array will get slower over time. This is how it works. When the HP virtual array is first turned on, it doesn't know which data needs to be in RAID 1+0 and which data needs to be in RAID 5. After a day or so of reading the access patterns, it will figure out which RAID level is best suited for which data. Once that happens, the performance will be as good as in any manually configured array. And, don't forget, if the environment should change after that, even slightly, the HP virtual array will adjust while a manually configured array can do nothing but keep plugging along in a degraded mode.

Question: *Are you saying that this technology is for everyone? I'm not looking for anything new.*

Answer: If it's new to you, then it is understandable that HP's virtual technology might seem unnecessary. After all, you've succeeded without it up till now. But sometimes improvements do come along that really do hold out the promise of a widespread advantage. The dilemma of our rapidly advancing technological age is to know which new technologies should be ignored and which are the ones that need to be grasped.

Consider: It is always easier to manage a smaller number of things than a larger. At first, bits were grouped into bytes. When there were too many bytes to keep track of, they had to be grouped into blocks. When there were too many blocks to manage, someone invented volumes and volume managers. What do you think the next logical step is when there are too many volumes? Answer: HP's virtual array.

summary

HP's virtual array with HP's patented Virtual Technology is the industry's most intelligent disk array. Because of its unique ability to greatly simplify storage management, it significantly reduces overall IT management costs and practically eliminates mistakes caused by human error. The Virtual Technology also has a significant positive impact on real-world performance by automatically eliminating "hot spots," and by performance tuning the array 24 hours per day, 365 days per year.

for more information

For additional information on HP virtual arrays and other HP storage products and solutions, please call your local HP sales representative or visit our Web site at **www.hp.com/go/storage**.