

Down to earth report: HP 3PAR StoreServ storage

Enrico Signoretti

March 2013

Index

Executive summary	3
Introduction	4
A brief history	4
The product family	5
Architecture	7
A brief introduction	7
Controller Node	8
Why it is important for you	9
The ASIC	10
What is an ASIC?	10
How it is implemented by 3PAR StoreServ	10
Why it is important for you	10
Software	12
What is software?	12
How it is implemented by 3PAR StoreServ	12
Why it is important for you	13
Wide striping	14
What is wide striping?	14
How it implemented by 3PAR	14
Why it is important for you	15
Thin provisioning	16
What is thin provisioning?	16
How it is implemented by 3PAR StoreServ	16
Why it is important for you	17
Snapshots	18
What is a snapshot?	18
How it is implemented by 3PAR StoreServ	18
Why it is important for you	19
Multi tenancy	20
What is multi tenancy?	20

Jukuconsulting.com	
How it is implemented by 3PAR StoreServ	20
Why it is important for you	21
Automation	22
What is automation?	22
How it is implemented by 3PAR StoreServ	22
Why it is important for you	23
Replication	24
What is data replication?	24
How it is implemented by 3PAR StoreServ	24
Why it is important for you	25
Storage federation	26
What is storage federation?	26
How it is implemented by 3PAR StoreServ	26
Why it is important for you	27
Integration	28
What is integration?	28
How it is implemented by 3PAR StoreServ	28
Why it is important for you	29
Bottom line	30
Juku	31
Why Juku	31
Author	31

Executive summary

Nowadays, where business runs at internet speed, many companies are looking at new ways to store and manage tier 1 data, especially from the efficiency and automation standpoints. Databases and core applications transactions need to be processed as fast as possible while integration with hypervisors and ease of management are the base of day-to-day work. Ordinary storage technology is no longer enough to sustain high demanding requirements at a reasonable cost, particularly when we speak in terms of ROI and TCO (total cost of ownership). Next generation architectures and technology are the only way to achieve the best of the breed in performance and efficiency.

Many vendors are working on new solutions to address customers' requests. Actually, the "enterprise storage" market is quite segmented and addressing many different targets: it ranges from commodity hardware storage to pure flash systems. At the same time the DNA of disk arrays is quite different too: some products adopt 100% commodity hardware while improving the capabilities of traditional solutions via software, others develop no-compromise architectures for achieving the best in terms of performance and efficiency. And of course, in the middle there are a lot of nuances.

Maturity in these kind of products is another important key factor. In fact, enterprises are looking at performance but, at the same time, reliability and availability are fundamental to run 7x24 businesses without disruptions.

Eventually, the initial investment to acquire a next generation storage solution is easy to absorb due to the impact of the increase of overall performance and diminished management efforts that can positively affect all the business processes.

HP 3PAR StoreServ Storage systems are to be considered top of the notch in next generation enterprise Tier 1 arrays, with well implemented features and very high overall efficiency. With the introduction of HP 3PAR StoreServ 7000, The HP 3PAR StoreServ family can also be adopted by any kind of end user ranging from large enterprises to SMBs.

Introduction

A brief history

The HP 3PAR StoreServ product family comes from an acquisition that HP made in 2010. 3PAR was originally founded in 1999 by Jeffrey Price, Ashok Singhal and a third partner, Robert Rogers, who left the company a couple years later. The P, A and R in the name come from the initials of the founders. The mission of the company was “to make storage solutions simple and efficient”.



3PAR

The first product, called “InServ storage server”, was shipped in 2002 and the company released its first thin provisioning implementation in 2003. In

fact, the company was widely recognized as a pioneer in thin provisioning and other efficiency techniques that we now find in many other storage systems. 3PAR's primary competitors were EMC, HDS and IBM.

3PAR went through an initial public offering in 2007 and was listed on NYSE. After the acquisition of 3PAR, HP decided to move some 3PAR key figures to occupy important management roles in HP's storage division. The most visible one was David Scott, former 3PAR CEO and president, who now is in charge of the overall enterprise storage division. Scott has been working on the new HP storage strategy and rationalization of the offering after the acquisitions that his new company made during the past years. 3PAR StoreServ, as the product is now known at HP, has a flagship role in this strategy but it's clear that the company will also continue to develop other important assets of its portfolio like: StoreVirtual (P4000) SMB storage, StoreOnce backup solutions and the recently announced StoreAll (Object storage) family of products.

The product family

[The HP 3Par StoreServ family](#) of products is comprised of two models: [7000](#) and [10000](#). The first one is aimed at addressing midrange deployments while the latter is intended for bigger and most demanding installations. They share the same operating system, every feature and some of the fundamental components like high performance custom ASIC chips. The main differences are the number of supported controllers, front-end/back-end connectivity, scalability and performance.

One of the most important advantages of HP 3PAR StoreServ is that the end user can easily find the model that best fits his actual needs without thinking about long term future requirements: the same operating environment and storage federation features will grant the capability to mix different models and hardware versions of the product for a seamless migration experience to more complex installations, avoiding to throw away former investments.

HP provides many software features for its 3PAR StoreServ as well as a lot of plug-ins and integrations for major hypervisors, operating systems and DB platform. [HP 3Par Operating System](#) is licensed on its relative hardware.



Additional software features, like snapshots or replicas, are often offered in commercial bundles. The software catalog contains many sophisticated management and reporting suites. Integration suites cover major Hypervisors, OSs and DBs (like Microsoft Windows servers, VMware, Hyper-V, Microsoft SQL or Oracle).

The following specifications table ([also available on HP.com website](#)) shows a comparison between the two models, each model could have different configurations:

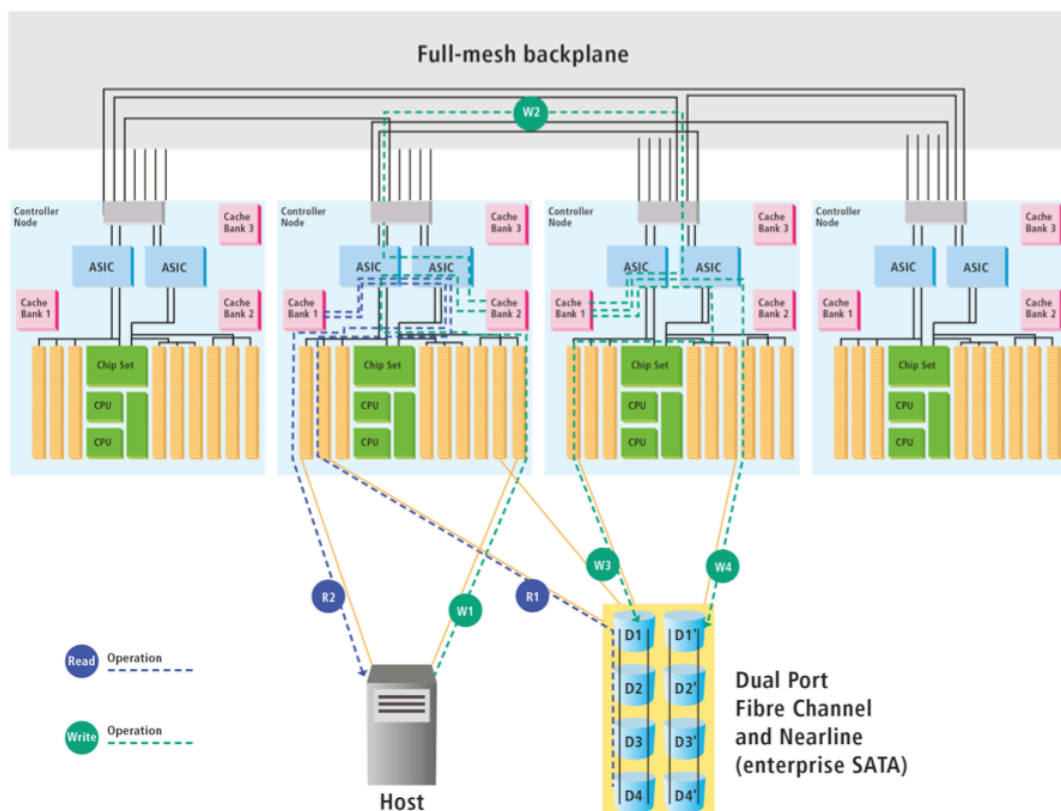
	HP 3PAR Storserv 10000	HP 3PAR Storserv 7000
Capacity	2.2PB	864TB
Drive type	LFF FC, LFF SATA, SFF SAS or LFF SAS	SFF SAS or LFF SAS
Host interface	8 Gb/sec Fibre Channel (192) Ports, 10 Gb/sec iSCSI (32) Ports maximum, depending on configuration/model	8 Gb/sec Fibre Channel (24) Ports, 10 GbE iSCSI (8) Ports Maximum supported, depending on configuration/model
Storage controllers	(8) Quad-core 2.8GHz P10000 Controller Node; Maximum, depending on model	(4) 7000 6-core 1.8GHz Controller Node; Maximum, depending on model
Availability features	Redundant power supplies and fans, Redundant batteries, A minimum of dual redundant controllers, RAID 1, RAID 5 and RAID MP for data protection, High availability cage and high availability magazine	Redundant power supplies and fans, A minimum of dual redundant controllers, with up to four controllers for added redundancy, RAID 1, RAID 5 and RAID 6 for data protection. The 4-node configuration provides cache persistence in the event of a controller failure.
Cache	512GB (+256GB control cache); Maximum supported, depending on model	64GB (+32GB control cache); Maximum supported, depending on model
Compatible OSs	MS Windows Server 2003/2008/2012, MS Windows Hyper-V, HP-UX, SUSE Linux Enterprise Server (SLES), Red Hat Enterprise Linux (RHEL), VMware ESX and ESXi, Oracle Solaris, Oracle UEK, Oracle Linux, Citrix XenServer, IBM AIX, HP OpenVMS	MS Windows Server 2003/2008/2012, MS Windows Hyper-V, HP-UX, SUSE Linux Enterprise Server (SLES), Red Hat Enterprise Linux (RHEL), VMware ESX and ESXi, Oracle Solaris, Oracle UEK, Oracle Linux, Citrix XenServer, IBM AIX, HP OpenVMS
warranty (parts/ labor/onsite)	3/3/3 years	3/3/3 years

Architecture

A brief introduction

The HP 3PAR StoreServ architecture combines commodity hardware with best in class performance components and innovative solutions. All 3PAR StoreServ models share a common design based on high speed, low latency, full mesh, passive backplane that connects all the controller nodes pairs together. The high speed, low latency interconnection links allow cache coherency and full coordination between nodes. In fact, controller nodes form a high available active cluster that presents itself as a single storage system to the hosts. Each controller node may have one or more available ports for external connectivity.

On the software side, the end user can create RAID protected volumes (called Virtual Volumes), exported as LUNs. Sophisticated virtualization and



internal mapping mechanisms of physical hard drives and SSDs grant a very high granularity and efficiency in provisioning and space management. A proprietary Gen4 ASIC chip is devoted to manage and speed up all these operations, offloading the controller nodes CPUs and enhancing performance.

The 3PAR StoreServ architecture is implemented in two different models: the high-end system 3PAR StoreServ 10000, and a midrange 3PAR StoreServ 7000. Each model could have different controller configurations: 10800 may accommodate up to eight controllers, the 10400 and 7400 are capable of supporting two or four controllers and the 7200 is the entry level system supporting only two controllers.

The quantity of the disks supported in the backend is a consequence of the number of controllers. 7000 is the newest model and uses next generation Intel exa-core CPUs and commodity 2.5" or 3.5" SAS trays. The 3PAR StoreServ 10000, which probably will receive an update in the next future, is still based on quad core CPUs and has both FC and SAS backend trays. HP recently added support for SAS drives to the 10000, which uses a highly available and very dense tray. Both 7000 and 10000 models support SSDs, 15K, 10K and 7.2K RPM disks. The size and format of the disks depends on the model.

The supported protocols are FC (8Gb/s) and iSCSI (10Gb/s), the number of available ports changes from model to model and it depends on the configuration of the system. FCoE is not supported at the moment, probably for lack of demand.

Controller Node

Controller nodes are key elements of the 3PAR StoreServ architecture. Each node is part of the active, highly available cluster. It offers computational power for data movements and IO connectivity. The cluster could be configured from two to eight nodes (it depends from the model and its configuration). The controller node has direct (2GB/sec) connections to all other nodes and it's coupled with another node to offer cache replication coherency and redundancy.

Controller nodes are equipped with Intel 4/6 cores CPUs and proprietary Gen4 ASICs (current products are using 4th generation chips). Control and data movement operations are processed simultaneously by CPUs and ASICs, each one with dedicated caches. This approach grants high speed

operations while avoiding bottlenecks. In larger configurations, each controller node can sport 36 ports for a total of 288 ports at maximum, 192 of these ports can be used for external connectivity.


Why it is important for you

3PAR StoreServ hardware architecture is simple and powerful. The low latency high speed direct connections between nodes enable high performance without the need of crossbar switches. Direct connection links could limit the overall size of the systems but I consider that the current 2.2PB of maximum capacity, 64 CPU cores and ASICs are substantially large numbers for the majority of end users. The huge amount of cache in these particular kind of systems is a must, but also helps to speed up the operations. Moreover, the simplicity and robustness of the design surely help the development of the operating system software and next generation features.

The ASIC

What is an ASIC?

ASIC (Application Specific Integrated Circuit) is a customized chip aimed to optimize specialized operations. This kind of chip has evolved a lot in the last few years and now



they can include microprocessors, memory, and other parts that we often find as separate components. The major advantages of these tailor-made integrated circuits are found in the cost reductions, big space and power savings. The significance is that you can fit more processing power in the same space at lower costs and complexity.

How it is implemented by 3PAR StoreServ

3PAR StoreServ is one of the few storage architectures in the market to develop its own ASIC, delivering fourth generation chips right now in its arrays. Fourth generation also means a certain level of maturity in the design and capabilities. 3PAR StoreServ's ASIC is a specialized component that performs very heavy tasks like RAID (5,6/MP) operations, CRC (Cyclic redundancy check) calculations, data movements (e.g. volume copies, snapshots management, etc.), thin provisioning/reclaim and zero block detection.

Why it is important for you

This chip does almost all the complex data calculations that you can find in a storage system. It helps to improve overall efficiency without having an impact on performance. In fact, many operations, as we are going to discuss

in the following chapters, are done at a very fine grained block size (16KB). Many vendors are avoiding the developing of these kinds of chips because of the high cost of R&D but the advantages to couple high specialized components like this to ordinary CPUs, is a no-compromise solution to obtain the best performance.

Using this powerful hardware, and subsequent fine granular block size, helps to achieve better overall space and performance efficiency of the systems giving the end user a better TCO.

Software

What is software?

When we speak about basic software in a modern storage system we can easily split the topic in two parts: the first one is the internal operating system and its management interface, the other is external management suites aimed at the management of one or more systems.

Storage vendors can propose many types of software titles to their customers (e.g.: plug-ins, connectors or other kind of integrations with server operating systems and applications) and we will discuss some of these in the following chapters.

The internal operating system of a modern disk array is the software that provides all the features, monitors the system's status, manages IO operations and data protection. The OS also provides the interfaces (APIs, CLI and GUI) for administration purposes.

The external management suites are more a comprehensive set of tools that usually manage more systems, collect data for performance analysis, do capacity planning, provide charge back functionalities and, often, much more than that.

Every vendor has its tools, many of these products are very focused on their arrays but, sometimes, they can manage third party arrays.

Software tools from different vendors have different licensing models and it could be quite difficult to compare features and costs of various offerings.

How it is implemented by 3PAR StoreServ

We have already described 3PAR StoreServ's hardware and its ASIC chips but the software has the same importance, especially because it is the enabler of the underlying hardware and this is what people see when they

manage the systems. All 3PAR StoreServ features, including those that rely on the ASIC chip, are available through its operating system: 3PAR OS (formerly InForm OS).

3PAR OS includes easy to use and coherent administration tools, like the 3PAR Management Console or the scriptable CLI, to help storage administrators managing their storage systems with limited effort. 3PAR OS also provides support to SNMP and SMI-S specifications.

3PAR Management Console is a Java based application that needs to reside on the same client of the CLI tools. It provides the GUI to manage all physical and virtual resources as well as a bunch of other basic tools for monitoring and reporting front-end and back-end activities. The software also controls the features that we are going to described in the next chapters (some of those features need to be enabled with an appropriate license key).

There are many additional HP 3PAR StoreServ software titles that can be installed to provide enhanced functionalities. They range from sophisticated monitor/reporting tools, to plug-ins that you can install on hosts to efficiently manage the interaction between OS/application and storage (e.g.: VMware recovery manager). The 3PAR StoreServ software catalog is extensive and still growing.

Why it is important for you

The application that every storage admin keeps open all the time on their desktop is the storage management tool: a coherent and easy to use interface is important to avoid errors, to have the right information in the right place and to speed up all management tasks.

At the same time sophisticated tools that enhance basic capabilities could be of great assistance in managing complex infrastructures as well as valid support for various business needs (e.g.: chargeback).

Last but not least, host based tools aimed at a proper integration between storage and upper software layers are fundamental elements to build modern and efficient infrastructures.

3PAR StoreServ has interesting solutions in all these three fields and HP supports all major software platforms with their integration tools.

Wide striping

What is wide striping?

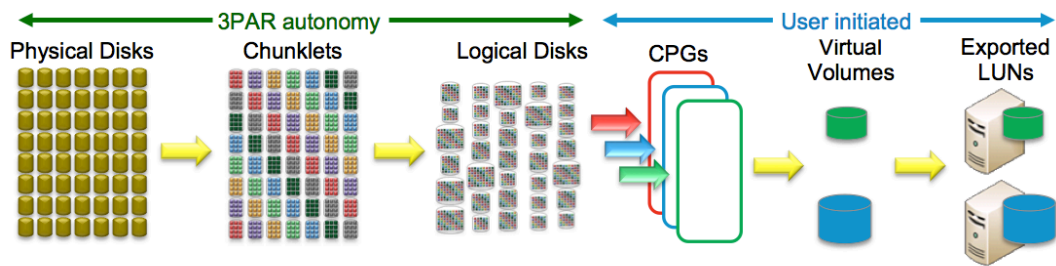
In legacy storage arrays, disks are organized in groups called RAID sets. Each RAID set/group has a limited number of available disks and they are formatted to act as a whole. Every written block (or maybe stripes of blocks) is split by the function of the number of the available disks and then a parity calculation is made. At the end, the result of this operation plus the parity will be written on physical disks. Some vendors still use this mechanism to protect and serve data but, RAID, and this way of working were born in the '80s and things have dramatically changed. In fact, LUNs that are presented to the hosts directly come from these groups and the performance of the single LUN is limited to the number of disks in the group. A second issue arises from the efficiency of the whole array: you can concatenate more LUNs (with an internal Volume Manager or at the host level) but you risk building a very complex structure to manage and with a lot of constraints.

More modern systems adopt another technique that I can sum up this way: every disk in the system or, at least every single type of disk, is part of pools. The disks are then segmented in chunks and are available to become part of the volumes. When the user allocates space for a LUN, a volume is created by picking the needed blocks from that particular pool. With this technique every single disk contributes with space and performance to the requested LUN; IO operations are spread on all array disks avoiding bottlenecks while making the most from the available resources.

How it implemented by 3PAR

3PAR StoreServ has developed a sophisticated disk virtualization layer that implements a true wide striping mechanism. Each physical disk is divided in small 1GB chunks (chunklets in 3PAR StoreServ naming convention) that are used to build LUNs and a distributed sparing mechanism.

The system automatically creates and manages Logical Disks (LDs) when needed: they are simply stripes of RAIDed chunklets and used for every kind of space request (Volumes, snapshots, etc.).



The user sees and can instantiate Virtual Volumes (VVs) from Common Provision Groups. CPGs are user created virtual pools of LDs and you can think of them as a profiling mechanism. Virtual Volumes are then exported to hosts as LUNs.

Explaining this mechanism is much harder than using it in day-to-day work!

Why it is important for you

3PAR StoreServ's disk virtualization layer brings many benefits to the end users. If compared to legacy arrays, a virtualized array is:

- Much easier to use,
- Granular in resources utilization,
- Particularly efficient in terms of performance.

This type of modern and elegant design brings a lot of advantages and allows implementation of innovative features without the limits and constraints that you find on traditional arrays. The advantages lie not only in performance or resources utilization but also, for example, in faster recoveries from failed disks. In fact, once that the system finds a failed disk, it proceeds to rebuild only the used blocks and spreads them on all the other disks instead of rebuilding the whole disk in a small RAID group.

Thin provisioning

What is thin provisioning?

Thin provisioning is a virtualization technique that gives the appearance of having more available space resources than are actually available. As opposed to fat (traditional) provisioning, it helps the user to not allocate unused resources by reserving space for potential future needs. Especially on legacy systems, the biggest problem of thin provisioning is the amount of controller's CPU power needed to manage it that can hurt the overall responsiveness of the array. Another big problem of some thin provisioning implementations is the lack of capability to stay thin. In fact, thin provisioning implementation needs to understand the type of written blocks and how to reclaim unused space to be effectively efficient. However some thin provisioning implementations have been vastly improved in the recent years and now, in some storage systems, they are the default option for creating volumes.

How it is implemented by 3PAR StoreServ

3PAR StoreServ is widely recognized as a pioneer in implementing thin provisioning (TP) techniques. The first 3PAR StoreServ thin provisioning feature was released 10 years ago and it's suddenly become the base of its product and marketing strategy.

The 3PAR StoreServ Gen4 ASIC chip has an important role in all the operations needed to manage TP: resources are provisioned and managed with a very small block size (16KB), on the fly and without any kind of preallocation. The ASIC detects zeroed blocks (blocks full of zeros) and doesn't actually write them on physical disks, saving even more space. 3PAR StoreServ Storage can also reclaim unused space: this feature is supported by some software vendors including Microsoft, VMware Oracle and Symantec.

Thin provisioning is implemented at the very core of the system, so that almost every service and feature (replication and snapshots, for example) extensively use it. It's also interesting to note that 3PAR StoreServ offers a free tool to import fat provisioned LUNs from legacy arrays and converts them into thin volumes during the importation.

Why it is important for you

Fat (preallocated) volumes are often the motive for scarce resource utilization, and many times it's easier to waste space instead of managing a good resource utilization policy.

In many environments, it is easy to find more than 50/60% of unused space due to this reason. This cost could severely impact the overall TCO of your storage system; wasted disk space means also wasted money in power, cooling and physical space in the datacenter.

3PAR's, ASIC based, thin provisioning is one of the best in class implementations because it is:

- Granular,
- Very efficient,
- Designed from the ground up to be and remain effectively thin.

Snapshots

What is a snapshot?

A snapshot is a point-in-time copy of a data volume. In practice it is a view of a volume at a given time. There are two kinds of snapshots: pointer based and copy-on-write and their implementation differ from vendor to vendor. Usually, pointer based snapshots are more efficient than C-O-W but a lot depends on how they are implemented. Nowadays snapshots are at the base of many integrations with operating systems, hypervisors and applications: they can contribute heavily to speed up many tasks like backups or the creation of copies of data for development/testing purposes. It is a must-have functionality for every enterprise storage system. Users need to understand their storage snapshots implementation before going in to production because limits and constraints of some implementations could severely hurt performance or data consistency.

How it is implemented by 3PAR StoreServ

3PAR StoreServ implements a copy-on-write snapshot (referred to as Virtual Copy) that uses Gen4 ASIC to obtain the same level of functionality as the best pointer based snapshots on the market. In fact, the ASIC does not suffer the slowdown introduced by the copy-on-write operations. Virtual Copy takes full advantage of the underlying virtualized layer. There is no need for space preallocation and the internal scheduler can automate many tasks (automation is discussed in the following chapters). Unlike many other storage arrays, 3PAR StoreServ snapshots may be writable, can be created from other snapshots and on different disk tiers. They can also be organized in consistency groups for a more reliable data consistency. HP has developed snapshot integrations with major software environments and OSs to obtain full copies (and restores) in a few clicks.


Why it is important for you

Snapshots are one of the most important features of a modern storage system. They can speed up many operations and save system administrators a lot of time, contributing to a faster ROI and a lower TCO. Efficient and fast snapshots system are at the core of a storage system's usability, and integration with upper layers is fundamental to get all the benefits: HP 3PAR Virtual Copy does it very well. 3PAR StoreServ also offers a CLI to help end users with the integration of its snapshots to legacy or custom software environments via scripts.

Multi tenancy

What is multi tenancy?

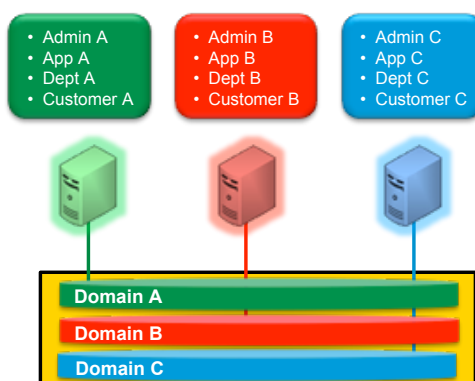
Multi tenancy is a hardware or software characteristic of IT systems. It refers to the capability of serving different clients (tenants) from a single host/system. In a multi tenant architecture the system is virtually or physically designed to be partitioned and serve each client with its own configuration, set of resources and data. Due to its characteristics these partitioning methods could also be considered a way to grant Quality of Service mechanisms to the clients.



Multi tenancy is one of the pillars of cloud computing architectures. There is no single way to implement multi tenancy, so it always results in competitive differentiation and discussions between vendors.

How it is implemented by 3PAR StoreServ

Unlike some other vendors 3PAR StoreServ implements a virtual partitioning mechanism called Virtual Domains. A super administrator can create up to 1024 partitions. Each partition acts as a “private virtual storage system” with its own resources: CPGs, volumes, hosts, etc. The super administrator can



enable administrators for virtual domains, allowing them to self-provision resources. All the Virtual Domains share ports, cache, CPUs, drives and other virtual resources so there is no need for preallocation to domains.

Why it is important for you

Storage partitioning could be implemented either in hardware or in software. Hardware partitioning is the best in terms of security, errors avoidance and higher availability. But, at the same time, it is very rigid and there is a concrete risk to waste useful resources due to physical limitations. In most of the cases hardware partitioning is also very limited in the number of partitions that can be actually realized on the system. On the other hand, software partitioning could be less effective in workloads and faults isolation but it is very elastic and gives the best results in terms of manageability and usability. Many cloud services providers have chosen 3PAR StoreServ for these reasons.

Automation

What is automation?

In recent years, especially on next generation storage systems, software features played a big role. One of the most interesting areas where software has been doing very well is the automation of repetitive tasks and processes. The kind of tasks that a storage system can do today based on time or event range from the capability to transparently move cold data on slower and cheaper disks, to automatic tuning or scheduling of some operations.

Obviously, the primary goal of automation is to minimize system management and lower TCO.

Some vendors developed these kind of features into their systems while others chose to deploy some functionalities through external appliances or software installed on external servers. Over time however the trend is to implement those features into the array likely due to the exceptional CPU power of modern CPUs, better reliability and ease of the development process. At the same time many software and operating systems vendors have exposed APIs to help storage systems in gathering events on the upper layers.

How it is implemented by 3PAR StoreServ

3PAR StorServ arrays implement many kinds of automations aimed to simplify a storage administrator's life. HP refers to this high level of automation as "autonomic" management. Some of those automatisms are embedded into flagship features. For example, space reclaim for thin provisioned volumes is automatically invoked when the operating system, or the hypervisor, says that some blocks were erased and they are no longer needed. Other features are available through an internal scheduler (snapshots creation and retention, for example).

One of the most interesting features of this family relies on automated storage tiering (called Dynamic Optimization and Adaptive Optimization). This implementation of automated tiering isn't the most granular of the market (it organizes data movements with 128MB blocks -regions- for user data and 32MB for Metadata) but benefits shown in real world cases by HP are quite good. Like other vendors, 3PAR

StoreServ offers an external software, called System Reporter, to collect historical data usage of all the resources (space, front-end and back-end IOPS, CPU utilization, and more). The data is stored in a DB and is available to generate the right information to move regions to the right kind of disks/RAID. The



movements are controlled by a series of algorithms that weigh IOPS quantity, latency, and used space in the volumes involved to find the best match. Movements can be scheduled or automated while the system is running and without service disruption.

Why it is important for you

Automation can be a huge time saver. The ability of the storage system to react to a changed workload and reorganize data to avoid hot spots may have tremendous positive impact on the quantity of time needed to manage the system. At the same time internal schedulers and other features help to maintain the system's efficiency at the higher levels with minimum efforts.

Replication

What is data replication?

Data replication features are at the base of many high availability and/or disaster recovery procedures. There are many ways to copy data on a second system but, when discussing hardware based replication, we have a few options that I could sum up with: synchronous, asynchronous and snapshots delivery.

The type of replication adopted depends on many business and technical factors and the results are slightly different. Async replication is the best in terms of RPO but it is also the most costly option and has rigid limits on distance and communication protocols between the storage systems. On the flip side, other replication methods achieve worse RPOs but they can be deployed on much longer distances (using TCP/IP) and at lower costs. In any case, replication impacts in both the primary and secondary system and needs to be carefully evaluated and configured before going in production.

One of the relevant aspects of remote replication also involves topology options: enterprise systems often offer N:N replication while mid-range systems are usually limited to 1:1. At the same time there are only a few vendors that are able to replicate data between their mid-range and enterprise class arrays.

Some vendors also require implementation services because of the complexity of the solutions.

How it is implemented by 3PAR StoreServ

The 3Par StoreServ replication capability (called Remote Copy) is implemented at the core of the system like other functionality that has been discussed in previous chapters. This means that it inherits all granularity and space efficiency of other features.

3PAR Remote Copy Software supports N:N replica configurations between all of its models (including older models) and it can do both sync and async replicas at the same time. All 3PAR StoreServ models are equipped with specific Ethernet ports and they can replicate data via FC or TCP/IP (using RCIP: remote copy IP). Having RCIP makes the lack of iSCSI non influential as supported protocol for replication.

3PAR StoreServ supports multi site traditional topologies (like cascading) but also a mix of synchronous and asynchronous replication, called long distance synchronous Remote Copy, that is more simple to deploy and maintain.

3PAR StoreServ systems extensively use snapshots for remote replication, not only as a primary volume for replication (optimizing data sent to the secondary site), but also to manage syncing and recovery operations with the automatic creation of temporary point-in-time backup copies before the execution of any important operation.

Why it is important for you

Remote Copy allows a great freedom of choice and ease of use for the end user.

The ability to couple smaller systems with bigger models and the few constraints imposed by the underlying technology makes this feature ideal in deploying complex configurations, and at long distances, with relatively minimal efforts.

Storage federation

What is storage federation?

Federation is a fairly new capability of modern storage systems. It's at the early stages of its development and every vendor is implementing it in different ways.

Some vendors have chosen to implement it through external appliances and building a virtualization layer on top of existing storage. Others are proposing federation as a feature of their systems. Each one of these approaches has its pros and cons and in many cases, the best solution depends on the user's needs.

The goal of storage federation can be summed up in the capability of two or more storage systems to transparently act as a whole. It's not like having 100% scale-out storage but storage federation offers some of the same benefits while maintaining all the characteristics of a scale-up approach. Seamless movement of data and workloads, automatic or manual load balancing, non disruptive migrations between two or more storage systems are among the main benefits of storage federation.

How it is implemented by 3PAR StoreServ

3PAR StoreServ storage federation functionality is called [Peer Motion](#). It's a software feature, supported on all models and doesn't need external hardware. Currently, Peer Motion allows only manual operations but expect a future smarter release with some form of automation.

Peer Motion user interface has been recently integrated into 3PAR Management Console and it needs access to CLI of the storage systems to run the needed scripts for data and workload movements. As you can imagine this is a pretty new software, but it works and has practical advantages.

Why it is important for you

The capability to federate two or more systems and manage data/workload movements between them, without service disruptions, has a great value for big end users who need to frequently manage migrations and resource reallocations.

HP 3PAR StoreServ is the first company to offer this functionality among all its storage systems, so that the end user may have the opportunity to think about future growth of the storage infrastructure in a more scale-up/scale-out way without making huge investments in advance.

Integration

What is integration?

Modern storage systems usually support many more features than legacy disk arrays but they are almost useless, especially on block storage, if they aren't integrated with software that resides in the servers.

Integration could mean many things and it depends on the host software your use. Obviously, the higher the software layer in on the stack, the more the integration involves a major number of components (in the middle of the stack) and the more complex the integration can be.

Storage software integration is a key element of a modern infrastructure, especially for virtualization. Hypervisor vendors provide a lot of APIs to integrate the storage layer with their software stack. It's quite important that the hypervisor and the storage system are aware of each other because this is the only way to get the best from both. Integration also means that the hypervisor knows its counterpart and trusts it, delegating very complex and CPU crunching data movement operations to the storage system.

In any case, both for traditional or virtualized infrastructures, the result of this kind of integration could represent a huge performance increase with great efficiency and savings.

How it is implemented by 3PAR StoreServ

3PAR StoreServ has a catalog of integration tools that covers all the major platforms from primary software vendors. These tools are basic plug-ins (almost all of them are free) and adaptors as well as more sophisticated software suites aimed to alleviate the Sys Admin from repetitive tasks or sensitive recovery operations.

HP 3PAR Recovery Manager is the commercial name of the above mentioned sophisticated software package. Recovery Manager addresses the integration between snapshots (and remote copy) with databases and hypervisors. The Recovery Managers orchestrates a number of operations that need to be done every time you need a consistent copy of your data. Because of the nature of 3PAR StoreServ snapshots and Recovery Manager software, you can take many copies of the data during the day, also scheduled at short intervals. If necessary, the same tool allows restoration of one snapshot with a few clicks to reinstate data at a given time or create a copy for other purposes.

Why it is important for you

If you have a very capable storage system but it is not integrated with host software then almost all its capabilities are useless. Fortunately this is not the case for 3PAR StoreServ.

HP 3PAR StoreServ software tools cover the most important operating systems, databases and hypervisors of the market. Recovery Managers can also be used with a CLI and included in scripts for a wider adoption in the end user's environment.

The ability to take snapshots of complex environments and create new full functional copies of those databases in a few clicks is priceless when you need to deal with development and test deployments or manage complex and time constrained backup policies.

Bottom line

HP 3PAR StoreServ Storage is a good option for any kind of block storage deployment from mid-range to high-end enterprise infrastructures.

Modern design, smart software features and ease of use are the key values of HP 3PAR StoreServ. Nonetheless, integration with major software vendors makes this solution even more usable while increasing its overall value.

These systems are designed from the ground up with efficiency and ease of use in mind and the architecture shows all you can expect from an enterprise tier-1 class storage system.

Efficiency and ease of use are also very important factors when the focusing on ROI and TCO. All 3PAR StoreServ features are designed to save resources while delivering the best performance from what is available. Another important contributor to ROI is the coherency of the entire platform, with 3PAR StoreServ, the customer can chose the right model for his needs and can migrate or put more systems together and take advantage of federation features also.

Juku

Why Juku

Jukus are Japanese specialized cram schools and our philosophy is the same. Not to replace the traditional information channels, but to help those who make decisions for their IT environments, to inform and discuss the technological side that we know better: IT infrastructure virtualization, cloud computing and storage.

Unlike the past, today those who live in IT should look around themselves: things are changing rapidly and there is the need to stay informed, learn quickly and to support important decisions, but how? Through our support, our ideas, the result of our daily interaction that we have globally on the web and social networking with vendors, analysts, bloggers, journalists and consultants. But our work doesn't stop there, the comparison and the search is global, but the sharing and application of our ideas must be local and that is where our daily experience, with companies rooted in local areas, becomes essential to provide a sincere and helpful vision. That's why we have chosen: "think global, act local" as a payoff for [Juku](#).

Author



Enrico Signoretti, Consultant, entrepreneur and passionate blogger (not necessarily in that order). Immersed into IT environments for over 20 years, his career began with Assembler in the second half of the 80's before moving on to UNIX platforms (but always with the Mac at heart) until now in the Cloud land. During these years his job has changed from deep technical roles to management and customer relationship management. In 2012 he founded Juku consulting SRL, a new consultancy and advisory firm very focused on supporting end

users, vendors and third parties in the development of their IT infrastructure strategy. He is constantly keeping an eye on market evolutions and continuously looking for new ideas and innovative solutions. You can find Enrico's social profiles here: <http://about.me/esignoretti>

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources Juku Consulting srl (Juku) considers to be reliable but is not warranted by Juku. This publication may contain opinions of Juku, which are subject to change from time to time. This publication is copyrighted by Juku consulting srl. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of Juku consulting srl, is in violation of Italian and European laws and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact Juku consulting srl (info@jukuconsulting.com).