



# Scalable Infrastructure with VMware® vStorage VMFS and 3PAR® Utility Storage

WHITE PAPER

## Overview

IT administrators today rely on infrastructure that is dynamically scalable and highly available. A scalable, resilient datacenter is achieved by augmenting or replacing hardware and software, based on the model of one application per server. This model not only results in underutilized resources but also creates hard-to-change boundaries that impair dynamic scalability and availability.

VMware®, with its partners such as 3PAR®, helps make it easier to transform a company's existing datacenter into a scalable and available infrastructure with server virtualization and dynamically tiered storage arrays. VMware vSphere™ 4 (“vSphere”) is the core virtualization platform from VMware that converts multiple servers and storage into a resource pool. VMware vStorage Virtual Machine File System (VMFS), which is included with vSphere, plays a key role in creating the foundation for virtualization-based distributed infrastructure services. It makes the virtual environment easy to provision, manage and scale by storing virtual machine files on the shared storage of a Fibre Channel or iSCSI SAN. This enables the IT administrator to spend less time managing the servers and storage environment, resulting in more time to complete other tasks and to innovate, which further reduces operating costs.

This paper focuses on how the advanced capabilities of both VMFS and 3PAR Utility Storage help simplify virtual machine provisioning and administration. In addition, it discusses how VMFS and Utility Storage together provide users with an ideal solution for expanding beyond the boundaries of a single server, to achieve the dynamically scalable and available IT infrastructure that is required by today's organizations.

## Overview of VMware vStorage VMFS

VMFS increases resource utilization by providing multiple virtual machines with shared access to a consolidated pool of clustered storage. Conventional file systems allow only one server to have read-write access to a specific file at a given time. This limits the scalability and the availability that are needed for this environment. In contrast, VMFS is a cluster file system that leverages shared storage to enable multiple instances of VMware ESX® 4 (VMware ESX) to read and write to the same storage concurrently. VMFS is optimized, tested and certified for a wide range of Fibre Channel and iSCSI SAN equipment, giving customers choice in leveraging block storage technology.

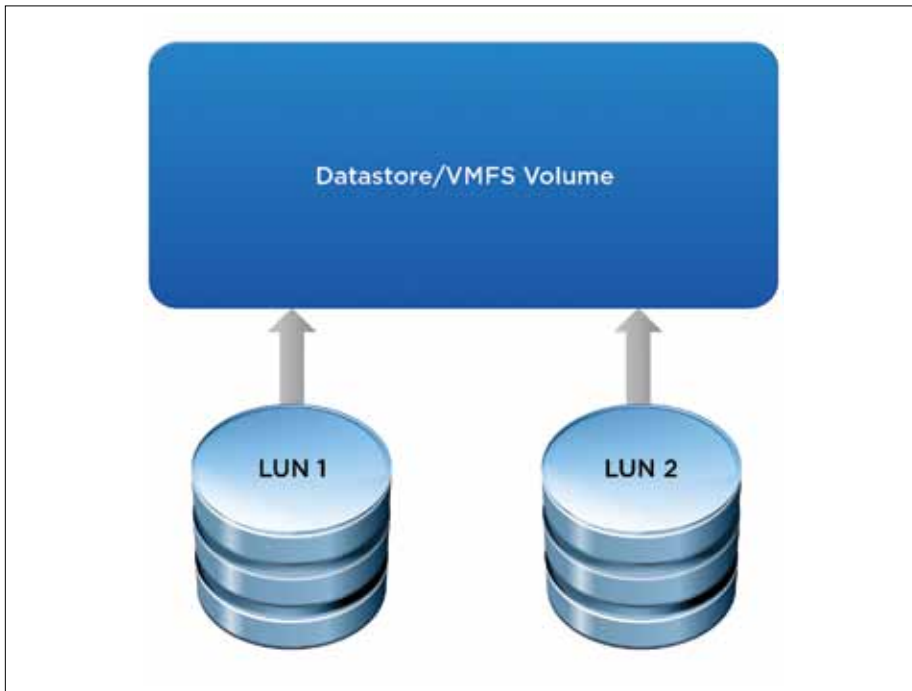
### Benefits of VMFS

With multiple instances of VMware ESX, VMFS enables access to the same virtual machine storage concurrently, enabling virtualization that scales beyond the boundaries of a single system. The distributed locking mechanism of VMFS facilitates the sharing of storage via these multiple instances of VMware ESX. This storage sharing enables virtual machines to dynamically and automatically migrate between VMware ESX servers.

VMFS is an automated file system that provides a hierarchical file system structure with user-friendly file-naming access. It automates the subdirectory naming process to make allocation, access and administration more functional. It efficiently stores the entire virtual machine state in a central location and can be created in advance, enabling instant provisioning of virtual machines by eliminating the need to provision storage for each new one.

### Scalable Infrastructure

VMware designed VMFS to provide management and scalability while maintaining optimal performance for the virtual environment. The single file system layer of VMFS (Figure 1) enables users to achieve higher server density from virtual machine to physical machine. VMware closely monitors any new virtual machine storage solution requirements and continually works to implement significant optimizations and advances to VMFS. The new features add to the high performance of VMFS and help ensure that customers and large enterprises can deploy even the most demanding applications in a virtual environment.



**Figure 1.** VMware vStorage VMFS manages and provides access to storage

VMFS uses adaptive block sizing for large I/Os, and subblock allocation for small files and directories. It supports large files while also simultaneously performing many small concurrent writes. This mixed-workload operation helps ensure optimal performance for any application, and it eliminates the need for an administrator to constantly tune the environment to maintain separation between different applications, providing superior service levels.

VMFS Volume Grow enables adding capacity to an existing datastore without disrupting running virtual machines. After a storage LUN that backs that datastore is expanded through the array management utility, the vSphere administrator can use VMFS Volume Grow to enlarge the VMFS extent on the expanded LUN. The newly available space then appears in the existing VMFS volume. To further enhance scalability, VMware vStorage Thin Provisioning enables overallocation of storage capacity for increased storage utilization, improved application uptime and simplified storage capacity management.

These VMFS capabilities provide the administrator with flexibility in managing growing storage capacity nondisruptively and efficiently. For more information on the VMFS components and how they work, refer to the VMware vStorage VMFS datasheet: <http://www.vmware.com/files/pdf/VMware-vStorage-VMFS-DS-EN.pdf>.

## Enhanced Management

Managing files in vSphere shared storage environments is similar to managing files and directories in a normal operating system. Most management functions that can be done to a file can also be done to the files that make up a virtual machine in a VMFS environment. VMFS enables you to copy, move and delete files the way you can in NAS systems. All of the files that make up a virtual machine are contained within a single subdirectory, which is called the VMhome. These files are the swapfiles, virtual disk files, snapshot files, configuration files and log file. Because VMhome is a single entity, it makes the management of all the virtual machines simple.

Functioning as a cluster volume manager, VMFS enables access to storage resources so that SAN storage (Fibre Channel and iSCSI) can be presented as volume on which virtual machines can reside. To store virtual disks, VMFS uses datastores, which are logical containers that hide specifics of storage from virtual machines.

Leveraging VMware vCenter™ Server, measurement of various aspects of the datastore space, including what storage capacity is allocated, not-shared and used, is made possible. Ease of use and navigation are enabled by the continued advancements in VMFS and their direct ties to VMware vCenter Server. This virtual machine and disk-level information is very transparent and makes it easier to troubleshoot and maintain a virtual environment in comparison to physical environments or other server virtualization environments. VMFS is optimized to support large files while performing many small concurrent writes. It helps reduce backup windows using virtual machine snapshots and it enables creation of a point-in-time copy of virtual machine data that can be used for testing, backup and recovery operations.

## Business Continuity

In SAN storage environments, VMFS enables portability and mobility of virtual machines across VMware ESX Servers, to provide high availability while lowering management overhead. As a clustered file system and clustered volume manager, its unique virtualization services enable the live migration of running virtual machines from one physical server to another. It also facilitates the automatic restart of a failed virtual machine on a separate physical server, VMware High Availability (HA).

VMFS provides on-disk locking to ensure virtual machine integrity and availability by preventing one from being run on more than one VMware ESX server at the same time. Furthermore, the locking mechanism of VMFS manages the storage access for multiple VMware ESX Servers and enables them to read and write to the same storage pool at the same time. It also provides the means by which VMware Distributed Resource Scheduler (DRS) can dynamically move an entire virtual machine from one VMware ESX Server to another in the same cluster to ensure that each one has access to appropriate compute resources at any point in time.

VMFS also provides snapshot functionality for point-in-time copy-on-write, which lays the foundation for advanced data protection capabilities for backup. Snapshots are leveraged by the VMware vStorage APIs for Data Protection and VMware Consolidated Backup, which is the foundation for scalable backups.

As VMFS stores a virtual machine's files in a single subdirectory (VMhome), disaster recovery, testing and cloning operations are greatly simplified, ensuring that the environment can be recovered during planned or unplanned failures. The entire state of the virtual machine can be remotely mirrored and easily recovered in the event of a disaster. VMware vCenter Site Recovery Manager enables IT administrators to automate the restart of virtual machines and also provides an easy way to test the disaster recovery solution.

FUNCTIONALITY	VMFS
Automated Filesystem w/ Journaling	Yes
Cluster File System	Yes
Cluster Volume Management	Yes
File Directories for Virtual Machines	Yes
Easily Copy, Move and Delete Virtual Machines	Yes
Thin Provisioning	Yes
Snapshot	Yes
Dynamic Volume Resizing	Yes

**Table 1.** Benefits with VMware vStorage VMFS

## Overview of 3PAR Utility Storage

Utility storage, a category of highly virtualized and dynamically tiered storage arrays, was developed as one approach to reduce the TCO of storage in utility computing and virtualized environments. The 3PAR InServ® Storage Server is an example of utility storage that 3PAR designed to be a scalable and efficient storage platform. The modular design of 3PAR Utility Storage enables it to scale from small to large deployments—addressing the needs from remote systems to large, centralized systems.

Through the use of virtualization management capabilities, 3PAR Utility Storage provides advantages in management, TCO, scalability and availability. These advantages help customers manage their storage environment with fewer human and capital resources while simultaneously providing the performance required by today's IT administrator.

### Reduce Storage Costs and Complexity

3PAR Utility Storage is designed to help users reduce many of the storage-related costs—both CAPEX and OPEX—while still providing the performance, scalability and reliability of block storage. To help lower storage costs and improve storage utilization, 3PAR Thin Provisioning applies dedicate-on-write technology to eliminate storage “waste” by enabling customers to safely overallocate storage capacity. Users can eliminate physical disk purchases immediately and postpone them until additional physical capacity is required. Building on thin provisioning, 3PAR Thin Conversion and 3PAR Thin Persistence help users migrate from “thick,” legacy volumes to thin volumes, and help them ensure that volumes stay thin over time by reclaiming unused or deleted blocks.

3PAR Autonomic Groups reduces the complexity associated with storage provisioning in clustered virtualized server environments. In just three steps, users can fully create and provision multiple volumes to multiple servers. They can create host, volume and even domain groups that automate and expedite storage provisioning, monitoring and data protection.

The 3PAR InForm® Operating System offers rapid provisioning to simplify and automate standard volume creation and management, relieving users of tedious preplanning chores and cutting standard storage provisioning to just two steps. 3PAR Utility Storage also reduces SAN costs by enabling users to simplify, delay or eliminate complex SAN infrastructures.

### Consolidate with Scalable Storage

3PAR storage addresses the need for a highly scalable storage system through a high-performance, modular architecture. 3PAR Utility Storage delivers consistent levels of performance as the number of volumes and drives increase up to the system's limits, by autonomically wide-striping data. The autonomic wide-striping from 3PAR automatically leverages the aggregate IOPS of all the physical spindles in the array. The resulting performance helps address the random I/O often associated with virtualized environments, where more applications are being deployed per physical servers than before.

3PAR InServ® arrays offer a mixed-workload architecture that enables it to scale with any workload placed on the storage. Through 3PAR InSpire® Architecture, transaction- and throughput-intensive workloads can coexist without contention on the same 3PAR storage resources. Control commands and data movement are processed separately, which eliminates the performance bottlenecks of many existing platforms when serving competing workloads. This capability is particularly important in environments where physical resources are extensively shared by different applications and workloads.

3PAR Dynamic Optimization helps the IT administrator scale the system while simultaneously maintaining optimal performance by providing a tool to nondisruptively spread an existing datastore across new physical resources. Users can alter a volume's underlying RAID protection level, drive type, stripe width and/or radial placement. 3PAR Policy Advisor for Dynamic Optimization adds further autonomic policy management and automation capabilities. Policy Advisor analyzes how the InServ's volumes use physical disk space and then automatically makes intelligent, nondisruptive adjustments to ensure optimal volume distribution and storage tiering across storage server resources.

Along with other VMware partners, 3PAR worked with VMware to support the VMware adaptive queue depth algorithm in vSphere. This feature dynamically adjusts the LUN queue depth in the VMkernel I/O stack, to minimize the impact of I/O congestion detected by the array. As a result, organizations can increase the number of virtual machines and add higher-performing applications to their VMware ESX Servers when attached to supported storage, such as 3PAR InServ Storage Servers. When combined, these performance capabilities help users scale out their environment as needed and maximize the utilization of their server and storage resources.

## Achieve Datacenter Protection

Both the hardware and software architectures of 3PAR storage are designed to offer high availability. Redundancy and online serviceability are built into every component. The 3PAR full-mesh, passive system backplane joins multiple controller nodes to form a cache-coherent, active-active cluster. Each controller node runs a separate instance of the operating system, providing software fault tolerance and ensuring availability of user data.

3PAR Persistent Cache is a resiliency feature built into the 3PAR InForm OS. It enables “always on” applications and virtual server environments, to gracefully handle a planned or unplanned controller failure. It eliminates the performance penalties associated with “write-through” mode, so that 3PAR arrays can maintain required service levels even in the event of a cache or controller node failure. 3PAR Persistent Cache preserves write caching by rapidly remirroring cache to the other nodes in the cluster in the event of a failure.

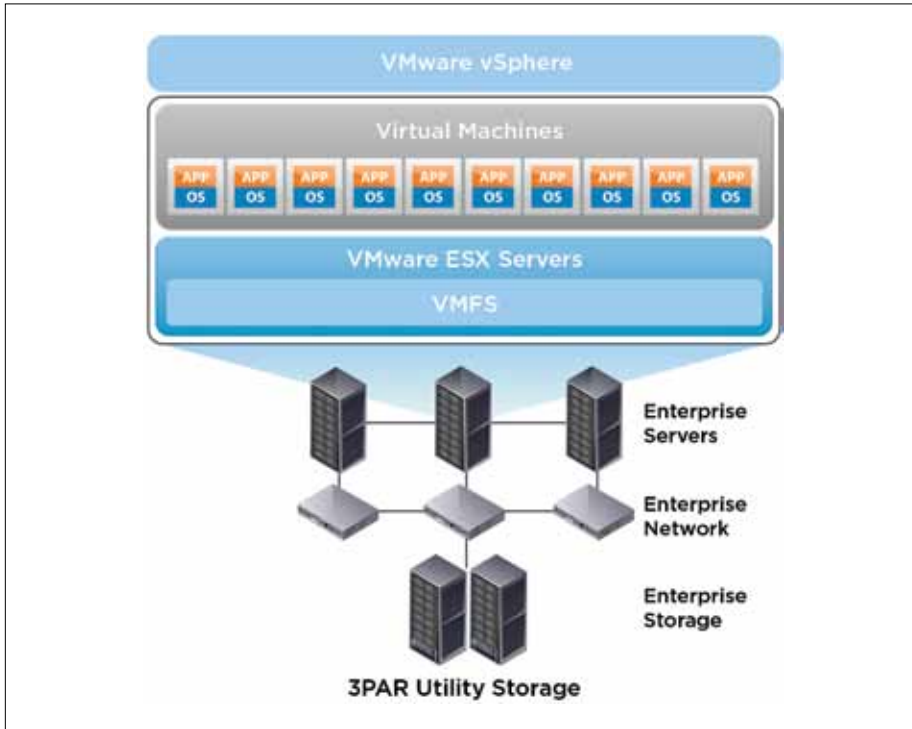
3PAR software solutions enable users to rapidly and efficiently create copies of data for protection of any application. 3PAR Virtual Copy is a reservationless, copy-on-write software product that enables users to create thin read-only and read-write snapshots. It is remote data replication software for disaster recovery and long-distance replication. Because it requires only a few commands and no professional services to enable, it facilitates rapid deployment. 3PAR Remote Copy integrates with VMware vCenter Site Recovery Manager (SRM) to help users simplify and automate disaster recovery across datacenters.

## Combining 3PAR Storage and VMware vStorage VMFS

The optimization and consolidation of many servers enabled by vSphere have increased resource utilization, which has led to new requirements when evaluating, managing and using physical servers and storage. Together, VMware vSphere and 3PAR Utility Storage help address the demanding requirements of virtual machines and the virtual datacenter with a scalable and available infrastructure that helps enhance both storage and server utilization. VMware and 3PAR will continue to innovate, enabling capabilities to meet the needs of today’s IT organizations.

Implementing a scalable infrastructure for an organization’s critical and noncritical applications remains one of the key goals of any IT organization. 3PAR storage is specifically architected to quickly process the various types of I/O that are common in a virtual datacenter, preventing it from being the bottleneck as the virtual datacenter grows. When it is used as the underlying storage for VMware vStorage VMFS, it provides optimal performance and virtual machine density. The combined solution helps organizations meet the scalability requirements for their mission-critical applications.

Simplified management is particularly important in today’s IT organizations, where the roles and responsibilities of the different IT administrators are beginning to blend. New infrastructure and virtualization administrators need to understand both their physical and virtual resources. The complementary capabilities of VMware VMFS and 3PAR Utility Storage—such as simplified provisioning, simplified management and nondisruptive expansion—help IT administrators quickly respond to changing business needs.



**Figure 2.** The VMware vStorage VMFS and 3PAR Utility Storage Stack

As virtualization drives up the efficiency of physical servers by sharing the CPU, memory and I/O resources across numerous virtual machines and applications, it is important that organizations maintain that efficiency at the storage layer as well—to achieve high utilization and minimized costs—a task not always possible with legacy storage systems and alternative hypervisor filesystems. The 3PAR wide-striping and high-parallelization capabilities help ensure that all of the physical storage resources are used equally by VMFS to maximize resource utilization. VMFS vStorage Thin Provisioning can be leveraged on its own or in conjunction with array-based thin provisioning, such as 3PAR Thin Provisioning.

Finally, with the high levels of consolidation and virtual machine density achievable with VMFS, the underlying SAN and block storage must offer complementary high levels of protection and availability in addition to those offered by vSphere. In a virtualized environment, a single storage array can contain hundreds, if not thousands, of virtual machines for an enterprise environment. The reliability and availability of the infrastructure are critical. The availability features of 3PAR Utility Storage, such as the 3PAR mesh-active architecture and 3PAR Remote Copy support for VMware vCenter Site Recovery Manager, and those of VMFS, such as VMware vMotion™ and VMware Storage vMotion, combine to provide the high availability demanded by an enterprise's virtual environment.

## Conclusion

The innovative solutions from both VMware and 3PAR are helping transform today's traditional datacenters into a scalable and available infrastructure. VMware vStorage VMFS simplifies virtual machine provisioning and administration with a high-performance cluster file system optimized for virtual machines. It offers users an ideal solution for expanding beyond the boundaries of a single server. By providing multiple virtual machines with shared access to a consolidated pool of storage, it also helps increase resource utilization. 3PAR Utility Storage can further improve these utilization benefits with advantages in management, TCO, scalability and availability. It combines with VMware vStorage VMFS to enhance a user's datacenter and deliver the infrastructure required by today's IT administrator.

### White Paper Authors

**Michael Haag**

Senior Alliance Manager, 3PAR

**Ibrahim "Ibby" Rahmani**

Senior Product Marketing Manager, VMware

