



Seven Myths Surrounding NFS in VMware Storage Environments

Introduction

There are two main technology approaches to operating a VMware based datastore in a virtual environment. The first is VMFS (virtual machine file system) — VMware's own purpose-built file system — running on a Fibre Channel SAN (storage area network and iSCSI storage array). The other is NFS (network file system) running on NAS (network attached storage) systems. Both have their respective pros and cons but NFS in particular has been dogged by a number of persistent myths largely based on historic issues rather than present day realities that have unfairly prejudiced its cause. This article discusses the relative merits of both, examining each of the NFS myths in detail, and explains why they do not stand up to serious scrutiny.

VMFS Vs NFS: a distorted playing field

As might be expected for a technology that has been built from the ground up to manage a VMware-based storage environment VMFS is robust, mature and well understood. It provides the VMware administrator with a fair amount of independence from the storage administrator for shared infrastructure. However, VMFS has some well-known shortcomings as follows:

1. Scaling requires careful, manually intensive planning when deploying many virtual machines.
2. Troubleshooting I/O performance issues on VMFS is a time consuming and often wearisome exercise because of the inherent lack of automation
3. As VMware virtualized ecosystems get bigger, issues 1 & 2 intensify at a far greater pace. A VMFS storage ecosystem is very inefficient for large numbers of VMs (hundreds to thousands).

By contrast the advantages of NFS datastores are well documented. An NFS based NAS storage solution is a much simpler, more automated, and altogether superior storage ecosystem for VMware datastores. It should also be acknowledged that there are also some substantial obstacles with most NFS storage solutions that need to be overcome before the promised benefits can be realised. However these are not insurmountable yet all too often it is the persistent, false mythology that surrounds virtualization environments and NFS datastores that discourages VMware administrators from pursuing a solution. In some cases the myth is based on past facts that no longer apply to VMware today. Other myths contain an element of truth but they cannot be universally applied to all NFS storage solutions and it is important to clear this up so that VMware administrators can form a balanced view of the choices available.

Myth 1. VMware does not support all advanced functions on NFS

This is simply not true. All VMware vSphere and ESX features are supported with NFS datastores. VMware has in fact fully supported NFS as a standard storage option since ESX 3.0. The myth is probably down to VMware's development cycle. VMware first develops and supports new advanced features on VMFS datastores with SAN storage. NFS support usually follows within a quarter or two. In exceptional circumstances the delay can be longer — for example Storage Vmotion™ and Site Recovery Manager™ (SRM) were delayed by up to a year because of the impact of a major new release (vSphere 4).

Myth 2. VMware performance is much slower with NFS datastores

This myth is perhaps the most enduring, pervasive and damaging since it is also quite false. The assertion is often repeated by some SAN vendors and experts alike to suit their own ends. However there is plenty of evidence to the contrary. VMware ran tests that showed NFS running on 1Gbps Ethernet performance delivered less than a 9 to 10 per cent decrease from



that of 4Gbps Fibre Channel. This test did not evaluate performance when 1Gbps Ethernet NICs are trunked on both the VMware server and the NFS storage system. Tests by Dell® as well as BlueArc have demonstrated that the performance of VMware using NFS running on 10Gbps Ethernet versus VMFS utilizing 8Gbps Fibre Channel is equivalent or even higher.

Myth 3. VMware CPU load is significantly higher with NFS

This myth originated following tests that compared the VMware CPU load when utilizing NFS datastores and TCP/IP protocol with the VMware CPU load of the Fibre Channel and iSCSI (software and hardware based) protocols. The tests showed NFS TCP/IP overhead ranged from 15 to 40 per cent more CPU overhead than the Fibre Channel protocol/drivers. A key problem with this test, and one that completely invalidates it, is that it did not measure the VMware CPU load for VMFS datastores when utilizing Fibre Channel or iSCSI protocols/drivers. It only measured the protocol/driver overhead. The overhead of VMFS was not measured. Since NFS datastores eliminate VMFS datastore overhead, comparing the overall CPU loads would have been far more insightful. Even so differences in the test results were not statistically significant considering Fibre Channel protocol is all silicon based, whereas the NFS and TCP/IP protocols were all software based. In any case solutions are now available that put both NFS and TCP/IP protocols in silicon thereby eliminating any, and all, CPU load issues.

Myth 4. VMware is limited to only 8 NFS datastores

This myth is based on ESX and vSphere NFS datastore setup defaults. In fact ESX supports up to 32 NFS datastores and vSphere supports up to 64.

Myth 5. VMware NFS datastores only scale to 16TB

The notion of a 16TB NFS datastore limitation is nothing to do with VMware; it's the typical file system size limit for most NFS storage systems. When the NFS storage system has this limit, then so do the NFS datastores. NFS file systems up to 256TB are available with a global name space that provide up to a 4PB name space. This means a datastore can grow up to 4PB or 256 times larger than a 16TB datastore.

Myth 6. NFS thin provisioned VMDKs automatically rehydrate when moved or cloned

This is true for ESX and but not for vSphere. Thin provisioning is the default setting for NFS VMDK datastores with vSphere 4 and no longer requires rehydration to be moved by Vmotion™ to a different NFS datastore or to be cloned.

Myth 7. Microsoft Windows® VMs can't boot or utilize NFS datastores

Microsoft Windows server doesn't really support NFS and cannot boot from NFS. Many people wrongly assume this means that when it runs as a VM in VMware that VMware cannot utilize NFS datastores. VMware built the NFS protocol into the ESX and vSphere virtualization layer. Microsoft Windows VMs never see the NFS protocol and are unaware that they are utilizing NFS datastores.

Other issues

As good as NFS datastores are for VMware, there are some other non-trivial issues associated with the way most (but not all) vendors implement NFS in their storage systems. Such systems will typically have:

1. A 16TB file system that limits the size of NFS datastores to 16TB
2. Software based NFS that severely limits the I/O performance of individual ports.
3. File system file object management limitations
4. Manual storage tiering

Conclusion

NFS based Networked Storage dramatically simplifies VM storage provisioning, data protection, management, task management, and troubleshooting. But issues with NFS datastore size, per port I/O performance and annoying file object limitations, not to mention manually intensive, application disruptive storage tiering and data migration, have tended to discourage VMware simplification. Companies should ignore the myths and look for solutions that combine NFS based NAS storage with intelligent data management software. In so doing they will be able to derive the full value of VMware and NFS while eliminating common operational limitations.



BlueArc Corporation
 Corporate Headquarters
 50 Rio Robles Drive
 San Jose, CA 95134
 t 408 576 6600
 f 408 576 6601
 www.bluearc.com

BlueArc UK Ltd.
 European Headquarters
 Queensgate House
 Cookham Road
 Bracknell RG12 1RB, United Kingdom
 t +44 (0) 1344 408 200
 f +44 (0) 1344 408 202