

# What Is New for Storage in Virtual Infrastructure 3 Release 3.5



Version 1.0



## Table of Contents

<b>What Is New for Storage in Virtual Infrastructure 3 release 3.5 .....</b>	<b>3</b>
<b>Storage VMotion .....</b>	<b>3</b>
<b>N-Port ID Virtualization .....</b>	<b>4</b>
<b>InfiniBand Support.....</b>	<b>4</b>
<b>SATA support .....</b>	<b>5</b>
<b>Round-Robin Load Balancing.....</b>	<b>5</b>
<b>VMotion with Local Swap Files .....</b>	<b>5</b>
<b>Network Based Storage Performance Enhancements.....</b>	<b>6</b>
<b>Conclusion.....</b>	<b>6</b>

## What Is New for Storage in Virtual Infrastructure 3 release 3.5

With the launch of VMware Infrastructure 3 version 3.5 comes a new version of VMware's hypervisor, ESX Server 3.5. One of the core focus areas for Virtual Infrastructure 3.5 is to enable greater uptime and agility in the management of the virtual infrastructure environment. VMware has addressed this by bringing new levels of flexibility and connectivity options to Virtual Infrastructure 3 version 3.5.

This document provides information about the following storage enhancements in VMware Infrastructure 3 version 3.5:

- Storage VMotion
- N-Port ID virtualization support
- InfiniBand support
- SATA support
- Round robin load balancing
- VMotion with swap files on local storage
- Network based storage performance enhancements

---

### Storage VMotion

Storage VMotion (SVM) enables live migration of virtual machine disks from one datastore to another with no disruption or downtime. Just as VMware VMotion allows IT administrators to minimize service disruption due to planned server downtime, Storage VMotion allows them to minimize disruption by reducing the planned storage downtime previously required for rebalancing or retiring storage arrays. Storage VMotion simplifies array migration and upgrade tasks, and reduces I/O bottlenecks by moving virtual machines while the VM remains up and running. It provides a hot migration of the storage location on which the vmhome resides.




Figure 1- Storage VMotion allows a hot migration of the storage location.

Bringing this live migration capability to the 3.5 release of ESX Server required a few tradeoffs. For one, the only fully supported storage source and target are Fibre Channel datastores. In addition, Storage VMotion is not fully integrated within Virtual Center and can be invoked only through the Remote Command Line option.

Options and instructions on how to download the RCLI and invoke Storage VMotion, can be found at the following link: <http://www.vmware.com/download/download.do?downloadGroup=VI-RCLI>

For more details on Storage VMotion, see the Basic System Administration guide located at:

[http://www.vmware.com/pdf/vi3\\_35/esx\\_3/r35/vi3\\_35\\_25\\_admin\\_guide.pdf](http://www.vmware.com/pdf/vi3_35/esx_3/r35/vi3_35_25_admin_guide.pdf)

---

## N-Port ID Virtualization

ESX Server 3.5 introduces support for N-Port ID Virtualization (NPIV) for Fibre Channel SANs, which enables each virtual machine to have its own World Wide Port Name (WWPN). This feature enables per-virtual-machine traffic monitoring using third party tools for chargeback. In a future release, NPIV is planned to enable per-virtual-machine LUN masking capabilities.

Both Emulex and QLogic offer NPIV-enabled HBAs that are supported with the default drivers included in the ESX Server 3.5 release. This feature is supported with Raw Disk Map devices and not integrated with VMFS in the current 3.5 release.

For more details on configuring and troubleshooting NPIV-enabled ports, please refer to:

[http://www.vmware.com/pdf/vi3\\_35\\_25\\_npiv\\_config.pdf](http://www.vmware.com/pdf/vi3_35_25_npiv_config.pdf)

---

## InfiniBand Support

ESX Server 3.5 also includes technical support for InfiniBand based HCAs (Host Channel Adapters) from Mellanox. This will be provided via a community support effort: VMware will take the level 1 support call, while all subsequent support will be handled by the HCA vendor (in this case, Mellanox). InfiniBand HCA will not be listed on VMware's I/O device HCL.

ESX Server 3.5 includes private extensions that allow an OFED InfiniBand stack to be plugged in and used. It is expected that these extensions will be available and supported for all VMware co-development partners in the future. Until that time, the extensions are supported for Mellanox, and potentially for others who work directly with VMware to test their offerings.

As with third party technologies not shipped directly by VMware, Mellanox's InfiniBand products will be covered under the third party hardware and software support policy, defined at:

<http://www.vmware.com/support/policies/ThirdParty.html>.

---

---

## SATA support

ESX Server 3.5 supports selected SATA devices connected to dual SAS/SATA controllers. For an up-to-date list of supported dual SAS/SATA controllers, see the guide to Compatibility Additions and Modifications for ESX Server 3i version 3.5; this guide is located at: [http://www.vmware.com/pdf/vi3\\_io\\_guide.pdf](http://www.vmware.com/pdf/vi3_io_guide.pdf)

At present (when this white paper is published), the current listing of dual SAS/SATA controllers supported by ESX Server 3.5 is as follows:

- mptscsi\_2xx (PCIE)      LSII078
- mptscsi\_pcie            LSII068E (LSISAS3442E)
- mptscsi\_pcix            LSII068 (SAS 5)
- aacraid\_esx30          IBM serveraid 8k SAS controller
- cciss                     Smart Array P400/256 controller
- megaraid\_sas            Dell PERC 5.0.1 controller

This SATA support is limited to ESX Server 3.5 and is not included with ESX Server 3.0.x.

---

---

## Round-Robin Load Balancing

ESX Server release 3.5 expands native load balancing by providing experimental support for round-robin load balancing of HBAs. Round-robin load balancing can be configured to switch to another path to the SAN; the switch can be based on the number of I/Os or megabytes sent down the path.

This feature can be enabled with default settings via the VI Client. While round robin load balancing can be selected through VirtualCenter, only default configuration options are available. More advanced options for policy based switching parameters are available via the CLI.

For more information on using this feature, see the *Round-Robin Load Balancing* technical note located at [http://www.vmware.com/pdf/vi3\\_35\\_25\\_roundrobin.pdf](http://www.vmware.com/pdf/vi3_35_25_roundrobin.pdf).

---

---

## VMotion with Local Swap Files

In previous versions of VMware Infrastructure, the operation of VMotion required swap files for virtual machines to be located on shared storage. VMware Infrastructure 3 release 3.5 now provides for swap files to be stored on local storage, while still facilitating VMotion migrations for these virtual machines. Users can configure a swap datastore policy at the host or cluster level, although the policy can be overwritten by the virtual machine configuration.

During a VMotion migration or a failover for virtual machines with swap files on local storage, if local storage on the destination is selected, the virtual machine swap file is recreated there. Creation time for the virtual machine swap file depends on local disk I/O. With VMware HA, creation time can also be affected if too many virtual machines are starting concurrently because of an ESX Server host failover.

VMotion migration of virtual machines with local swap files is supported only across ESX Server 3.5 hosts and later, and with VirtualCenter 2.5 and later.

---

## Network Based Storage Performance Enhancements

Virtual Infrastructure release 3.5 has several networking enhancements that, although fully supported in the networking itself, are not currently fully supported for network storage. These features include jumbo frames, TCP segmentation offload, and 10 gigabit Ethernet. Each of these features is described in more detail in the Performance Enhancements In Virtual Infrastructure 3 Release 3.5 Technical Paper posted on VMware Technology Network (VMTN) (<http://www.vmware.com/resources/techresources/>).

Each of these new networking features provides increased efficiency and throughput for network traffic and can offer significant improvements in storage network access for Ethernet based datastores. At this writing, these features are considered experimentally supported for use in datastores on NAS and iSCSI. This will be further tested by VMware's QA group and should be fully supported for network based storage in a future Virtual Infrastructure release.

---

## Conclusion

The new release of VMware Infrastructure 3 release 3.5 brings many new capabilities that increase the flexibility and ease of management for virtualization environments. These enhancements provide new means to increase the performance and uptime for workloads run in virtual machines.

This paper has provided an overview of new features that are specific to the storage interface of ESX Server. More details of how these features can be installed, as well as the benefits gained by leveraging these new capabilities, can be found in the ESX Server documentation. While this paper has addressed only some of the enhancements, many other features are documented in similar tech notes, which can be found on the VMware Technology Network (VMTN) at: <http://www.vmware.com/resources/techresources/>



VMware, Inc. 3401 Hillview Drive Palo Alto CA 94304 USA Tel 650-427-5000

[www.vmware.com](http://www.vmware.com)

© 2008 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156, 6,795,966, 6,880,022, 6,961,941, 6,961,806 and 6,944,699; patents pending. VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

WP-038-PRD-01-01 – Paul Manning, January 2008

