

# IBM System Storage SAN Volume Controller (SVC) and IBM Storwize V9000

*iSER host-attach configuration  
guidelines for VMware vSphere ESXi*

---

## Overview

### Challenge

System performance depends on how fast the data is served. The time required to serve the data is commonly referred as latency. Latency depends on the interconnect used and its software I/O stack in data packets processing. For Ethernet, the latency depends on the software processing in the implementation and the TCP/IP protocol stack which directly affect latencies. There are additional software processing delays in the iSCSI software stacks which add latency. This increased latency affects the overall system performance.

### Solution

Remote Direct Memory Access (RDMA) gives much more benefit in latency, IOPs, and bandwidth on Ethernet compared to traditional iSCSI. Using iSER as a host-attach protocol gives all RDMA benefits to the host subsystem such as VMware, where mission-critical workloads are hosted.

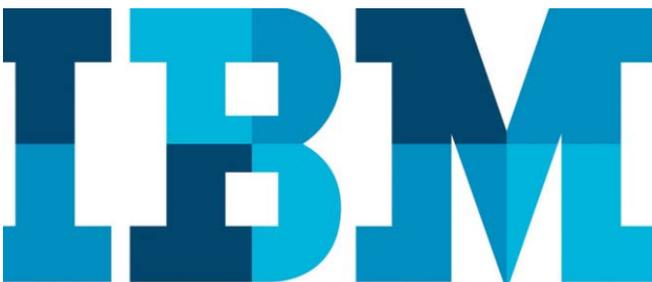
## Introduction

IBM® Spectrum Virtualize V8.2 delivers support for IBM FlashSystem® 9100 hardware enclosures and iSCSI Extensions for RDMA (iSER) host support. This white paper attempts to brief iSER host-attach configuration guidelines for VMware vSphere ESXi.

### All-flash storage

An all-flash storage array is a storage infrastructure that contains only flash memory drives instead of spinning-disk drives. The primary motive for the data centers to adapt the all-flash storage was an attempt to meet the performance and capacity demands of modern workloads

Fiber Channel has traditionally been the preferred interconnect for all-flash storage. However, 21st century data center paradigms such as cloud, analytics, software defined storage, and so on are driving a definitive shift towards Ethernet infrastructure that includes Ethernet connectivity for both server and storage. As the speed of Ethernet rapidly increase to 25 Gb/40 Gb/50 Gb/100 Gb and so on, it becomes more and more important as an interconnect to all-flash storage where lower latency performance is needed. While traditional iSCSI has gained significant ground as Ethernet interconnect to storage, inefficiencies in the TCP/IP stack do not allow it to be the preferred interconnect for all flash storages.



---

## Architectural overview

### Software

- IBM Spectrum Virtualize V8.2 or later code base
- VMware ESX 6.7 operating system

### Hardware

- IBM FlashSystem 9100 hardware enclosures
- iSCSI Extensions for RDMA (iSER) host support

### Network

- A 25 GB network adapter on both the initiator and target systems
  - Cisco 3232C switches for interconnect
- 

## iSER

To improve the performance of data-intensive workloads in the data center, you can use the iSER protocol on a higher speed Ethernet. iSER is a new standard for extending iSCSI with Remote Direct Memory Access (RDMA). iSER increases performance and reduces network latency by removing the need for Transmission Control Protocol/Internet Protocol (TCP/IP).

iSER provides an interconnect that is very capable of rivaling Fiber Channel as the all-flash interconnect choice. It leaves the administrative framework of iSCSI untouched while mapping the data path over RDMA. As a result, management applications such as VMware® vCenter, OpenStack, and so on continue to work as is, while the iSCSI data path gets a speed boost from RDMA. Depending upon how RDMA is implemented and depending upon the use cases for which it is deployed, there are two variants of iSER available in industry.

- **iWARP (Internet Wide-area RDMA Protocol)** - Implements RDMA for efficient data transfer over Internet Protocol (IP) networks. It is layered on Internet Engineering Task Force (IETF)-standard congestion-aware protocols such as TCP and Stream Control Transmission Protocol (SCTP).
- **RDMA over Converged Ethernet (RoCE)** - Implements RDMA over converged Ethernet. There are two version of this protocol:
  - RoCE v1- Ethernet link layer protocol which allows communication between any two hosts in the same Ethernet broadcast domain
  - RoCE v2 - Internet layer protocol which means that RoCE v2 packets can be routed

IBM SVC and Storwize V9000 IBM System Storage® SAN Volume Controller (SVC), IBM FlashSystem V9000, and IBM Storwize® family deliver expanded connectivity with increased performance through two new 25 Gb Ethernet iSCSI adapter cards with either iWARP RDMA or RoCE protocol support.

The fact that the family supports both iWARP and RoCE contributed in boosting the popularity of the product line.

## VMware vSphere ESXi

With the release of vSphere 6.7, VMware added iSER as a native supported storage protocol to ESXi. With iSER run over iSCSI, users can boost their vSphere performance just by replacing the regular Network Interface Cards (NICs) with RDMA-capable NICs.

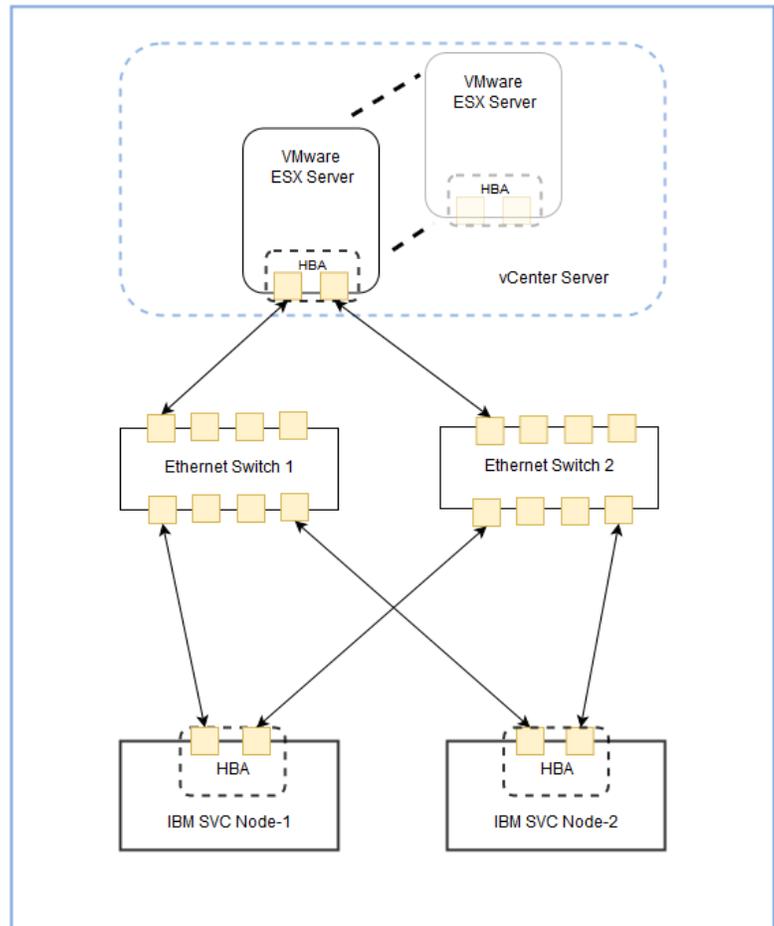


Figure 1. Configuration example

## Setup overview

Use the `esxcli` command to enable the VMware iSER adapter. Prerequisites to set up an ESX server for iSCSI and iSER:

- Make sure that your iSCSI storage supports the iSER protocol.
- Install the RDMA-capable adapter on your ESXi host.
- Use the RDMA-capable switch.

Procedure to set up the ESX server for iSCSI and iSER:

1. Run the following command in ESXi Shell to enable the iSER adapter  
`esxcli rdma iser add`
2. Verify that the VMware iSER adapter has been added.
3. Navigate to the host.
4. Click the **Configure** tab.
5. Under **Storage**, click **Storage Adapters**, and review the list of adapters.  
If you enabled the adapter, the adapter (vmhba#) appears on the list under the **VMware iSCSI over RDMA (iSER) Adapter** category.

## Configuring the Storwize FS9100 system

Perform the following tasks to configure the Storwize FS9100 system:

1. Use the management GUI to configure the Ethernet ports for iSER host-attached connectivity.
2. Use the management GUI to create iSER host objects.
3. Use the management GUI to create volumes and map those volumes to the iSER host objects created earlier.

Refer to [Configuring the Storwize FS9100](#) in IBM Knowledge Center for detailed instructions on these tasks.

## Setting up the ESX network for iSER

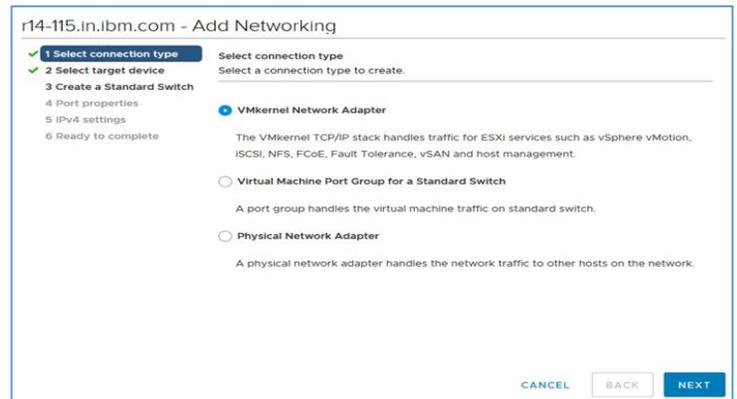
VMware iSER adapters depend on the VMkernel networking. Configuring the network connection involves creating a virtual VMkernel adapter for each physical network adapter. You can use a 1:1 mapping between each virtual and physical network adapter.

### Create a single VMkernel adapter for iSER

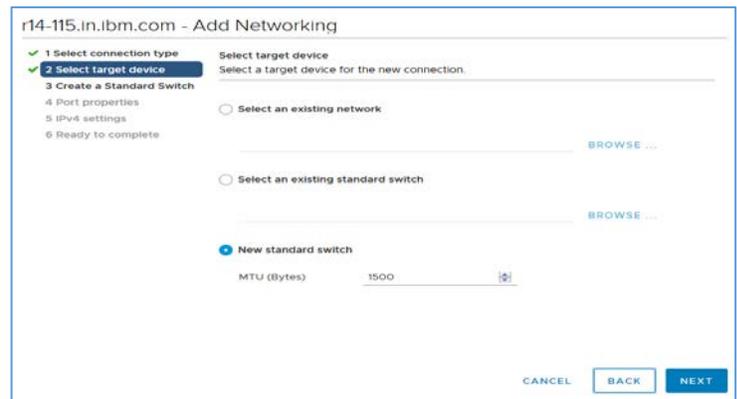
Connect the VMkernel, which runs services for iSER storage, to a physical network adapter.

Procedure to create a VMkernel adapter:

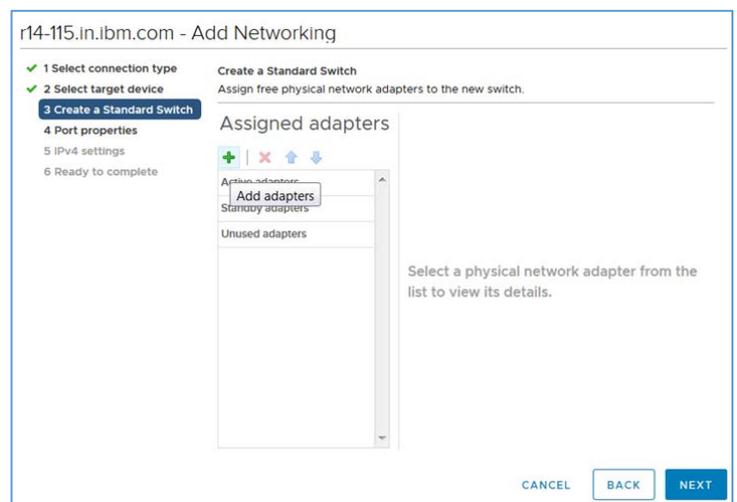
1. Navigate to the host.
2. Right-click the host and click **Add Networking**.
3. On the Select connection type page, select **VMkernel Network Adapter** and click **Next**.



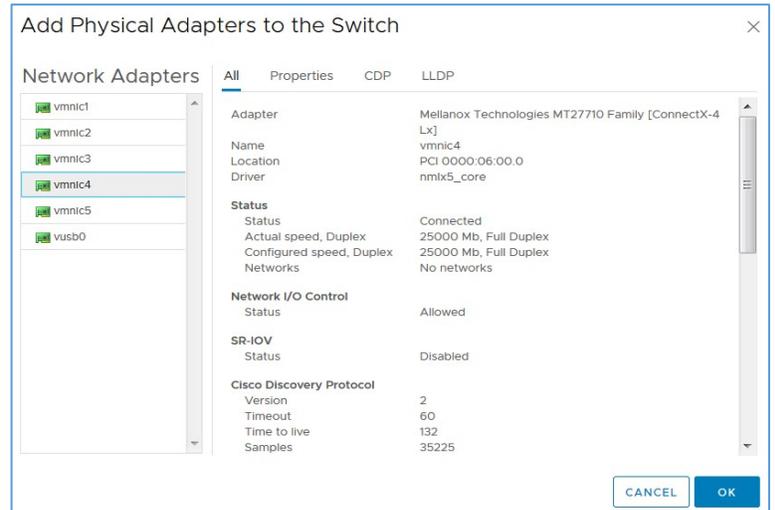
4. Select **New standard switch** to create a vSphere Standard switch and click **Next**.



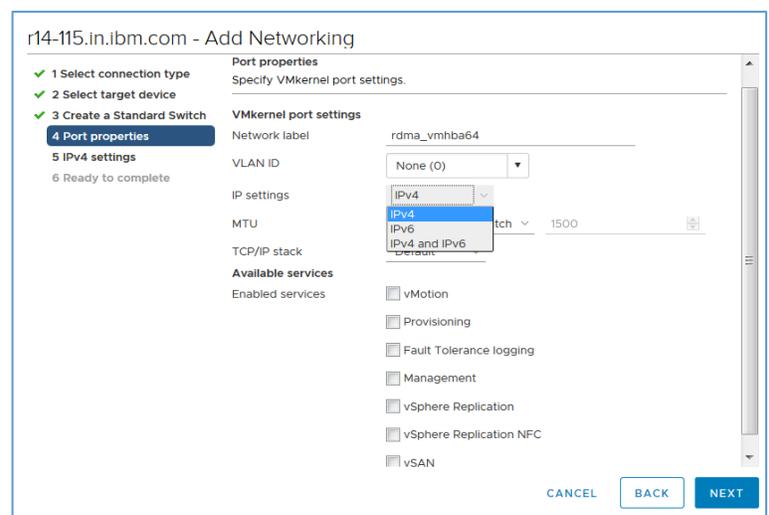
5. On the Create a Standard Switch page, click the **Add adapters** icon and select the network adapter (vmnic#) to use for iSCSI.



6. Make sure to use an appropriate RDMA-capable vmnic and click **Next**.



7. On the Port properties page, enter a network label. A network label is a friendly name that identifies the VMkernel adapter that you are creating, for example, **rdma\_vmhab64**. VLAN ID is optional. The selection depends upon the customer environment as well as network topology and Virtual network configurations. On the same page, specify the IP settings as well.



8. Specify the IP settings and click **Next**.

The screenshot shows the 'r14-115.in.ibm.com - Add Networking' wizard. The progress bar on the left indicates that step 5, 'IPv4 settings', is the current step. The main content area is titled 'IPv4 settings' and contains the following fields:

- Obtain IPv4 settings automatically:
- Use static IPv4 settings:
- IPv4 address: 192.168.65.199
- Subnet mask: 255.255.255.0
- Default gateway:  Override default gateway for this adapter  
9.113.56.1
- DNS server addresses: 9.0.146.50

Buttons at the bottom right include 'CANCEL', 'BACK', and 'NEXT'.

9. On the Ready to complete page, review the information you have specified and click **Finish**.

The screenshot shows the 'r14-115.in.ibm.com - Add Networking' wizard at the 'Ready to complete' step. The progress bar on the left indicates that step 6, 'Ready to complete', is the current step. The main content area is titled 'Ready to complete' and contains the following information:

Review your settings selections before finishing the wizard.

New port group	rdma_vmhabs64
Standard switch	vSwitch1
VLAN ID	None (0)
vMotion	Disabled
Provisioning	Disabled
Fault Tolerance logging	Disabled
Management	Disabled
vSphere Replication	Disabled
vSphere Replication NFC	Disabled
vSAN	Disabled

**NIC settings**

MTU	1500
TCP/IP stack	Default

**IPv4 settings**

IPv4 address	192.168.65.199 (static)
Subnet mask	255.255.255.0

Buttons at the bottom right include 'CANCEL', 'BACK', and 'FINISH'.

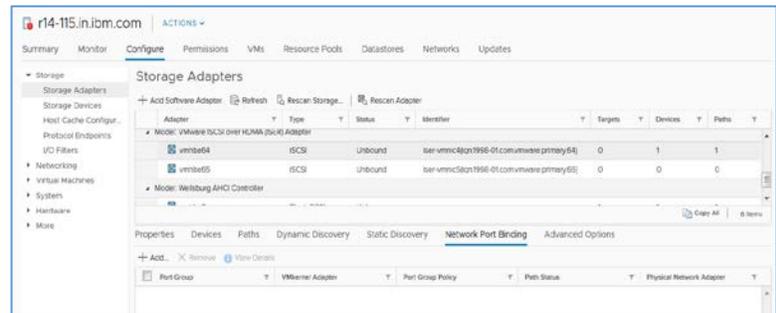
## Bind iSER and VMkernel adapters

Prerequisites to bind iSER and VMkernel adapters:

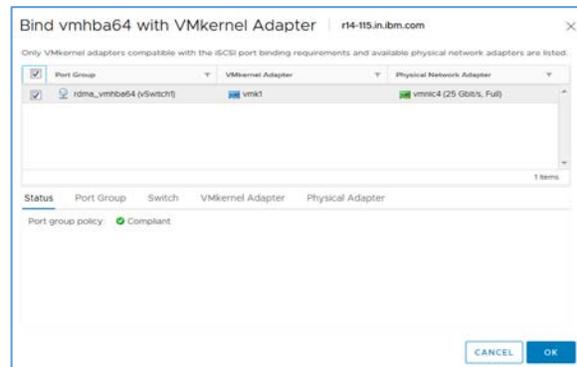
Create a virtual VMkernel adapter for each physical RDMA-capable network adapter on your host.

Procedure to bind iSER and VMkernel adapters:

1. Navigate to the host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters** and select the appropriate iSCSI adapter from the list.
4. Click the **Network Port Binding** tab and click **Add**.



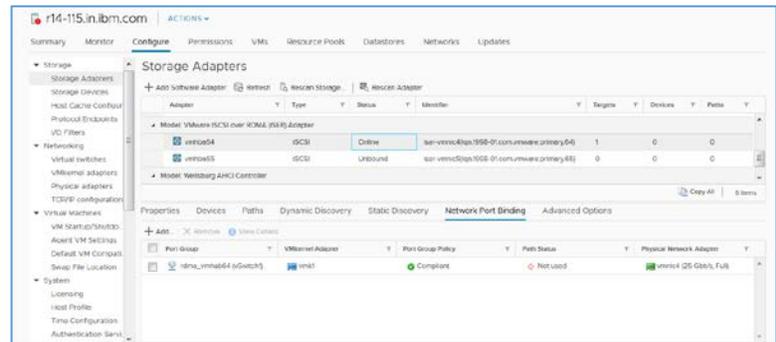
5. Select a VMkernel adapter to bind with the iSCSI adapter and click **OK**.



## Review networking details of the VMkernel adapter that is bound to iSER adapter

Perform the following steps to review the networking details of the VMkernel adapter that is bound to the iSER adapter:

1. Navigate to the host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters** and select the appropriate iSCSI adapter from the list.
4. On the **Network Port Binding** tab, select the corresponding port group check box and click **View Details**.



5. Review the VMkernel adapter information by switching between the available tabs.

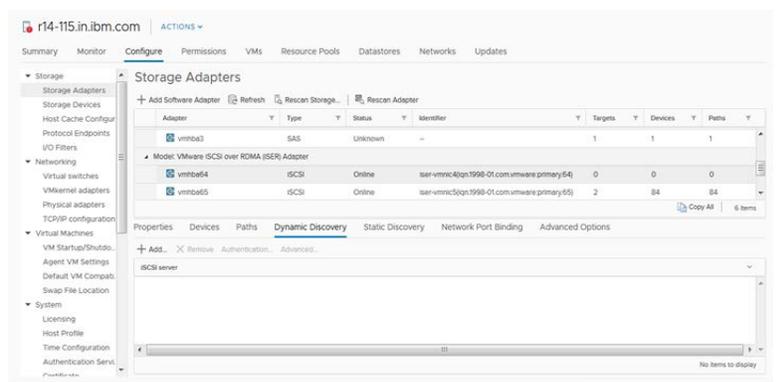
## Configure dynamic or static discovery for iSER

In dynamic discovery, each time the initiator contacts a specified iSCSI storage system, it sends the SendTargets request to the system. The iSCSI system responds by supplying a list of available targets to the initiator. In addition to the dynamic discovery method, you can use the static discovery method and manually enter information for the targets.

When you set up static or dynamic iSER discovery, you can only add new iSCSI targets. You cannot change any parameters of an existing target. To make changes, remove the existing target and add a new one.

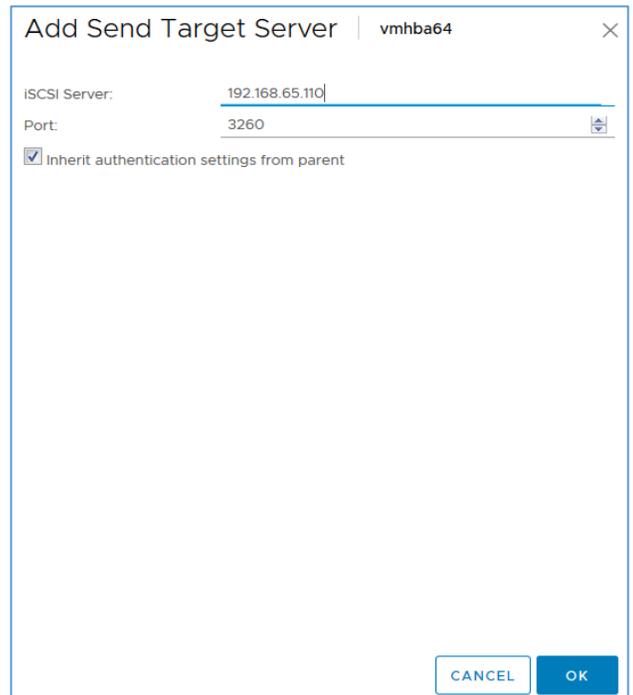
Procedure to configure dynamic and static discovery for iSER:

1. Navigate to the host.
2. Click the **Configure** tab.
3. Under **Storage**, click **Storage Adapters** and select the adapter (vmhba#) to configure.



4. Configure the dynamic discovery method.

- a. Click **Dynamic Discovery** and click **Add**.
- b. Enter the IP address or the Domain Name System (DNS) name of the storage system and click **OK**.



The screenshot shows a dialog box titled "Add Send Target Server" for the device "vmhba64". It contains the following fields and options:

- iSCSI Server:** 192.168.65.110
- Port:** 3260
- Inherit authentication settings from parent
- CANCEL** and **OK** buttons at the bottom right.

- c. Rescan the iSCSI adapter.  
After establishing the SendTargets session with the iSCSI system, your host populates the static discovery list with all newly discovered targets.

5. Configure the static discovery method.

- a. Click **Static Discovery** and click **Add**.
- b. Enter the target server's information and click **OK**.

The screenshot shows a dialog box titled "Add Static Target Server" with a close button (X) in the top right corner. The dialog contains the following fields and options:

- iSCSI Server: 192.168.65.110
- Port: 3260
- iSCSI Target Name: iqn.1986-03.com.ibm:2145.fab3|node1
- Inherit authentication settings from parent

At the bottom right, there are two buttons: "CANCEL" and "OK".

- c. Rescan the iSCSI adapter.

## Creating VMFS datastore and raw device mapping on ESX server

The volumes mapped from the Spectrum Virtualize FS9100 GUI in the step can be used as a VMware virtual machine file system (VMFS) datastore and raw device mapping (RDM).

VMFS datastores serve as repositories for virtual machines. You can set up VMFS datastores on the iSER storage device mapped to the host. Refer to the VMware documentation, [Create a VMFS Datastore](#) for more information on creating a VMFS datastore.

Using the RDM mechanism, a virtual machine can have direct access to a logical unit number (LUN) which is mapped from the Spectrum Virtualize FS9100 system. Refer to the VMware documentation, [Raw Device Mapping](#) for detailed description on RDM.

## Summary

The paper explained RDMA, the iSER host-attach feature for IBM Spectrum Virtualize products and its host attachment details with VMware ESX hosts. The key points addressed here are the most used configurations in the field by the customers.

## Get more information

To learn more about the product/service/function contact your IBM representative or IBM Business Partner or visit the IBM Knowledge Center websites.

Note:

- RHEL 6.5 operating system as a virtual machine is not supported.
- Increase the *SCSI time out* of the RDMs and virtual disk attached to the Linux guest OS to 60 sec (from the default 30 sec) to avoid I/O timeouts.

For details, visit <https://www-03.ibm.com/systems/support/storage/ssic/interoperability.wss#>

## About the author

**Rajsekhar Bharali** is a test specialist in IBM Systems and Development Labs Organization. He has more than 8 years' experience working with the storage platforms. You can reach Rajsekhar at [rabharal@in.ibm.com](mailto:rabharal@in.ibm.com)

**Rohan Salvi** is a development engineer in IBM Systems and Development Labs Organization. He has more than 4 years' experience working with the storage platforms. You can reach Rohan at [rosalvi1@in.ibm.com](mailto:rosalvi1@in.ibm.com)

**Sarvesh S. Patel** is a test architect in IBM Systems and Development Labs Organization. He has more than 10 years' experience working with the storage platforms. He is also a Master Inventor and a member of the TEC-I chapter. You can reach Sarvesh at [sarvepat@in.ibm.com](mailto:sarvepat@in.ibm.com)



---

© Copyright IBM Corporation 2019  
IBM Systems  
3039 Cornwallis Road  
RTP, NC 27709

Produced in the United States of America

IBM, the IBM logo and [ibm.com](http://ibm.com) are trademarks or registered trademarks of the Internal Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked items are marked on their first occurrence in the information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at [ibm.com/legal/copytrade.shtml](http://ibm.com/legal/copytrade.shtml)

VMware, the VMware logo, VMware Cloud Foundation, VMware Cloud Foundation Service, VMware vCenter Server, and VMware vSphere are registered trademarks or trademarks of VMware, Inc. or its subsidiaries in the United States and/or other jurisdictions.

Other product, company or service names may be trademarks or service marks of others.

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.



Please recycle