QSAN
Official Document

RDMA (Remote Direct Memory Access)

White Paper

May 2026

ANNOUNCEMENT

Copyright

© Copyright 2026 QSAN Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from QSAN Technology, Inc.

QSAN believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Trademarks

- QSAN, the QSAN logo, QSAN.com, XCubeFAS, XCubeSAN, XCubeNXT, XCubeNAS, XCubeDAS, XEVO, SANOS, and QSM are trademarks or registered trademarks of QSAN Technology, Inc.
- Microsoft, Windows, Windows Server, and Hyper-V are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.
- Linux is a trademark of Linus Torvalds in the United States and/or other countries.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Mac and OS X are trademarks of Apple Inc., registered in the U.S. and other countries.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
- VMware, ESXi, and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other countries.
- Citrix and Xen are registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries.
- Other trademarks and trade names used in this document to refer to either the entities claiming the marks and names or their products are the property of their respective owners.

TABLE OF CONTENTS

Announcement i
Notices iv
Preface vi
Technical Support vi
Information, Tip, and Caution vi
1. Introduction 8
1.1. What is RDMA 8
1.2. Key Feature of RDMA 9
2. Theory of Operations 10
2.1. Mechanism of RDMA 10
3. RDMA Configuration 12
4. Test Results 13
5. Conclusion 15
6. Appendix 16
6.1. Apply To 17
6.2. Reference 17

©2026 QSAN Technology, Inc. All rights reserved.

FIGURES

Figure 1-1 RDMA Data Transfer Diagram 8
Figure 2-1 Data I/O Task under RDMA 10

Announcement
©2026 QSAN Technology, Inc. All rights reserved.
www.qsan.com

TABLES

Table 4-1 XN4226D iSCSI Protocol Test Matrix.....14
Table 4-2 XN4226D NVMe-oF Protocol Test Matrix.....15

Announcement
©2026 QSAN Technology, Inc. All rights reserved.
www.qsan.com
Official Document

NOTICES

Information contained in this document has been reviewed for accuracy. But it could include typographical errors or technical inaccuracies. Changes are made to the document periodically. These changes will be incorporated in new editions of the publication. QSAN may make improvements or changes in the products. All features, functionality, and product specifications are subject to change without prior notice or obligation. All statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

©2026 QSAN Technology, Inc. All rights reserved.

PREFACE

Technical Support

Do you have any questions or need help trouble-shooting a problem? Please contact QSAN Support, we will reply to you as soon as possible.

- Via the Web: https://www.qsan.com/technical_support
- Via Telephone: +886-2-77206355
- (Service hours: 09:30 - 18:00, Monday - Friday, UTC+8)
- Via Skype Chat, Skype ID: qsan.support
- (Service hours: 09:30 - 02:00, Monday - Friday, UTC+8, Summer time: 09:30 - 01:00)
- Via Email: support@qsan.com

Information, Tip, and Caution

This document uses the following symbols to draw attention to important safety and operational information.



INFORMATION

INFORMATION provides useful knowledge, definition, or terminology for reference.



TIP

TIP provides helpful suggestions for performing tasks more effectively.

©2026 QSAN Technology, Inc. All rights reserved.



CAUTION

CAUTION indicates that failure to take a specified action could result in damage to the system.

©2026 QSAN Technology, Inc. All rights reserved.

1. INTRODUCTION

This document provides technical guidance for RDMA, which enhances the overall IT environment by offloading data transfers to RDMA, reducing CPU overhead, and ensuring optimal data paths for storage performance.

1.1. What is RDMA

RDMA (Remote Direct Memory Access) is a networking technology that allows one system to access another system's memory directly over a network. By offloading data transfers from the CPU and bypassing the OS, RDMA significantly reduces latency and system overhead.

Offloading data movement to RDMA-capable NICs also frees up CPU resources, allowing more efficient use of hardware and improving scalability across distributed environments. RDMA supports zero-copy transfers, minimizing memory usage and system overhead. As a result, it enhances performance while reducing power consumption, making it a key enabler for modern storage protocols like iSCSI, NFS, and NVMe-oF. In essence, RDMA transforms data communication into a faster, leaner, and more scalable process—well-suited for next-generation IT infrastructure.

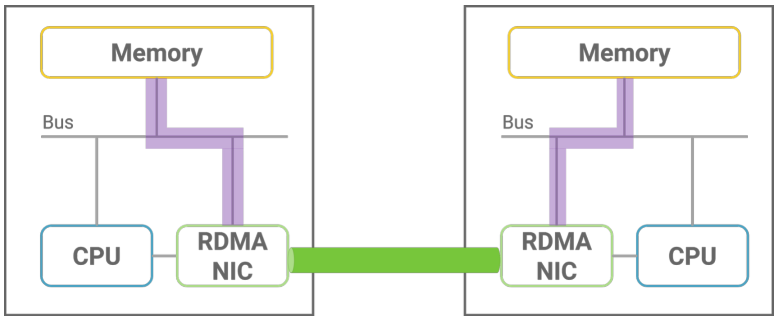


Figure 1-1 RDMA Data Transfer Diagram

©2026 QSAN Technology, Inc. All rights reserved.

1.2. Key Features of RDMA

RDMA enables high-speed, low-latency data transfers by allowing systems to access each other's memory directly without CPU involvement or kernel intervention. This architecture eliminates traditional bottlenecks by supporting zero-copy communication, avoiding repeated memory copies between user and kernel space. Applications can prepare and transmit data directly from user space, achieving faster response times and freeing up valuable CPU resources, especially on the remote side.



INFORMATION

When adopting an RDMA environment, please check all the hardware, including the host, NIC, and ensure that your storage version supports RDMA architecture. In addition, the storage system must be equipped with at least 16 GB of memory.

©2026 QSAN Technology, Inc. All rights reserved.

2. THEORY OF OPERATIONS

This chapter will introduce the theory of RDMA in detail.

2.1. Mechanism of RDMA

RDMA enables direct memory access between two systems across a network without involving the CPU or OS. Unlike traditional networking stacks that rely on TCP/IP, RDMA provides a low-latency, high-throughput transport mechanism that operates beneath modern storage protocols. Technologies such as NVMe-oF, iSCSI, and even NFS over RDMA leverage RDMA to achieve faster, more efficient data movement. In this architecture, RDMA is not the storage protocol itself. It is the underlying transport layer that accelerates how those protocols deliver data.

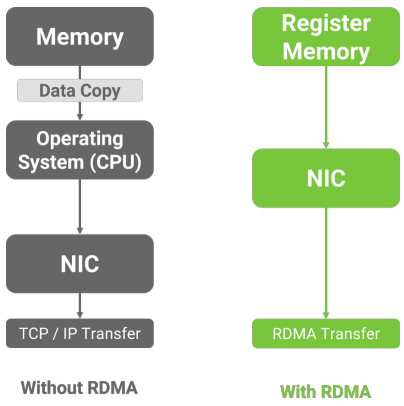


Figure 2-1 Data I/O Task under RDMA

To initiate an RDMA operation, applications first register memory regions with the local NIC. These pinned regions are made accessible to the RDMA engine, enabling direct transfers to or from remote memory. Communication between endpoints is managed through queue pairs that handle send/receive or read/write commands, using pre-established memory keys to ensure access control and address mapping. RDMA optimizes the data transfer path between systems. In addition, features like message-based transactions and zero-copy transfers minimize overhead and enable efficient, hardware-driven communication.

©2026 QSAN Technology, Inc. All rights reserved.

2.1.1. Message-Based Transactions

RDMA communication is built on a message-based model, where each operation—such as a send, receive, read, or write. It is treated as a discrete, self-contained transaction. Unlike stream-based protocols like TCP, which transmit a continuous flow of bytes without preserving message boundaries, RDMA maintains the integrity of each message from sender to receiver.

This approach simplifies application design and eliminates the need for manual parsing, reassembly, or byte tracking. The RDMA engine delivers each message directly to its corresponding registered memory buffer, ensuring accurate and efficient delivery. This is particularly beneficial in environments where structured data exchange is critical, such as storage protocols or tightly-coupled distributed systems.

Message-based transactions also improve determinism, as each operation maps cleanly to a memory region and a completion event. Combined with RDMA's use of queue pairs and direct memory access, this design enables faster, more predictable communication with minimal software overhead—ideal for latency-sensitive applications in storage, AI, and virtualization.

2.1.2. Zero-Copy Transfers

One of RDMA's core advantages is its ability to perform zero-copy transfers—moving data directly between the memory of two systems without passing through intermediate buffers or involving the CPU. In traditional network communication, data must be copied multiple times: from user space to kernel space, then from kernel space to NIC buffers, and again on the receiving side. Each copy introduces latency, consumes memory bandwidth, and increases CPU load.

RDMA eliminates these steps by allowing the NIC to access application-registered memory regions directly. Once memory is pinned and registered, the RDMA engine performs the data transfer directly between these regions over the network. This approach not only reduces latency and CPU utilization but also minimizes cache pollution and memory bandwidth consumption.

©2026 QSAN Technology, Inc. All rights reserved.

3. RDMA CONFIGURATION

The following are tutorials for configuring RDMA environments. Follow the steps and you can learn them quickly.

For iSCSI Environment

- **How to Setup RDMA over iSCSI**

For NFS Environment

- **How to Setup RDMA over NFS**

©2026 QSAN Technology, Inc. All rights reserved.

4. TEST RESULTS - iSCSI

This test verifies the dramatic performance improvement offered by RDMA. We test the RDMA with iSCSI and NVMe Protocol. Please check the following test result for a clear vision of the benefit from RDMA.

Test Equipment and Configurations

- Server
  - Model: GIGABYTE R181-342
    Memory: 64 GB
    iSCSI HBA: Mellanox ConnectX-4 Lx
    OS: Ubuntu 22.02Equipment
- Storage
  - Model: XN4226D-8C
    Memory: 32 GB per controller
    Host Connection: iSCSI: HCQ25GS3-40
    Firmware: QSM 4.1.6
    SSD: QSAN SD43T840-00 3.84 TB * 26

©2026 QSAN Technology, Inc. All rights reserved.

4.1.1. iSCSI Test Configuration

The first test item is the iSCSI protocol over RDMA data transfer. The iSCSI over RDMA test showed an average 10% performance gain, with a notable 60% improvement in 128K sequential writes.

Table 4-1 XN4226D iSCSI Protocol Test Matrix

Test Item	Without RDMA	With RDMA	Improvement
4K Random Read (IOPS)	551,722	591,306	7%
4K Random Write (IOPS)	401,957	445,308	11%
128K Sequential Read (MB/s)	6,481	7,764	20%
128K Sequential Write (MB/s)	4,821	7,721	60%

©2026 QSAN Technology, Inc. All rights reserved.

5. TEST RESULTS – NVME-OF

5.1.1. NVMe-oF Test Configuration

The next test evaluates the NVMe-oF over RDMA data transfer method for the all-flash model. The most significant improvement was observed in 128K sequential write throughput, which increased by 56% with RDMA enabled, demonstrating its advantage in accelerating large block sequential transfers.

Table 5-1 XN4226D NVMe-oF Protocol Test Matrix

Test Item	Without RDMA	With RDMA	Improvement
4K Random Read (IOPS)	967,357	916,059	7%
4K Random Write (IOPS)	521,900	485,000	11%
128K Sequential Read (MB/s)	14,094	12,626	12%
128K Sequential Write (MB/s)	9,049	5,817	56%

©2026 QSAN Technology, Inc. All rights reserved.

6. CONCLUSION

As the test result shows, adopting RDMA can greatly improve system performance by making data transfers faster and more efficient. It allows two systems to move data directly between their memory without going through the CPU or operating system, which reduces delay and frees up system resources. RDMA also avoids unnecessary copying and handles data in clear, structured messages. By integrating RDMA into QSAN's storage solutions, enterprises can accelerate critical workloads, reduce latency, and ensure infrastructure readiness for AI, virtualization, and large-scale data demands. When used in the right setup, RDMA helps systems run faster, smoother, and more reliably.

In particular, RDMA proves especially valuable in workloads such as real-time financial trading, AI/ML training and inference, virtualization clusters, and high-performance databases, all of which benefit significantly from its ability to minimize CPU involvement and accelerate communication between nodes.

©2026 QSAN Technology, Inc. All rights reserved.

7. APPENDIX

7.1. Apply To

- QSM firmware 4.1.5 and later
- XEVO firmware 3.0.0 and later

7.2. Reference

Document

- QSM 4 Software Manual
- XEVO 3 Software Manual

©2026 QSAN Technology, Inc. All rights reserved.