QSAN
Official Document

QScale (Scale-out)

White Paper

September 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, and QSAN.com 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 v
Preface vi
Technical Support vi
Information, Tip, and Caution vi
1. Introduction 1
2. Theory of Operation 2
2.1. Cluster Component Overview 2
2.2. Data Distribution and Access 3
2.3. QScale File Access Architecture 3
2.4. Cluster Component Status Matrix 5
2.5. QScale Formation Requirements 6
3. Test Results 8
3.1. Test Environment and Results 8
3.2. Explanation of Linear Scaling 10
4. Conclusion 11
5. Appendix 12
5.1. Apply To 12
5.2. Reference 12

©2026 QSAN Technology, Inc. All rights reserved.

FIGURES

Figure 2-1 Cluster Data Access Topology 3
Figure 2-2 QScale File Access Architecture 4

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

TABLES

Table 2-1 Node Status Table 5
Table 2-2 Cluster Volume Status Table 5
Table 2-3 Volume Status Table 6
Table 3-1 Test Result 9

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

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, Teams / 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

The rapid growth of unstructured data is placing increasing pressure on traditional storage infrastructures. As organizations expand their data-intensive workloads, a single storage system may eventually become constrained by capacity limits, performance bottlenecks, and operational complexity.

To address these challenges, QSAN introduces QScale, a scale-out storage architecture that aggregates multiple storage systems into a unified cluster environment. By combining resources across multiple nodes, QScale enables organizations to expand capacity and performance incrementally while simplifying management and maintaining service continuity.



INFORMATION

To utilize the QScale feature, please check the latest XInsight environment and the license active status before configuring.

QSAN QScale is designed to aggregate multiple storage systems into a unified service architecture. Through scale-out service technology and cluster-level management, enterprises can expand file storage capacity and distribute workloads across multiple nodes.

This solution aims to provide the following core values:

- Linear capacity expansion through the addition of storage nodes.
- Improved aggregation performance by distributing workloads across multiple systems.
- Hot and cold data placement at the file system level across nodes.
- Node-level redundancy using designated standby or mirrored resources.
- Cluster-aware fault handling for node or extent anomalies.
- Centralized QScale management through a single window.

This document introduces the advanced architecture, operating model, deployment considerations, use cases, and validation framework of QSAN QScale archive storage.

©2026 QSAN Technology, Inc. All rights reserved.

2. THEORY OF OPERATION

QScale is a horizontal scale-out storage architecture that aggregates QSAN storage systems into a unified cluster environment. Within the cluster, storage resources from different nodes can be combined to create cluster volume, enabling capacity and performance to scale beyond the limitations of a single system.

The architecture is designed to support growing datasets and concurrent workloads while providing consistent access experience through centralized cluster management.

2.1. Cluster Component Overview

QScale consists of the following key components for running a horizontally scaled architecture, including:

Cluster

Defines the collection of participating storage nodes managed as a single logical environment.

Primary Node

The primary node serves as the control point responsible for cluster-wide coordination, configuration consistency, and management operations.

Node

A storage system participating in the cluster and contributing capacity and resources to the overall environment.

Cluster Volume

A logical storage layer that aggregates resources from multiple nodes and presents them as a unified storage space.

©2026 QSAN Technology, Inc. All rights reserved.

2.2. Data Distribution and Access

QScale provides a unified access layer across all participating nodes in the cluster. Each node maintains its own local storage resources while simultaneously integrating storage resources from other nodes. This distributed architecture enables capacity aggregation across multiple systems while presenting a consistent storage environment to users and applications.

When data is written through a node, QScale prioritizes local resource utilization to maximize efficiency. As capacity utilization increases, additional data can be distributed across available resources within the cluster, allowing organizations to scale storage capacity without disrupting existing workflows.

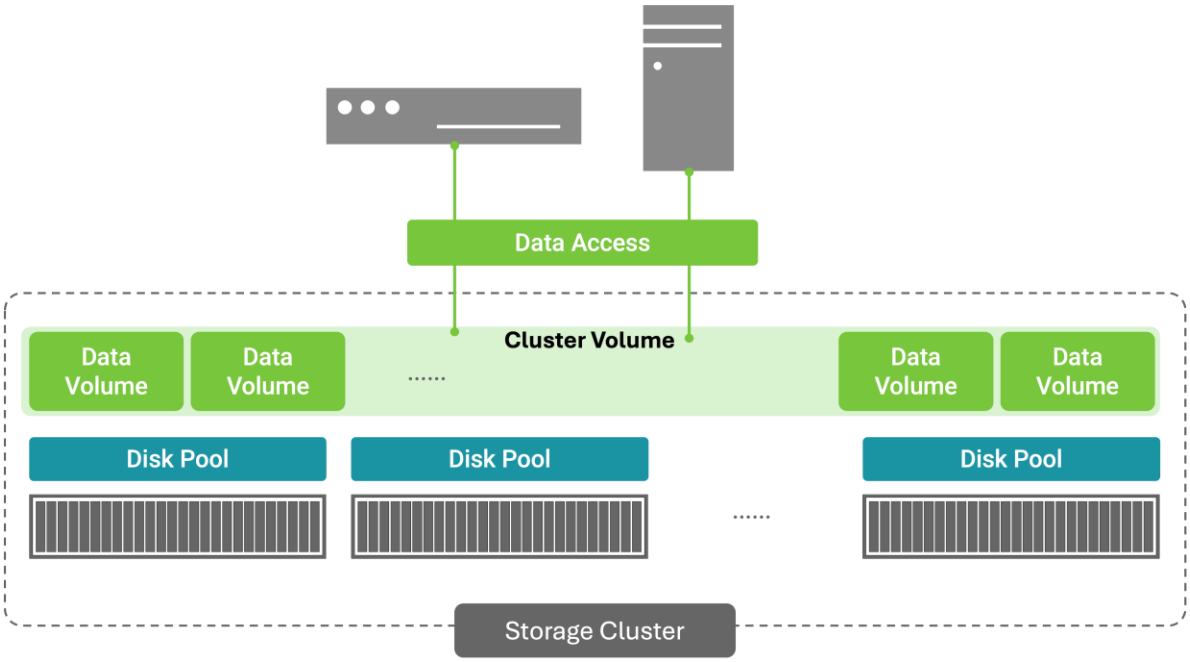


Figure 2-1 Cluster Data Access Topology

2.3. QScale File Access Architecture

Taking QScale for File as an example, the file access architecture is designed to provide a unified file access layer across multiple NAS nodes, allowing user clients to access the clustered file disk area through any participating node. As shown in the architecture diagram, each node provides its own data service and exposes a dedicated cluster IP address. Clients can access the clustered

©2026 QSAN Technology, Inc. All rights reserved.

file disk area through cluster IP address 1 on node 1, cluster IP address 2 on node 2, or cluster IP address 3 on node 3.

Internally, each node maintains a local file system volume and mounts the file systems of other nodes via NFS. These local and remote file systems are then integrated through a federated mount layer and presented as a single clustered file disk area.

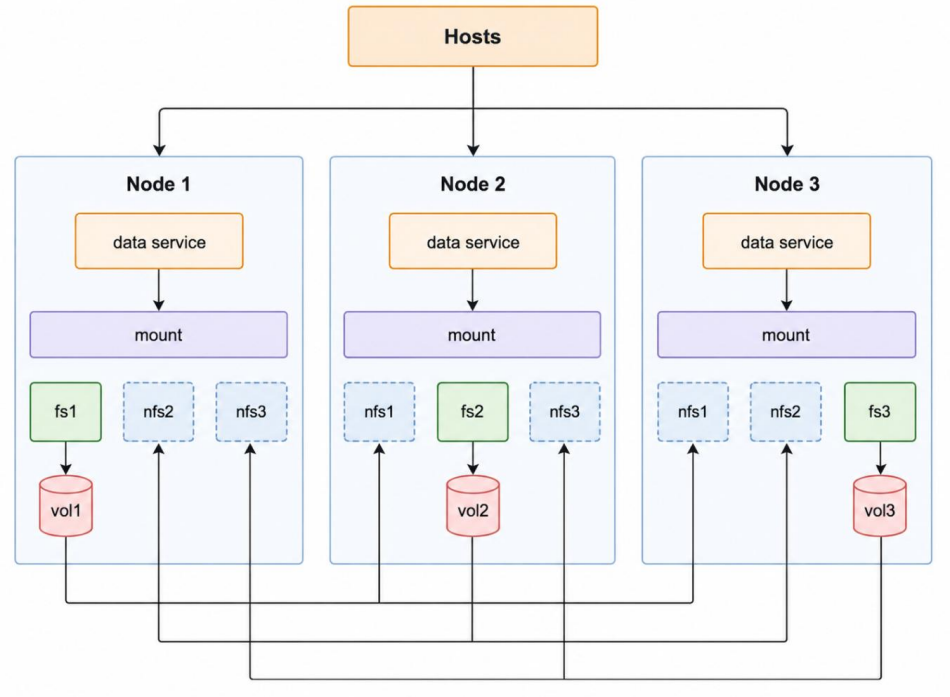


Figure 2-2 QScale File Access Architecture

When a client writes data to a shared folder on a clustered file system, the system prioritizes writing the data to the local file system of the node the client is connected to. For example, if a client accesses the shared folder via cluster IP 1 on node 1, the new file will first be written to node 1's local file system. Once node 1 reaches its available capacity limit, the federated mount layer can redirect new write operations to the file systems provided by node 2 or node 3 via NFS. This mechanism allows the clustered file system to aggregate capacity from multiple nodes while maintaining a simple and consistent access path for clients.

Shared Management Practices

Once the cluster is established, the administrator should perform all subsequent operations through XInsight UI. Avoid performing operations through the native interface of a single node, as these operations may bypass the cluster synchronization logic. All cluster-aware operations

©2026 QSAN Technology, Inc. All rights reserved.

must be initiated from XInsight UI to ensure that the configuration and state of all nodes within the cluster are synchronized and consistent.

2.4. Cluster Component Status Matrix

QScale continuously monitors cluster health and resource availability to help administrators identify and respond to operational events.

Node Status

Table 2-1 Node Status Table

Status	Description
OK	Node operating normally
Disconnected	Node unreachable through the expected network path

Cluster Volume Status

Table 2-2 Cluster Volume Status Table

Status	Description
OK	Cluster operating normally
Not Ready	Administrative action required
Degraded	One or more resources unavailable but service remains operational

©2026 QSAN Technology, Inc. All rights reserved.

Volume Status

Table 2-3 Volume Status Table

Status	Description
OK	Volume operating normally
Disconnected	Volume or network path unavailable
Unavailable	Volume exists but cannot be accessed

2.5. QScale Formation Requirements

Before a cluster is established, all participating nodes must meet a set of defined prerequisites and consistency checks.

Required Software Version

- QSM firmware 4.2.0 and later
- XEVO firmware 3.2.0 and later
- XInsight 1.6.0 and later
- QScale subscription license for each node

Node Requirements

- A maximum of 8 nodes, including the primary node, are supported in the cluster creation workflow.
- Nodes already belonging to other clusters cannot be added.
- A file volume can only belong to one cluster file volume.
- Windows AD (Active Directory) domain configuration must be completed before joining QScale.
- WORM (Write Once Read Many) volume is not supported.

©2026 QSAN Technology, Inc. All rights reserved.

Cluster Initialization Requirements

- The system will designate one node as the primary node.
- Administrator credentials are required during cluster establishment to perform authentication and security key exchange.
- All candidate nodes with dual controllers must be accessible via the management IP.

Identity and Configuration Consistency

Before successfully establishing a cluster, all nodes must maintain the following consistent settings:

- Domain status
- Date and time alignment
- Time zone alignment
- 2FA (Two-Factor Authentication) for cluster participation is disabled

Parameter Check

- The maximum number of users: 2,048
- The maximum number of groups: 1,024
- The maximum number of file volumes: 8,192
- The maximum number of shares: 8,192

©2026 QSAN Technology, Inc. All rights reserved.

3. TEST RESULTS

3.1. Test Environment and Results

We design experiments and test performance. One group consisted of a cluster of two nodes, while the other control group had only one system.

Test Environment

- Server
  - Model: ASUS Server
- Storage
  - Model: 2 x XN rackmount storage system
    Firmware: QSM 4.2.0
  - SSD: 8 x SSD

Test Topology

- QScale Configuration
  - XN (SSD x4) + XN (SSD x4)
- Control Group
  - XN (SSD x8)

©2026 QSAN Technology, Inc. All rights reserved.

Test Result

Table 3-1 Test Result

METRIC	QSCALE GROUP XN + XN	CONTROL GROUP XN
Throughput (Read)	4,400 MB/s	2,650 MB/s
Throughput (Write)	2,900 MB/s	2,200 MB/s
IOPS (Read)	28K IOPS	16K IOPS
IOPS (Write)	8K IOPS	5K IOPS
4 x AJA (Read)	2,700 MB/s	1,600 MB/s
4 x AJA (Write)	1,700 MB/s	1,200 MB/s

Summary

Benchmark results demonstrate that QSAN QScale significantly improves performance when scaling from a single XN node to a two-node scale-out configuration. In this validation, the two-node configuration achieved higher performance across all test categories, with performance improvements ranging from 31.8% to 75.0% depending on the workload type.

The performance improvements were most significant for read workloads, including a 75.0% increase in read IOPS, a 68.8% increase in AJA read performance, and a 66.0% increase in read throughput. Write workloads also showed substantial performance improvements, including a 60.0% increase in write IOPS, a 41.7% increase in AJA write performance, and a 31.8% increase in write throughput.

These results validate the architectural advantages of QScale: by aggregating multiple XN nodes into a unified scale-out file environment, enterprises can significantly improve availability, performance, and capacity. This design is particularly important for file-intensive environments, such as media production, collaborative content workflows, monitoring data repositories, and general enterprise file services.

©2026 QSAN Technology, Inc. All rights reserved.

Performance scaling is affected by workload patterns, client distribution, network topology, disk configuration, and protocol behavior. Therefore, test results should be viewed as practical validation of horizontal scaling efficiency, rather than fixed performance guarantees. Even so, the test results confirm that QScale provides a scalable foundation for organizations seeking incremental growth without redesigning their file access architecture.

3.2. Explanation of Linear Scaling

The design goal explicitly points to linear growth in performance and available capacity. In practice, this means that the working storage domain can be expanded by adding nodes and file volumes without forcing a complete upgrade to build a larger system.

However, linear scaling should be understood as an architectural goal that depends on:

- Workload locality
- Network path quality
- Cluster file volume composition
- Protocol behavior and application concurrency patterns

From a performance perspective, the optimal deployment model is to distribute client access across multiple nodes. In the best-case scenario, three clients access the same cluster file disk area via cluster IP 1, cluster IP 2, and cluster IP 3, respectively. Since each client writes to the local file system of its connected node, the workload is distributed across three independent nodes. In this scenario, the overall performance is expected to approach linear scaling, reaching approximately three times the performance of a single node, provided that client workload, network bandwidth, disk configuration, and file system behavior do not become bottlenecks.

©2026 QSAN Technology, Inc. All rights reserved.

4. CONCLUSION

As data volumes continue to grow across modern IT environments, organizations require storage architectures that can scale beyond the limitations of a single system. QScale addresses this challenge by providing a unified cluster architecture that enables incremental expansion of capacity and resources while simplifying administration through centralized management.

By combining distributed storage resources, unified access, and cluster-aware operations, QScale helps organizations build scalable storage infrastructures that can adapt to evolving business requirements. From media production and surveillance archives to AI-driven data environments, QScale delivers a flexible and efficient approach to long-term storage growth.

©2026 QSAN Technology, Inc. All rights reserved.

5. APPENDIX

5.1. Apply To

- QSM firmware 4.2.0 and later
- XEVO firmware 3.2.0 and later
- XInsight 1.6.0 and later
- QScale subscription license for each node

5.2. Reference

Product Page

- QScale

Data Sheet

- QSM 4 Data Sheet
- XEVO 3 Data Sheet

Document

- QSM 4 Software Manual
- XEVO 3 Software Manual

©2026 QSAN Technology, Inc. All rights reserved.