QSAN
Official Document

QDP (Dynamic Pool)

White Paper

October 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. QDP Architecture 2
2.2. Data Protection and Rebuild 3
2.3. Comparison of Traditional RIAD and QDP 5
3. Deployment and Operations 8
3.1. Prerequisites 8
3.2. Capacity Planning 9
3.3. Operations 14
3.4. Best Practices 16
4. Conclusion 18
5. Appendix 19
5.1. Apply To 19
5.2. Reference 19

©2026 QSAN Technology, Inc. All rights reserved.

FIGURES

Figure 2-1 Stripe Distribution of Traditional RAID and QDP 3
Figure 2-2 Rebuild of Traditional RAID and QDP 4
Figure 3-1 Disk Group Capacity Composition and Pool Planning 11
Figure 3-2 Online Disk Group Expansion and Automatic Data Balancing 12

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

TABLES

Table 2-1 QDP Composition 2
Table 2-2 Key Differences between Traditional RAID and QDP 6
Table 2-3 Rebuild Time Comparison 7
Table 3-1 Architectural Limitations 8
Table 3-2 Functional Limitations 9
Table 3-3 QDP Stripe Width Parameters 9
Table 3-4 Online Expansion Behavior 13
Table 3-5 Migrating from Traditional RAID 16

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, 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

RAID technology has remained structurally unchanged since it was first proposed in 1988. Stripes are bound to a fixed set of disks, available capacity is determined by the smallest disk in the group, and after failure, the entire rebuild workload is funneled to a single hot spare. When the disk capacity is a few hundred gigabytes, these tradeoffs matter little.

With the capacity of a single disk now exceeding 20TB, three structural problems have emerged:

- Rebuild time becomes too long. All rebuild writes are absorbed by a spare disk, so rebuilds can run for tens of hours. A second failure during this window puts data at risk.
- Capacity is not fully utilized. When a cabinet holds drives of different capacities, traditional RAID must work on the smallest drive, with everything above that threshold being shelved.
- Hot spare disks have no role in normal operation. A dedicated standby device provides neither capacity nor I/O; it exists only for the moment of failure.

QSAN QDP (Dynamic Pool) solves these problems through declustered RAID architecture. Disk space is divided into multiple partitions, and each stripe dynamically selects a fixed number of partitions from the disk pool so that data and parity are evenly distributed on each disk in the group. Spare space is also distributed evenly across each disk, rather than concentrated on an idle hot spare. This directly brings the following benefits.

1. Disks of different capacities can be mixed freely.
2. The reconstruction work is shared by all healthy disks.
3. Disks that retain free space can also provide capacity and I/O during normal operation.

©2026 QSAN Technology, Inc. All rights reserved.

2. THEORY OF OPERATION

2.1. QDP Architecture

QDP (DynamicPool, formerly known as V-RAID) is QSAN's own implementation of declustered RAID. The essential difference from traditional RAID is that traditional striping is permanently bound to a set of disks, while declustered striping dynamically selects a fixed number of disks from a larger population.

2.1.1. Composition of a Stripe Set

QDP divides the space on each disk into uniformly sized chunks and then assembles these chunks into internal stripe sets. Each stripe set consists of two block types.

Table 2-1 QDP Composition

BLOCK TYPE	COUNT	PURPOSE
Data Block	D — specified by the user at creation	Stores actual data; determines the usable capacity ratio
Parity Block	P — 1 for QDP 5, 2 for QDP 6	Provides fault tolerance

Therefore, the strip width is D + P. No matter how many disks exceed the stripe width, there is spare space: QDP requires that at least one disk's spare space be reserved in advance, but that spare disk is not a separate disk - it is evenly distributed on each disk in the group. This is the most fundamental departure from the data protection concept of traditional RAID. In traditional RAID, the spare disk is an idle hot spare disk; in QDP, the spare disk exists on each working disk.

©2026 QSAN Technology, Inc. All rights reserved.

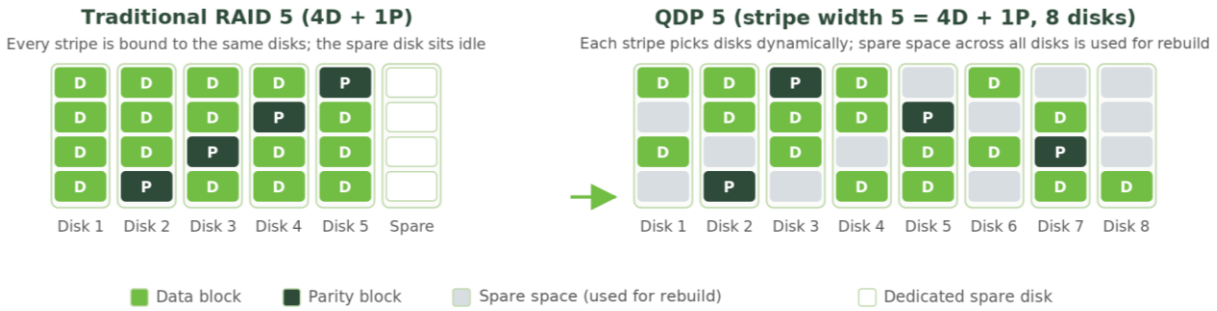


Figure 2-1 Stripe Distribution of Traditional RAID and QDP

As shown in the figure above, all four stripes of a traditional RAID 5 group are located on disks 1 through 5, while the sixth disk (the hot spare) remains completely idle during normal operation. QDP 5 provides the same level of protection while spreading four stripes across eight disks, each stripe still contains four data blocks and a parity block, but the disk groups vary from stripe to stripe. Unused blocks (including spare space) are distributed across all disks and used directly during rebuilding.

2.1.2. Architecture Compatibility

QDP only changes how partitions are composed. Everything above the VD (Virtual Disk) layer remains unchanged. This design decision has important practical implications. Existing volumes, LUNs, file services, and data services will continue to work without modifications to the QDP. Once a volume is created, the management experience is the same as traditional RAID.

TIP

QDP is selected at the storage pool level. A system can host both a QDP pool and a traditional RAID pool without interference.

2.2. Data Protection and Rebuild

©2026 QSAN Technology, Inc. All rights reserved.

Fault Tolerance

Fault tolerance is determined by the number of parity blocks, in line with traditional RAID. QDP 5 can tolerate one disk failure, and QDP 6 can tolerate two disk failures. This does not change with the number of disks or stripe width. Environments with higher protection level requirements should choose QDP 6.

Distributed Rebuild

When a disk in the group fails or is removed, the system automatically starts rebuilding. No action is required on the part of the operator, and there is no need to wait for the disk to be replaced. The reconstructed data is written directly to the spare space on the remaining healthy disk. The premise is that the disk group has reserved spare space for at least one disk in advance. If the number of disks is equal to the stripe width (no spares), the group will not be automatically rebuilt after degrading, and a hot spare must be assigned.

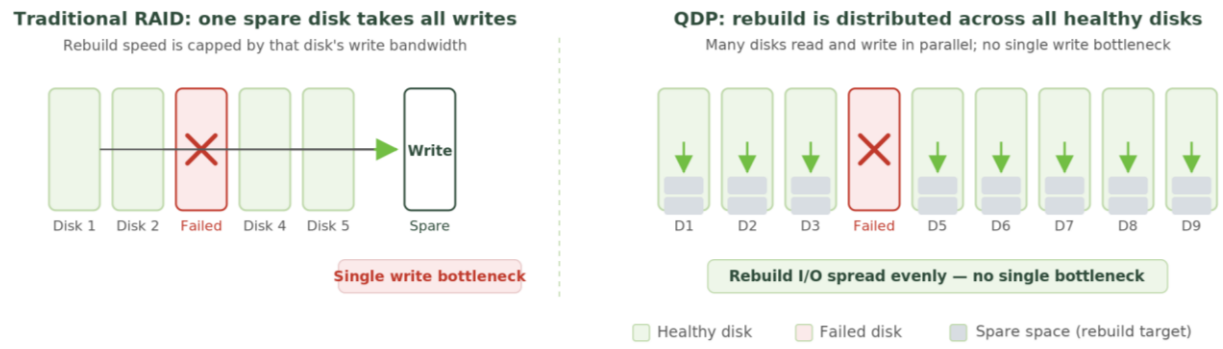


Figure 2-2 Rebuild of Traditional RAID and QDP

This is the advantage of declustered RAID. In traditional RAID, each rebuild write is written to a hot spare disk, so the rebuild speed is limited by the write bandwidth of that disk. In QDP, many healthy disks are read and written simultaneously, and there is no single point of contention.

2.2.1. Reserve Spare Space in Advance

Automatic rebuilding in QDP relies on reserved spare space within the group. The method is simple. When creating the disk group, the number of disks should be at least one more than the stripe width; one or two more is recommended.

For example, a disk group containing 16 disks with a stripe width of 14, for QDP 5, this equates to 13 data blocks plus 1 parity block. This effectively distributes the spare space of two disks

©2026 QSAN Technology, Inc. All rights reserved.

across the disk group, allowing rebuilding to begin immediately in the event of a failure, with sufficient margin to handle potential secondary failures during the rebuild process.



TIP

Customers migrating from a "traditional RAID 5 with a dedicated hot spare" setup simply need to add this spare disk as the spare disk to the QDP disk group. This way, the spare disk's capacity is not idle but distributed as spare space across all disks; it is used for I/O during normal operation and provides protection during rebuilding.



CAUTION

If the data block count is set to the maximum value in the creation wizard, making the disk count equal to the stripe width, the group will have no remaining space and will not be automatically rebuilt after a disk failure.

2.3. Comparison of Traditional RIAD and QDP

The following table is the key differences from traditional RIAD and QDP.

©2026 QSAN Technology, Inc. All rights reserved.

Table 2-2 Key Differences between Traditional RAID and QDP

ASPECT	TRADITIONAL RAID	QDP
Stripe Placement	Bound to a fixed set of disks	Each stripe selects disks dynamically, spread over all disks
Disk Capacity	Limited by the smallest disk	Full capacity of every disk is used; mixing is allowed
Rebuild Target	Writes to a single hot spare disk	Writes distributed to the spare space on all healthy disks
Spare	A whole dedicated hot spare disk must be reserved and sits idle	At least one disk's worth of spare must be reserved in advance; it is spread over all member disks, which keep serving I/O
Capacity Expansion	Difficult to add disks incrementally	Disks can be added online at any time
Usage Balancing	Not supported	Data is automatically redistributed within the group after disks are added
EE Spare / Sub-groups	Configured manually	No separate setting: spare = disk count − stripe width; sub-groups are handled automatically

2.3.1. Rebuild Time Comparison

The following table shows the rebuild time after a single disk failure in an approximately 1 PB all-flash NVMe SSD configuration. Because the number of disks and capacity per disk differ between configurations, rebuild time per GB (rebuild time / capacity of a single disk) is given as a basis for comparison; lower is faster.

©2026 QSAN Technology, Inc. All rights reserved.

Test Conditions

The background rebuild rate is the factory default. Rebuilding is triggered by removing a single member disk; there is no front-end I/O operation during rebuilding. Traditional RAID 5 and RAID 5EE use a single 1 PB thick volume, QDP 5 disk groups reserve a disk-sized spare space, and the pool utilization is approximately 82%.

Table 2-3 Rebuild Time Comparison

ARCHITECTURE	DISK CONFIGURATION	REBUILD TIME	REBUILD TIME PER GB
Traditional RAID 5	21 × 61.44 TB	77.7 hours	4.9 seconds
RAID 5EE	21 × 61.44 TB	57.8 hours	3.6 seconds
QDP 5	16 × 57 TB	48.2 hours	3.3 seconds

Test Summary

QDP 5 rebuilds approximately 33% faster than traditional RAID 5 on a per GB basis. RAID 5EE, QSAN's existing distributed standby RAID, also lets each member disk share rebuild writes, so the gap to QDP 5 is even smaller (~10%). What differentiates QDP from RAID EE is the ability to mix disk capacities and expand online one disk at a time.

©2026 QSAN Technology, Inc. All rights reserved.

3. DEPLOYMENT AND OPERATIONS

3.1. Prerequisites

This feature has architectural limitations, as detailed below.

Table 3-1 Architectural Limitations

ITEM	CURRENT LIMITATION
RAID Levels	Only QDP 5 and QDP 6 are offered
Stripe Width	3 to 16; cannot be modified after creation
Data and Parity Disks	At most 16 disks per group participate in data and parity; additional disks provide free space only
Disks per Group	Up to 63
Disk Removal	Disk groups cannot be shrunk; disks cannot be removed once added
Disk Mixing	SSDs and HDDs must not be mixed; JBOF and expansion enclosure disks must not be mixed
Provisioning	Thin provisioning only; thick provisioning is not offered

And has functional limitations, as detailed below. Users must be aware of these rules before using.

©2026 QSAN Technology, Inc. All rights reserved.

Table 3-2 Functional Limitations

ITEM	CURRENT LIMITATION
Deduplication and Compression	Not supported; the option is hidden in the management interface
RAID Migration	Bidirectional migration between traditional RAID and QDP is supported at the underlying layer, but the graphical interface does not yet provide a workflow for it; the command line must be used. Interface support is planned for a later release
Monitoring	Use existing pool monitoring; no QDP specific monitoring or alert items are provided

3.2. Capacity Planning

The stripe width is a parameter set by the user when creating a QDP disk group. The management interface will prompt for the block count; the stripe width is equal to this value plus the number of parity blocks.

Table 3-3 QDP Stripe Width Parameters

RAID TYPE	PARITY BLOCKS P	STRIPE WIDTH D + P	DATA BLOCK COUNT D
QDP 5	1	3 ~ 16	2 ~ 15
QDP 6	2	4 ~ 16	2 ~ 14

©2026 QSAN Technology, Inc. All rights reserved.



TIP

The block count does not specify which physical disks store data. The QDP partitioning mechanism is independent of specific disks; this value determines the ratio of data to protection in the stripe. This value cannot be modified after the disk group is created, so it must be properly planned from the beginning.

3.2.1. Calculating Usable Capacity

QDP fully utilizes the total capacity of each disk, rather than using only the minimum capacity. The available capacity is calculated as follows.

Usable capacity = Raw capacity of the disks participating in data and parity × D / (D + P)

Taking five disks (2 TB, 2 TB, 4 TB, 4 TB, 8 TB, total original capacity 20TB) as an example.

CONFIGURATION	STRIPE WIDTH	CALCULATION	USABLE CAPACITY
QDP 5	5 (4D + 1P)	20 TB × 4 / 5	16 TB
QDP 6	5 (3D + 2P)	20 TB × 3 / 5	12 TB
Traditional RAID 5	5 (4D + 1P)	Smallest disk 2 TB × 4	8 TB

Using the same disk array, QDP 5 offers twice the usable capacity of traditional RAID 5. This difference stems from the fact that traditional RAID requires a minimum disk capacity; data exceeding 2TB on 4TB and 8TB disks becomes unusable, while QDP has no such limitation. This example is for illustrative purposes only; in real-world deployments, the number of disks must be at least one more than the stripe width to provide spare capacity.

©2026 QSAN Technology, Inc. All rights reserved.

3.2.2. Disk Group Size and Pool Planning

A single disk group can hold a maximum of 63 disks, but only a maximum of 16 of them are used for data storage and parity checking. Therefore, placing a large number of disks in a disk group does not proportionally increase the available capacity: the extra disks actually provide usable space.

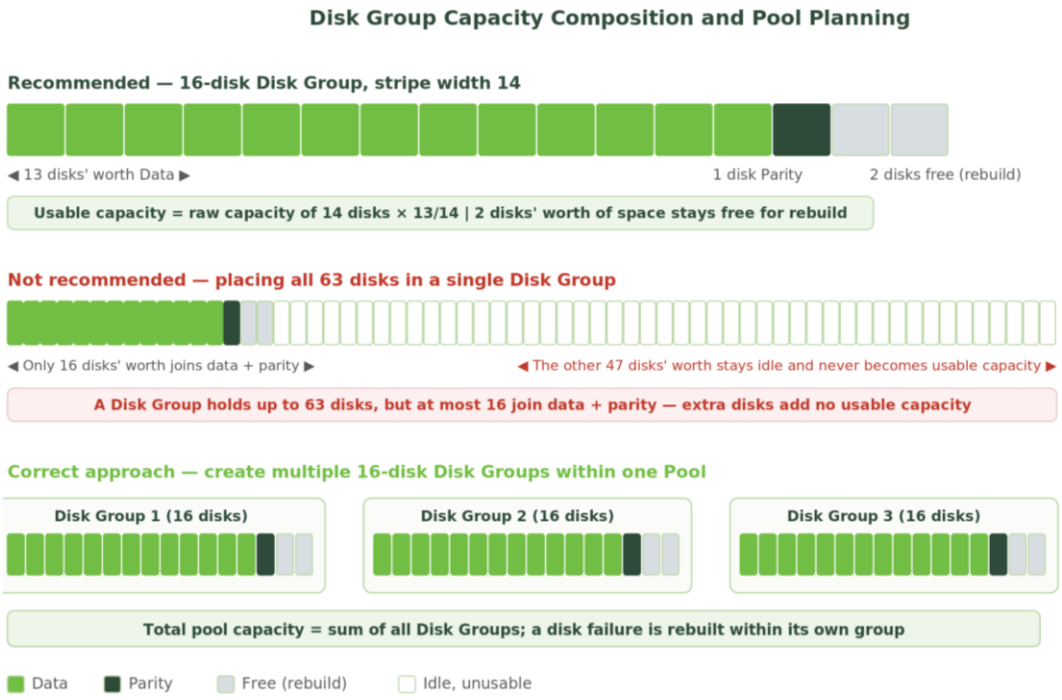


Figure 3-1 Disk Group Capacity Composition and Pool Planning

When a storage pool needs more available capacity, the correct approach is to create multiple disk groups within the same pool, each containing 16 disks. The total capacity of the storage pool is the sum of the capacities of all disk groups, and rebuilding after a disk failure is limited to the affected disk group.

3.2.3. Recommended Configuration

The following are recommended configurations for your reference.

- Disks per group: 16
- Stripe width: 14 (QDP 5: 13D + 1P; QDP 6: 12D + 2P)
- Spare space: The equivalent of two disks (at least one), spread across the group for rebuild

©2026 QSAN Technology, Inc. All rights reserved.

- When capacity runs short: Add a further disk group to the same pool rather than adding disks to an existing group

3.2.4. Online Expansion

Disks can be added to the QDP disk group at any time. There is no need to rebuild the entire disk pool or interrupt the service.

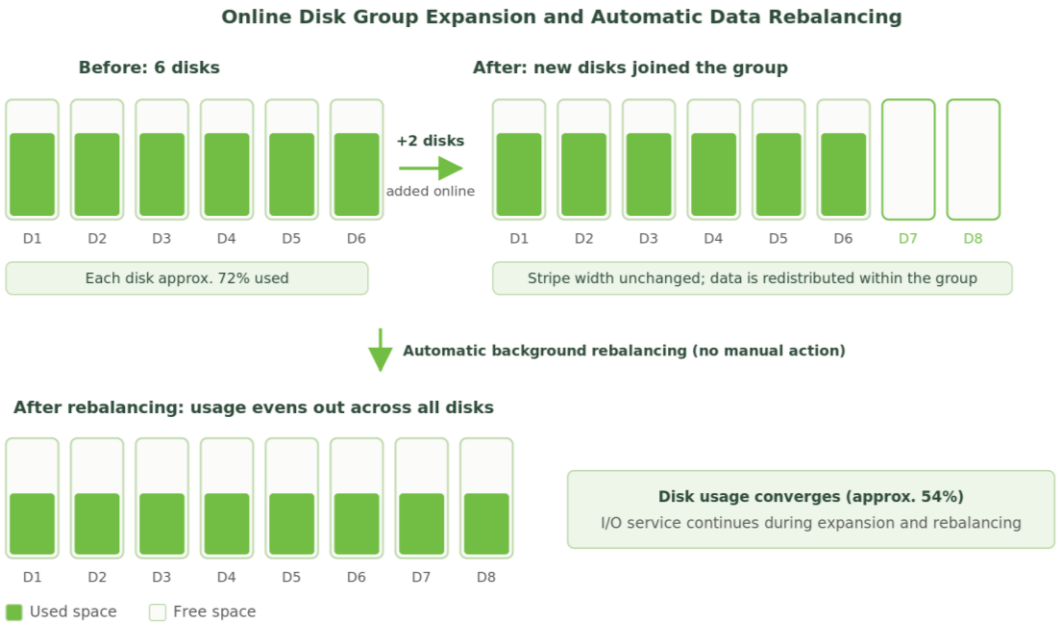


Figure 3-2 Online Disk Group Expansion and Automatic Data Balancing

The following describes the actions taken to expand online.

©2026 QSAN Technology, Inc. All rights reserved.

Table 3-4 Online Expansion Behavior

ITEM	BEHAVIOR
Disks Added at Once	Multiple disks may be added in a single operation
Service Interruption	None. I/O continues normally throughout the expansion
Stripe Width	Unchanged; expansion does not alter it
Disk Type	Only disks of the same type as the existing group (SSD groups accept SSDs, HDD groups accept HDDs)
Encryption	If the pool is encrypted, only SED disks may be added
Maximum Count	A single disk group must not exceed 63 disks after the addition
Shrinking	Not supported. Once added, a disk cannot be removed from the group

Automatic Data Balancing

After a new disk is added, the system automatically redistributes the data within the disk group, ensuring it is evenly distributed across all disks, including the new disk. No manual triggering or intervention is required. This balancing serves two purposes.

1. To ensure the new disk handles its appropriate I/O load.
2. To maintain a balance of available space among disks, so that in the event of a future failure, every disk has rebuildable space available.

The **Disk Information** dialog box in the management interface displays the used capacity of each disk in the QDP disk group, allowing you to observe the balancing progress and actual load distribution.

©2026 QSAN Technology, Inc. All rights reserved.



CAUTION

When a disk group has reached its maximum number of disks for data processing and parity checking (16 disks), adding more disks will provide available space. To increase available capacity, add another disk group to the same disk pool.

3.3. Operations

Creating a QDP Pool

In the pool creation wizard, QDP appears as the third pool type (alongside thick and thin pools) and is a default option. The creation process includes the following steps.

Create Pool

1. Select **QDP + Thin Provisioning** pool.
2. Input the pool name, encryption (SED) and auto tiering.

Create Disk Group

1. Select disks (capacities may be mixed freely), choose **QDP 5** or **QDP 6**.
2. Set the **Data Block Count** (the difference between disk count and stripe width is the spare). The interface shows the selected disk count against the limit, together with the expected capacity calculated by the backend.

Second Disk Group

1. Appears only when auto tiering is enabled; disks from a different tier must be selected. Behavior matches traditional RAID.

Create Volume

1. Identical to traditional RAID; no differences.

©2026 QSAN Technology, Inc. All rights reserved.

Review Summary

1. Pool type is shown as **QDP + Thin Provisioning** and the RAID level as **QDP 5** or **QDP 6**.



INFORMATION

After selecting QDP, deduplication and compression, EE spare, and subgroup settings will be hidden. Deduplication and compression are not supported; the number of spare disks is determined by the number of disks and data blocks (the wizard defaults to reserving disk capacity); subgroups are handled automatically by the system.

3.3.1. Post-Creation Management

The disk group panel on the QDP pool details page differs from a traditional RAID pool in the following ways.

- The RAID field is displayed as **QDP 5** or **QDP 6**.
- The **Data Block Count** field displays the value selected during creation.
- The disk count is presented as **used / limit**, with an **Expand Disks** button
- The **EE Spare** and **Sub-group** fields are hidden.
- The **Disk Information** dialog gains a **Used** column showing how much has been written to each disk within the group.

3.3.2. Monitoring

QDP pools use existing pool monitoring mechanisms, including dashboards, event logs, and alarm settings. Their operation is the same as traditional RAID pools, so administrators do not need to learn new monitoring workflows.

©2026 QSAN Technology, Inc. All rights reserved.

3.4. Best Practices

Disk Group Configuration

- Create a disk group containing 16 disks with a stripe width of 14, and reserve space equivalent to two disks' capacity within the group. This is the recommended initial configuration, balancing capacity and rebuild capability; the number of disks must be at least one more than the stripe width, otherwise the disk group will not be able to rebuild automatically after degradation.
- When more capacity is needed, add a new disk group to the same disk pool, rather than continuing to add disks to the existing disk group—adding more disks beyond 16 disks will not increase available capacity.
- For environments with the highest availability requirements, select QDP 6 and reserve more spare space in the disk group to cope with a second failure during rebuild.
- The stripe width cannot be modified after the disk group is created, so it must be planned correctly from the beginning.

Migrating from Traditional RAID

For customers performing a "traditional RAID 5 plus a dedicated hot spare disk" setup, the simplest migration method is to include the spare disk in the QDP disk group:

Table 3-5 Migrating from Traditional RAID

BEFORE MIGRATION	AFTER MIGRATION
RAID 5: five data disks plus one dedicated hot spare	QDP 5: all six disks in a single disk group (stripe width 5, one disk’s worth of spare)
The spare sits idle, providing neither capacity nor I/O	That disk’s capacity becomes spare space spread over six disks, serving I/O in normal operation
Rebuild writes concentrate on the spare disk	Rebuild is distributed across all healthy disks

©2026 QSAN Technology, Inc. All rights reserved.

Disk Selection and Feature Pairing

- There are no restrictions on the ratio of mixed-capacity disks, so you can safely place disks of different capacities in the same disk group; however, SSDs and HDDs cannot be mixed.
- QDP requires that at least one disk capacity be reserved as spare capacity, distributed across all disks in the disk group. After migration, the existing spare disk should be added to the disk group as the required spare capacity.
- Applications requiring deduplication and compression should continue to use thick provisioning pools, as these features are incompatible with QDP.
- Disk groups cannot be shrunk, and disks cannot be removed once added; therefore, the configuration should be verified before submitting disks.

©2026 QSAN Technology, Inc. All rights reserved.

4. CONCLUSION

QDP addresses three major structural problems inherent in the era of high-capacity hard drives. Excessive rebuild times, idle capacity caused by mixing hard drives of different capacities, and hot spares that contribute nothing during normal operation. These are all solved at the architectural level. By distributing each stripe across the entire disk group, the rebuild work is shared by all healthy disks, fully utilizing the capacity of each disk. Previously idle spare disks have become available space distributed across all disks for handling I/O requests. In an approximately 1 PB all-flash configuration, QDP 5's rebuild speed per GB is about 33% faster than traditional RAID 5.

Equally important is what parts QDP preserves. Since the changes are limited to the partition layer, everything above the virtual disk remains unchanged. Snapshots, replication, remote replication, and archive services continue to function normally, and administrators do not need to change system operation methods. QDP introduces a new parameter, strip width and a prerequisite, requiring at least one more disk than the stripe width as a spare. Here is a planning principle, build disk groups of 16 disks and add additional disk groups when capacity is insufficient.

©2026 QSAN Technology, Inc. All rights reserved.

5. APPENDIX

5.1. Apply To

- QSM firmware 4.3.0 and later
- XEVO firmware 3.3.0 and later
- XInsight 2.0.0 and later

5.2. Reference

Video

- QDP - Declustered RAID with Distributed Rebuild and No Hot Spare

©2026 QSAN Technology, Inc. All rights reserved.