

Lenovo RAID Introduction

Reference Information

Using a Redundant Array of Independent Disks (RAID) to store data remains one of the most common and cost-efficient methods to increase server's storage performance, availability, and capacity. RAID increases performance by allowing multiple drives to process I/O requests simultaneously. RAID can also prevent data loss in case of a drive failure by reconstructing (or rebuilding) the missing data from the failed drive using the data from the remaining drives.

This guide describes the RAID technology and its capabilities, compares various RAID levels, introduces Lenovo RAID controllers, and provides RAID selection guidance for ThinkSystem, ThinkServer, and System x servers.

Introduction to RAID

RAID array (also known as *RAID drive group*) is a group of multiple physical drives that uses a certain common method to distribute data across the drives. A *virtual drive* (also known as *virtual disk* or *logical drive*) is a partition in the drive group that is made up of contiguous data segments on the drives. Virtual drive is presented up to the host operating system as a physical disk that can be partitioned to create OS logical drives or volumes.

This concept is illustrated in the following figure.




Figure 1. RAID overview

The data in a RAID drive group is distributed by segments or strips. A *segment* or *strip* (also known as a *stripe unit*) is the portion of data that is written to one drive immediately before the write operation continues on the next drive. When the last drive in the array is reached, the write operation continues on the first drive in the block that is adjacent to the previous stripe unit written to this drive, and so on.

The group of strips that is subsequently written to all drives in the array (from the first drive to the last) before the write operation continues on the first drive is called a *stripe*, and the process of distributing data is called *striping*. A strip is a minimal element that can be read from or written to the RAID drive group, and strips contain data or recovery information.

The particular method of distributing data across drives in a drive group is known as the *RAID level*. The RAID level defines a level of fault tolerance, performance, and effective storage capacity because achieving redundancy always lessens disk space that is reserved for storing recovery information.

Recommendation: It is always recommended to combine into a single RAID drive group the drives of the same type, rotational speed, and size.

RAID levels

There are basic RAID levels (0, 1, 5, and 6) and spanned RAID levels (10, 50, and 60). *Spanned RAID arrays* combine two or more basic RAID arrays to provide higher performance, capacity, and availability by overcoming the limitation of the maximum number of drives per array that is supported by a particular RAID controller.

RAID 0

RAID 0 distributes data across all drives in the drive group, as shown in the following figure. RAID 0 is often called *striping*.




Figure 2. RAID 0

RAID 0 provides the best performance and capacity across all RAID levels, however, it does not offer any fault tolerance, and a drive failure causes data loss of the entire volume. The total capacity of a RAID 0 array is equal to the size of the smallest drive multiplied by the number of drives. RAID 0 requires at least two drives.

RAID 1

RAID 1 consists of two drives, and data are written to both drives simultaneously, as shown in the following figure. RAID 1 is also known as *mirroring*.




Figure 3. RAID 1

RAID 1 provides very good performance for both read and write operations, and it also offers fault tolerance. In case of a drive failure, the remaining drive will have a copy of the data from the failed drive, so data will not be lost. RAID 1 also offers fast rebuild time. However, the effective storage capacity is a half of total capacity of all drives in a RAID 1 array. The capacity of a RAID 1 array is equal to the size of the smallest drive.

RAID 5

RAID 5 distributes data and parity across all drives in the drive group, as shown in the following figure. RAID 5 is also called *striping with distributed parity*.

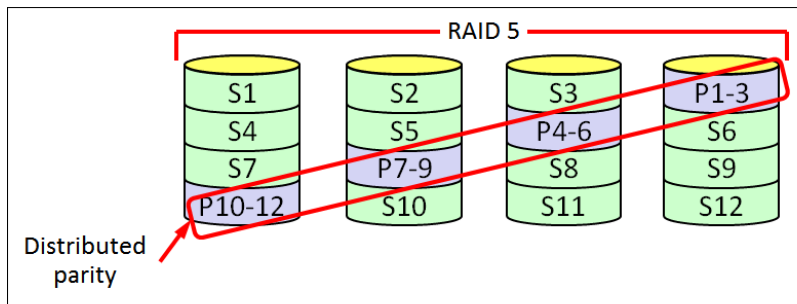


Figure 4. RAID 5

RAID 5 offers fault tolerance against a single drive failure with the smallest reduction in total storage capacity; however, it has slow rebuild time. RAID 5 provides excellent read performance similar to RAID 0, however, write performance is satisfactory due to the overhead of updating parity for each write operation. In addition, read performance can also suffer when the volume is in degraded mode, that is, in case of a drive failure.

The capacity of a RAID 5 array is equal to the size of the smallest drive multiplied by the number of drives less the capacity of the smallest drive. RAID 5 requires minimum three drives.

RAID 6

RAID 6 is similar to RAID 5, except that RAID 6 writes two parity segments for each stripe, as shown in the following figure, enabling a volume to sustain up to two simultaneous drive failures. RAID 6 is also known as *striping with dual distributed parity*.




Figure 5. RAID 6

RAID 6 offers the highest level of fault tolerance, and it provides excellent read performance similar to RAID 0, however, write performance is satisfactory (and slightly slower than RAID 5) due to the additional overhead of updating two parity blocks for each write operation. In addition, read performance can also suffer when the volume is in degraded mode, that is, in case of a drive failure.

The capacity of a RAID 6 array is equal to the size of the smallest drive multiplied by the number of drives less the twice capacity of the smallest drive. RAID 6 requires minimum four drives.

RAID 10

RAID 10 is a combination of RAID 0 and RAID 1 where data is striped across multiple RAID 1 drive groups, as shown in the following figure. RAID 10 is also known as *spanned mirroring*.




Figure 6. RAID 10

RAID 10 provides fault tolerance by sustaining a single drive failure within each span, and it offers very good performance with concurrent I/O processing on all drives. The capacity of a RAID 10 array is equal to a half of the total storage capacity. RAID 10 requires at least four drives.

RAID 50

RAID 50 is a combination of RAID 0 and RAID 5 where data is striped across multiple RAID 5 drive groups, as shown in the following figure. RAID 50 is also known as *spanned striping with distributed parity*.




Figure 7. RAID 50

RAID 50 provides fault tolerance by sustaining a single drive failure within each span, and it offers excellent read performance. It also improves write performance compared to RAID 5 due to separate parity calculations in each span. The capacity of a RAID 50 array is equal to the size of the smallest drive multiplied by the number of drives less the capacity of the smallest drive multiplied by the number of spans. RAID 50 requires minimum six drives.

RAID 60

RAID 60 is a combination of RAID 0 and RAID 6 where data is striped across multiple RAID 6 drive groups, as shown in the following figure. RAID 60 is also known as *spanned striping with dual distributed parity*.




Figure 8. RAID 60

RAID 60 provides best fault tolerance by sustaining two simultaneous drive failures within each span, and it offers excellent read performance. It also improves write performance compared to RAID 6 due to separate parity calculations in each span.

The total capacity of a RAID 60 array is equal to the size of the smallest drive multiplied by the number of drives less the capacity of the smallest drive multiplied by twice the number of spans. RAID 60 requires minimum eight drives.

Hot Spares

Hot Spares are designated drives that automatically and transparently take the place of a failed drive in the fault tolerant RAID arrays. This helps minimize the amount of time when the array remains in the degraded mode after the drive failure by automatically starting the rebuild process to restore the data from the failed drive on the hot spare.

Hot spare drives can be global or dedicated. A global hot spare drive can be used to replace the failed drive in any fault tolerant drive group as long as its capacity is equal or larger than the capacity of the failed drive used by the RAID array. A dedicated hot spare can only replace the failed drive in the designated drive group.

RAID level comparison

The selection of a RAID level is driven by the following factors:

- Read performance
- Write performance
- Fault tolerance
- Degraded array performance (for fault tolerant RAID levels)
- Effective storage capacity

The following table summarizes the RAID levels and their characteristics to assist you with selecting the most appropriate RAID level for your needs.

Table 1. RAID level comparison

Feature	RAID 0	RAID 1	RAID 5	RAID 6	RAID 10	RAID 50	RAID 60
Minimum drives	2†	2	3	4	4	6	8
Maximum drives*	32	2	32	32	16**	192***	192***
Tolerance to drive failures	None	1 drive	1 drive	2 drives	1 drive per span	1 drive per span	2 drives per span
Rebuild time	None	Fast	Slow	Slow	Fast	Slow	Slow
Read performance	Excellent	Very good	Excellent	Excellent	Very good	Excellent	Excellent
Write performance	Excellent	Very good	Satisfactory	Satisfactory	Very good	Good	Good
Degraded array performance	None	Very good	Satisfactory	Satisfactory	Very good	Good	Good
Capacity overhead	None	Half	1 drive	2 drives	Half	1 drive per span	2 drives per span

† Some implementations of RAID-0 allow only 1 drive, however this is not strictly RAID.

* The maximum number of drives in a RAID array is controller and system dependent.

** Up to 8 spans with 2 drives in each span.

*** Up to eight 24-drive external expansion enclosures; up to 8 spans with up to 32 drives in each span (the number of drives in the span should be the same across all spans).

Notes:

- The terms Excellent, Very good, Good, and Satisfactory are relative performance indicators for comparison purposes and do not represent any absolute values: Excellent is better than Very good, Very good is better than Good, and Good is better than Satisfactory.
- The terms Fast and Slow are relative rebuild time indicators for comparison purposes and do not represent any absolute values: Fast is faster than Slow.
- Degraded array performance means the performance of a RAID array with a failed drive or drives.

The following table summarizes benefits and drawbacks of each RAID level.

Table 2. RAID level benefits and drawbacks

RAID level	Benefits	Drawbacks
RAID 0	Fastest performance with full usable capacity of an array.	No data protection.
RAID 1	High read and write performance and data protection (tolerates a single drive failure) with fast rebuild time.	Highest loss in usable capacity (a half of total array capacity) with limited number of drives (2 drives).
RAID 5	High read performance and data protection (tolerates a single drive failure) with lowest loss in usable capacity.	Slow write performance and slow rebuild time.
RAID 6	High read performance and highest data protection (tolerates two drive failures).	Slow write performance and slow rebuild time.
RAID 10	High read and write performance and data protection (tolerates a single drive failure per span) with fast rebuild time.	Highest loss in usable capacity (a half of total array capacity) with limited number of drives (up to 16 drives).
RAID 50	High read performance and large array capacity (scales beyond 32 drives) with data protection (tolerates a single drive failure per span).	Slow write performance and slow rebuild time.
RAID 60	High read performance and large array capacity (scales beyond 32 drives) with highest data protection (tolerates two drive failures per span).	Slow write performance and slow rebuild time.

The following figure might help you select a RAID level based on your data storage requirements.

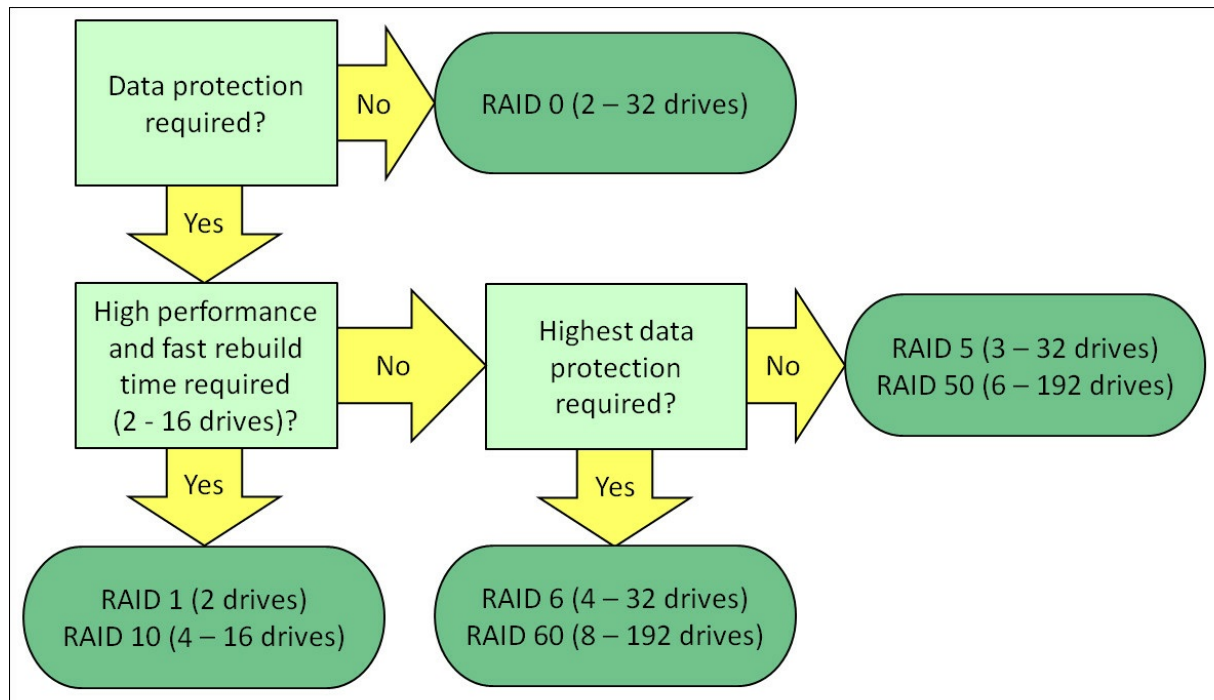


Figure 9. RAID selection guidance

Lenovo RAID controllers

Lenovo offers portfolio of RAID controllers for ThinkSystem, ThinkServer, and System x servers that is optimized to deliver the value and performance demanded by the ever-growing storage I/O requirements of today's enterprises. Lenovo RAID controllers are segmented into three categories based on price, performance, and features:

- Entry (Embedded Software RAID)
- Value (Basic hardware RAID)
- Advanced (Advanced hardware RAID)

The following table compares RAID controller categories.

Note: The terms Low, Medium, and High are relative price indicators for comparison purposes and do not represent any absolute values: Low is less expensive than Medium and Medium is less expensive than High.

Table 3. RAID controller category comparison

Feature	Embedded software RAID	Basic hardware RAID	Advanced hardware RAID
RAID processing	Host resources (CPU and memory)	Dedicated I/O controller + Host resources (CPU and memory) for RAID 5, 50	Dedicated I/O controller full offload
RAID levels	0, 1, 10, 5	0, 1, 10, 5, 50	0, 1, 10, 5, 50, 6, 60
Drive scalability	Limited (4 - 8 drives)	Moderate (8 - 32 drives)	High (192 drives)*
Advanced RAID features**	Limited	Yes	Yes
Cache	No	No	Yes
Cache protection	No	No	Yes
Price	Low	Medium	High

* With up to eight 24-drive external expansion enclosures.

** Advanced RAID features include Online Capacity Expansion, Online RAID Level Migration, Patrol Read, and Consistency Check.

The following table positions the Lenovo RAID controllers based on the interface and category.

Table 4. RAID controller positioning

Interface	Server Internal storage			External storage expansion
	Embedded software RAID	Basic hardware RAID	Advanced hardware RAID	Advanced hardware RAID
ThinkSystem RAID controllers				
12 Gbps SAS PCIe 4.0			<ul style="list-style-type: none"> ● 940-8i 4GB Flash ● 940-8i 8GB Flash ● 940-16i 8GB Flash ● 940-32i 8GB Flash 	
12 Gbps SAS PCIe 3.0		<ul style="list-style-type: none"> ● 530-8i ● 730-8i 1GB Cache* ● 730-8i 2GB Flash† 	<ul style="list-style-type: none"> ● 930-8i 2GB Flash ● 930-16i 4GB Flash ● 930-16i 8GB Flash ● 930-24i 4GB Flash 	<ul style="list-style-type: none"> ● 930-8e 4GB Flash
6 Gbps SATA	<ul style="list-style-type: none"> ● Onboard SATA (Intel RSTe) 			
System x ServeRAID controllers				
12 Gbps SAS PCIe 3.0		<ul style="list-style-type: none"> ● M1210 ● M1215 	<ul style="list-style-type: none"> ● M5210 ● M5210e 	<ul style="list-style-type: none"> ● M5225-2GB
6 Gbps SATA	<ul style="list-style-type: none"> ● C110 			
ThinkServer RAID controllers				
12 Gbps SAS PCIe 3.0		<ul style="list-style-type: none"> ● 520i 	<ul style="list-style-type: none"> ● 720i ● 720ix 	<ul style="list-style-type: none"> ● 9380-8e
6 Gbps SAS		<ul style="list-style-type: none"> ● 510i 		<ul style="list-style-type: none"> ● 9286CV-8e
6 Gbps SATA	<ul style="list-style-type: none"> ● 110i ● 121i 			

* ThinkSystem RAID 730-8i 1GB Cache adapter is available worldwide except US and Canada

† The 730-8i 2GB adapter also supports RAID 6 and 60.

Related publications and links

For more information, refer to the following publications:

- Lenovo ThinkSystem RAID Adapter and HBA Reference
<https://lenovopress.com/lp1288>
- Lenovo RAID Management Tools and Resources
<http://lenovopress.com/lp0579>

Related product families

Product families related to this document are the following:

- [RAID Adapters](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result 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.

© Copyright Lenovo 2024. All rights reserved.

This document, LP0578, was created or updated on May 5, 2020.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP0578>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP0578>.

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®

ServeRAID

System x®

ThinkServer®

ThinkSystem®

The following terms are trademarks of other companies:

Intel® is a trademark of Intel Corporation or its subsidiaries.

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