

Synology RAID F1 White Paper



Table of Contents

Executive Summary	01
Introduction to Synology RAID F1	02
RAID F1 Performance	
RAID Rebuild	
Reliability compared with RAID 5	
Conclusion	04

Find your information

Synology publishes a wide range of supporting documentation.

In **Knowledge Base**, you will find useful **Help** and **FAQ** articles, as well as **video tutorials** breaking up processes into handy steps.

In **Synology Documentation**, you can find **User's Guides**, **Solution Guides**, brochures, and **White Papers**. Experienced users and administrators will find answers and guidance in technical **Administrator's Guides** and **Developer Guides**.

Got a problem and unable to find the solution in our official documentation? Search hundreds of answers by users and support staff in **Synology Community** or reach **Synology Support** through the web form, email or telephone.

Executive Summary

RAID is widely used in enterprise storage products. It does not only prevent data loss in the event of a disk failure but also increases performance by spreading data across multiple disks. In the past, RAID implementations always assume that spinning disks are used as storage media instead of SSDs or flash disks. However, SSDs have a finite number of program-erase (P/ E) cycles. If traditional RAID is used for random write workloads, multiple SSDs will probably be worn out and fail simultaneously, resulting in a crashed RAID and data loss. **Synology RAID F1** algorithm tackles the problem by writing more parity bits into a specific SSD to avoid all SSDs from being worn out all at once and making one system-assigned SSD to be worn out first. With this approach, RAID F1 is expected not to crash as data are unevenly written to SSDs. Synology RAID F1 enhances the endurance of RAID compared to other RAID algorithms, an important consideration for enterprise flash storage products.

/ Introduction to Synology RAID F1

RAID F1 is the new RAID type for SSDs. **F** stands for flash, and **1** stands for 1-disk resiliency and 1-parity. The layout of RAID F1 is based on RAID 5. The difference between these two RAID types is that RAID F1 has an additional parity block per cycle. In the illustration below, the blocks with "P" stand for the parity blocks, and the rest are data blocks.



Since each write-operation involves writing onto a parity block, it is expected that the parity block will be the first to wear out. This uneven parity distribution causes an SSD to reach its lifespan earlier than others, rather than ending all SSDs' lifespan at the same time. When one SSD fails, the user can replace it with a new one.

The characteristics of RAID F1 is similar to RAID 5. Parity blocks are XOR'ed of all other data blocks. One block is used as a parity block within each stripe, so the usable capacity of a RAID F1 array is N-1 times of the smallest drive, where N is the stripe width or the number of disks.

RAID F1 Performance

In brief, RAID F1 provides the best balance between reliability and performance. The RAID F1 parity assignment, compared with RAID 4, provides more IOPS. The read and write performance of RAID F1 is close to that of RAID 5 with minimum performance differences because of extra parity blocks on RAID F1 volumes. There is no notable effect on the CPU utilization between RAID F1 and RAID 5.

RAID Rebuild

The goal of RAID F1 is to minimize the possibility of data loss. To achieve this goal, the system will find out the most aged SSD to rebuild and reallocate the parity layout. For example, when disk **A** has been removed, the array should be repaired by installing disk F. Later on the system will find out that disk **B** is the most aged disk, it will be set as the new system-assigned SSD into which more parity blocks will be written. This approach ensures that only **one** "most aged SSD" stores the most parity blocks.



Reliability compared with RAID 5

RAID 5 uses up all SSDs' life at the same pace because the parity is evenly distributed to all disks in the array. In this case, the possibility of data loss on a worn out SSD becomes very high along with RAID aging. In contrast, RAID F1 consistently minimizes data loss probability by making one of the most aged SSD reaches the end of its lifespan first.

Conclusion

In summary, Synology RAID F1 benefits from the following:

- Performance is close to RAID 5
- Intelligent RAID Rebuild
- Flash endurance superior to traditional RAID algorithm
- High storage efficiency with low capacity overhead of 8% (RAID F1 with 12 SSDs)

Synology RAID has a flash protection mechanism, which represents a giant leap in flash storage technology. RAID F1 is tailor-made and has the unique characteristics of flash storage, enabling Synology all-flash arrays to provide better data endurance and avoiding data loss. Providing IT administrators confidence to deploy Synology FS series products in their production and business critical environments.



SYNOLOGY

INC. 9F, No. 1, Yuan Dong Rd. Banqiao, New Taipei 220545 Taiwan Tel: +886 2 2955 1814

SYNOLOGY AMERICA CORP.

3535 Factoria Blvd SE, Suite #200, Bellevue, WA 98006 USA Tel: +1 425 818 1587

SYNOLOGY UK LTD.

Unit 5 Danbury Court, Linford Wood, Milton Keynes, MK14 6PL, United Kingdom Tel.: +44 (0)1908048029

SYNOLOGY FRANCE SARL

102 Terrasse Boieldieu (TOUR W) 92800 Puteaux France Tel: +33 147 176288

SYNOLOGY GMBH

Grafenberger Allee 295 40237 Düsseldorf Deutschland Tel: +49 211 9666 9666

SYNOLOGY SHANGHAI

200070, Room 201, No. 511 Tianmu W. Rd., Jingan Dist., Shanghai, China

SYNOLOGY JAPAN CO., LTD.

4F, No. 3-1-2, Higashikanda Chiyoda-ku, Tokyo, 101-0031 Japan





synology.com

Synology may make changes to specifications and product descriptions at any time, without notice. Copyright © 2021 Synology Inc. All rights reserved. I Synology and other names of Synology Products are proprietary marks or registered trademarks of Synology Inc. Other products and company names mentioned herein are trademarks of their respective holders.