

LSFS device technical description

2025

StarWind Documents



Trademarks

“StarWind”, “StarWind Software” and the StarWind and the StarWind Software logos are registered trademarks of StarWind Software. “StarWind LSFS” is a trademark of StarWind Software which may be registered in some jurisdictions. All other trademarks are owned by their respective owners.

Changes

The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, StarWind Software assumes no liability resulting from errors or omissions in this document, or from the use of the information contained herein. StarWind Software reserves the right to make changes in the product design without reservation and without notification to its users.

Technical Support and Services

If you have questions about installing or using this software, check this and other documents first - you will find answers to most of your questions on the [Technical Papers](#) webpage or in [StarWind Forum](#). If you need further assistance, please [contact us](#) .

About StarWind

StarWind is a pioneer in virtualization and a company that participated in the development of this technology from its earliest days. Now the company is among the leading vendors of software and hardware hyper-converged solutions. The company’s core product is the years-proven StarWind Virtual SAN, which allows SMB and ROBO to benefit from cost-efficient hyperconverged IT infrastructure. Having earned a reputation of reliability, StarWind created a hardware product line and is actively tapping into hyperconverged and storage appliances market. In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms” following the success and popularity of StarWind HyperConverged Appliance. StarWind partners with world-known companies: Microsoft, VMware, Veeam, Intel, Dell, Mellanox, Citrix, Western Digital, etc.

Copyright ©2009-2018 StarWind Software Inc.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of StarWind Software.

IMPORTANT NOTE: Starting from 14869 build, the LSFS device feature was excluded and deprecated. Please remove all Asynchronous replicas and LSFS devices before the update. See more information in the release notes:
<https://www.starwindsoftware.com/release-notes-build>

Technical Properties

1. [LSFS](#) is a log-structured file system that keeps track of the changes that are made in a journal. This file system keeps no data, changes only.
2. Old data are stored on the file system until online defragmentation cleans it up.
3. [LSFS](#) is a [snapshot based file system](#) and uses [snapshots](#) based on the concept of journaling file system. Every snapshot is incremental and occupies additional space that is equal to changes made since previous snapshot creation.
4. The journal divides into file segments for convenience. The minimum size of file segment is 128 MB, maximum – 512 MB. LSFS always keeps one additional empty file-segment.
5. Junk rate predetermines maximum allowed LSFS growth (overprovisioning) comparison to declared LSFS size. The default rate is 60% causing LSFS file-segments might use 2.5 times more space than initial LSFS size. Additionally metadata occupies up to 20% of initial LSFS size resulting a total overprovisioning of LSFS devices at 200%.
6. A contribution of metadata is usually small, but some patterns cause high metadata growth and it may occupy as much space as the useful data. The “worst” pattern is 4k random write. Nevertheless, when the whole disk is full, the ratio between metadata and useful data should be stabilized and 3 times growth is maximum possible. Information about useful data size, metadata and fragmentation can be found in StarWind Management Console when [LSFS](#) device is selected.
7. No matter which data access pattern is used, underlying storage always receives 4 MB blocks. Read patterns may vary starting from 4k blocks. According to mentioned above underlying storage may be a cheap spinning drive.
8. If available disk space on the physical storage is low, the LSFS uses more aggressive defragmentation policy and slows down access speed for the end user. If there is no space on the physical drive, then LSFS device becomes ‘read-only’.
9. 4k is the block, that LSFS uses. This means that write speed will be low if write requests are not aligned and deduplication will not be performed. But it shouldn't be the problem since all brand new drives use 4k blocks. Starting from Windows Vista it is possible to use such drives. Hyper-V aligns VHDX, but VHD might not be aligned. ESX cannot work with 4k drives, but VMFS should be aligned as well, thus there is no problem.

Limits And Requirements

1. Required free RAM (it is not related to L1 cache):
 - 4.6 GB of RAM per 1 TB initial LSFS size (with deduplication disabled)
 - 7.6 GB of RAM per 1 TB initial LSFS size (with deduplication enabled)
2. Maximum size of LSFS device is 64 TB.
3. Over-provisioning is 200% (LSFS files can occupy 3 times more space compared to initial LSFS size). Snapshots require having additional space to store them.
4. The physical block size of LSFS device is 4k. 512b or 4k are the sizes of virtual blocks. It is recommended to use 4k blocks for Microsoft Windows Environments and 512b blocks for VMware ESX.

Features And Benefits








1. Defragmentation works continuously in a background. File-segment will be defragmented, when data capacity is out of allowed value. Data from the old file-segmented will be moved to another empty file segment and the old fragmented file will be deleted. Maximum allowed junk rate before defragmentation process is 60%. This value can be changed using the context menu of the device in StarWind Management Console.
2. It is possible to run full defragmentation via device context menu (for standalone devices only). Manually started defragmentation ignores default defragmentation rate and cleans up all junk blocks inside file-segments.
3. HA based on LSFS uses snapshots for synchronization purposes. Snapshots are created every 5 minutes by default and then removed. HA synchronizes only latest changes after any failure, because each HA partner has healthy snapshot before the failure and there is no need to synchronize all the data but latest changes only that were made after the snapshot was created. Full synchronization takes part only after initial replica creation. Even in case of full synchronization, only useful data is replicated skipping junk data.
4. [LSFS](#) creates restore points during a usual job. This is a journal part, which will be used in the case of failure. Restore points can be observed via Snapshot Manager in Device Recovery Mode.
5. [Deduplication](#) analyzes unique chunks of data identifying and storing them. As the analysis continues, other chunks are being compared to the stored copy and whenever a match occurs, the redundant chunk is replaced with a reference shortcut that points to the stored chunk. Given that the same byte pattern may occur dozens, hundreds, or even thousands of times (the matching frequency is dependent on the chunk size), the amount of data to be stored or transferred can be greatly reduced. For example, similar

VMs that reside on top of LSFS device will be deduplicated well, but pagefile of each VM will not. According to our observations deduplication is not applicable to pagefiles, thus the size of every pagefile will be added as unique data. Long story short, 10 similar VMs 12GB each (2GB for pagefile) occupy $(12-2)+2*10=30\text{GB}$.

Request A Product Feature

To request a new product feature or to provide feedback on a StarWind product, please email to our support at support@starwind.com and put "Request a Product Feature" as the subject.

Contacts

US Headquarters	EMEA and APAC
 +1 617 829 44 95	 +44 2037 691 857 (United Kingdom)
 +1 617 507 58 45	 +49 800 100 68 26 (Germany)
 +1 866 790 26 46	 +34 629 03 07 17 (Spain and Portugal)
	 +33 788 60 30 06 (France)

Customer Support Portal: <https://www.starwind.com/support>

Support Forum: <https://www.starwind.com/forums>

Sales: sales@starwind.com

General Information: info@starwind.com



StarWind Software, Inc. 100 Cummings Center Suite 224-C Beverly MA 01915, USA
www.starwind.com ©2025, StarWind Software Inc. All rights reserved.