

# IBM Spectrum Archive Enterprise Edition V1.2.6 Installation and Configuration Guide

Larry Coyne

Khanh Ngo

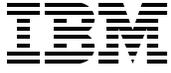
Illarion Borisevich

Markus Schäfer



Storage





International Technical Support Organization

**IBM Spectrum Archive Enterprise Edition V1.2.6:  
Installation and Configuration Guide**

June 2018

**Note:** Before using this information and the product it supports, read the information in “Notices” on page ix.

**Sixth Edition (June 2018)**

This edition applies to Version 1, Release 2, Modification 6 of IBM Spectrum Archive Enterprise Edition (product number 5639-LP1).

© Copyright International Business Machines Corporation 2014, 2018. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

<b>Notices</b> .....	ix
Trademarks .....	x
<b>Preface</b> .....	xi
Authors .....	xi
Now you can become a published author, too! .....	xiii
Comments welcome .....	xiii
Stay connected to IBM Redbooks .....	xiii
<b>Summary of changes</b> .....	xv
June 2018, Sixth Edition Version 1.2.6 .....	xv
January 2018, Fifth Edition Version 1.2.5.1 .....	xv
August 2017, Fourth Edition Version 1.2.4 .....	xvi
February 2017, Third Edition Version 1.2.2 minor update .....	xvi
January 2017, Third Edition Version 1.2.2 .....	xvii
August 2016, Second Edition Version 1.2.1 .....	xvii
June 2016, First Edition, Version 1.2 .....	xviii
January 2015, Third Edition Version 1.1.1.2 .....	xix
November 2014, Second Edition Version 1.1.1.1 .....	xix
<b>Chapter 1. IBM Spectrum Archive Enterprise Edition</b> .....	1
1.1 Introduction .....	2
1.1.1 Operational storage .....	4
1.1.2 Active archive .....	5
1.2 IBM Spectrum Archive EE functions .....	6
1.3 IBM Spectrum Archive EE components .....	7
1.3.1 IBM Spectrum Archive EE terms .....	7
1.3.2 Hierarchical Storage Manager .....	11
1.3.3 Multi-Tape Management Module .....	11
1.3.4 IBM Spectrum Archive Library Edition Plus component .....	13
1.4 IBM Spectrum Archive EE cluster configuration introduction .....	13
<b>Chapter 2. IBM Spectrum Archive overview</b> .....	15
2.1 Introduction to IBM Spectrum Archive and LTFS .....	16
2.1.1 Tape media capacity with IBM Spectrum Archive .....	18
2.2 Comparison of the IBM Spectrum Archive products .....	19
2.2.1 IBM Spectrum Archive Single Drive Edition .....	20
2.2.2 IBM Spectrum Archive Library Edition .....	21
2.2.3 IBM Spectrum Archive Enterprise Edition .....	23
2.3 IBM Spectrum Scale .....	25
2.3.1 Overview .....	26
2.3.2 Storage pools .....	26
2.3.3 Policies and policy rules .....	27
2.3.4 Migration or premigration .....	28
2.3.5 Active File Management .....	29
2.4 OpenStack SwiftHLM .....	30
2.5 IBM Spectrum Archive EE dashboard .....	32
2.6 IBM Spectrum Archive EE REST API .....	34
2.7 Types of archiving .....	34

<b>Chapter 3. Planning</b> .....	37
3.1 System requirements .....	38
3.1.1 Limitations .....	39
3.2 Required software .....	39
3.3 Hardware and software setup .....	41
3.4 Sizing and settings .....	43
3.4.1 IBM Spectrum Archive EE metadata file system .....	43
3.4.2 Redundant copies .....	44
3.4.3 Performance .....	44
<b>Chapter 4. Installation</b> .....	49
4.1 Installing IBM Spectrum Archive EE on a Linux system .....	50
4.2 Installation prerequisites for IBM Spectrum Archive EE .....	50
4.2.1 Installing the host bus adapter and device driver .....	51
4.2.2 Installing the IBM tape device driver .....	51
4.3 Installing IBM Spectrum Archive EE .....	57
4.3.1 Extracting binary rpm files from an installation package .....	58
4.3.2 Installing, upgrading, or uninstalling IBM Spectrum Archive EE .....	60
4.4 Installing a RESTful server .....	66
4.5 Quick installation guide for IBM Spectrum Archive EE .....	71
4.6 Library replacement .....	71
4.6.1 Library replacement procedure .....	72
4.6.2 Pool relocation procedure .....	74
<b>Chapter 5. Upgrading from version 1.1.x</b> .....	79
5.1 Overview of the IBM Spectrum Archive EE upgrade from version 1.1.x .....	80
5.2 The <code>ltfsee_config_save</code> command .....	81
5.3 The <code>ltfsee_config_upgrade</code> command .....	82
5.4 Upgrading from IBM Spectrum Archive Enterprise Edition version 1.1.x .....	82
5.4.1 Prerequisite tasks for the upgrade .....	82
5.4.2 Upgrading your operating system .....	84
5.4.3 Upgrading IBM Spectrum Archive Enterprise Edition .....	84
5.4.4 Post-upgrade tasks .....	86
5.5 Upgrade processing output examples .....	86
5.5.1 Example output from <code>ltfsee_config_save</code> .....	86
5.5.2 Example of creating and moving the <code>upgrade_offline.lst</code> in the home directory to <code>&lt;metadata_filesystem_name&gt;/ltfsee/upgrade</code> .....	88
5.5.3 Example output from <code>ltfsee_install --upgrade</code> .....	88
5.5.4 Example output from <code>ltfsee_config_upgrade</code> .....	90
<b>Chapter 6. Configuration</b> .....	95
6.1 Configuration prerequisites .....	96
6.1.1 Configuration worksheet tables .....	96
6.1.2 Obtaining configuration information .....	100
6.1.3 Configuring key-based login with OpenSSH .....	102
6.1.4 Preparing the IBM Spectrum Scale file system for IBM Spectrum Archive EE ..	102
6.1.5 Enabling control paths in the logical library .....	104
6.2 Configuring IBM Spectrum Archive EE .....	104
6.2.1 The <code>ltfsee_config</code> utility .....	104
6.2.2 Configuring a single node cluster .....	106
6.2.3 Configuring a multiple-node cluster .....	110
6.2.4 Configuring a multiple-node cluster with two tape libraries .....	113
6.2.5 Modifying a multiple-node configuration for control node redundancy .....	114
6.3 First-time start of IBM Spectrum Archive EE .....	116

6.4	Configuring IBM Spectrum Archive EE with IBM Spectrum Scale AFM . . . . .	118
6.4.1	Configuring a Centralized Archive Repository solution . . . . .	118
6.4.2	Configuring an Asynchronous Archive Replication solution . . . . .	121
<b>Chapter 7.</b>	<b>Operations . . . . .</b>	<b>129</b>
7.1	Overview . . . . .	131
7.1.1	IBM Spectrum Archive EE command summaries . . . . .	131
7.1.2	Using the command-line interface . . . . .	134
7.2	Status information . . . . .	134
7.2.1	IBM Spectrum Scale . . . . .	134
7.2.2	IBM Spectrum Archive Library Edition Plus component . . . . .	136
7.2.3	Hierarchical Space Management . . . . .	136
7.2.4	IBM Spectrum Archive EE . . . . .	138
7.3	Upgrading components . . . . .	139
7.3.1	IBM Spectrum Scale . . . . .	139
7.3.2	IBM Spectrum Archive LE+ component . . . . .	143
7.3.3	Hierarchical Storage Management . . . . .	143
7.3.4	IBM Spectrum Archive EE . . . . .	143
7.4	Starting and stopping IBM Spectrum Archive EE . . . . .	143
7.4.1	Starting IBM Spectrum Archive EE . . . . .	143
7.4.2	Stopping IBM Spectrum Archive EE . . . . .	144
7.5	IBM Spectrum Archive EE database backup . . . . .	145
7.6	IBM Spectrum Archive EE automatic node failover . . . . .	145
7.6.1	IBM Spectrum Archive EE monitoring daemon . . . . .	146
7.7	Tape library management . . . . .	147
7.7.1	Adding tape cartridges . . . . .	147
7.7.2	Moving tape cartridges . . . . .	149
7.7.3	Formatting tape cartridges . . . . .	150
7.7.4	Removing tape drives . . . . .	153
7.7.5	Adding tape drives . . . . .	154
7.8	Tape storage pool management . . . . .	154
7.8.1	Creating tape cartridge pools . . . . .	155
7.8.2	Deleting tape cartridge pools . . . . .	155
7.9	Pool capacity monitoring . . . . .	156
7.10	Migration . . . . .	159
7.10.1	Managing file migration pools . . . . .	161
7.10.2	Threshold-based migration . . . . .	161
7.10.3	Manual migration . . . . .	166
7.10.4	Replicas and redundant copies . . . . .	170
7.10.5	Data Migration . . . . .	173
7.10.6	Migration hints and tips . . . . .	175
7.11	Premigration . . . . .	177
7.11.1	Premigration with the ltfsee premigrate command . . . . .	178
7.11.2	Premigration running the mmapplypolicy command . . . . .	178
7.12	Preserving file system objects on tape . . . . .	179
7.12.1	Saving file system objects with the ltfsee save command . . . . .	179
7.12.2	Saving file system objects with policies . . . . .	180
7.13	Restoring non-empty regular files and file system objects from tape . . . . .	182
7.14	Recall . . . . .	183
7.14.1	Transparent recall . . . . .	184
7.14.2	Selective recall . . . . .	184
7.14.3	Recalling files with the ltfsee recall command . . . . .	185
7.14.4	Read Starts Recalls: Early trigger for recalling a migrated file . . . . .	186

7.15	Repairing files to their resident state . . . . .	187
7.16	Handling of write-failure tapes . . . . .	189
7.17	Handling of read-failure tapes . . . . .	189
7.18	Reconciliation . . . . .	189
7.19	Reclamation . . . . .	192
7.19.1	Reclamation considerations . . . . .	192
7.19.2	Reclamation performance . . . . .	193
7.20	Checking and repairing . . . . .	195
7.21	Importing and exporting . . . . .	196
7.21.1	Importing . . . . .	196
7.21.2	Exporting . . . . .	199
7.21.3	Import/Export enhancements . . . . .	202
7.22	Drive Role settings for job assignment control . . . . .	203
7.23	Tape drive intermix support . . . . .	204
7.23.1	Objective for WORM tape support . . . . .	206
7.23.2	Function overview for WORM tape support . . . . .	206
7.23.3	The effects of file operations on immutable and appendOnly files . . . . .	207
7.24	Obtaining the location of files and data . . . . .	210
7.25	Obtaining inventory, job, and scan status . . . . .	210
7.26	Cleaning up a scan or session . . . . .	212
7.27	Monitoring the system with SNMP . . . . .	212
7.28	Configuring Net-SNMP . . . . .	214
7.28.1	Starting and stopping the snmpd daemon . . . . .	214
7.28.2	Example of an SNMP trap . . . . .	215
7.29	IBM Spectrum Archive REST API . . . . .	215
7.29.1	Pools endpoint . . . . .	215
7.29.2	Tapes endpoint . . . . .	218
7.29.3	Libraries endpoint . . . . .	221
7.29.4	Nodegroups endpoint . . . . .	221
7.29.5	Nodes endpoint . . . . .	221
7.29.6	Drives endpoint . . . . .	222
<b>Chapter 8. Hints, tips, and preferred practices . . . . .</b>		<b>225</b>
8.1	Preventing migration of the .SPACEMAN and metadata directories . . . . .	227
8.2	Maximizing migration performance with redundant copies . . . . .	227
8.3	Changing the SSH daemon settings . . . . .	228
8.4	Setting mmapplypolicy options for increased performance . . . . .	229
8.5	Setting the inode size for the GPFS file system for increased performance . . . . .	231
8.6	Determining the file states for all files within the GPFS file system . . . . .	231
8.7	Memory considerations on the GPFS file system for increased performance . . . . .	233
8.8	Increasing the default maximum number of inodes in IBM Spectrum Scale . . . . .	234
8.9	Configuring IBM Spectrum Scale settings for performance improvement . . . . .	235
8.10	Real world use cases for mmapplypolicy . . . . .	235
8.10.1	Creating a traditional archive system policy . . . . .	235
8.10.2	Creating active archive system policies . . . . .	236
8.10.3	IBM Spectrum Archive EE migration policy with AFM . . . . .	238
8.11	Capturing a core file on RHEL with abrt . . . . .	238
8.12	Antivirus considerations . . . . .	239
8.13	Automatic email notification with rsyslog . . . . .	239
8.14	Overlapping IBM Spectrum Scale policy rules . . . . .	240
8.15	Storage pool assignment . . . . .	241
8.16	Tape cartridge removal . . . . .	242
8.16.1	Reclaiming tape cartridges before you remove or export them . . . . .	242

8.16.2	Exporting tape cartridges before physically removing them from the library. . .	242
8.17	Reusing LTFS formatted tape cartridges . . . . .	242
8.17.1	Reformatting LTFS tape cartridges through ltfsee commands . . . . .	242
8.18	Reusing non-LTFS tape cartridges . . . . .	244
8.19	Moving tape cartridges between pools . . . . .	244
8.19.1	Avoiding changing assignments for tape cartridges that contain files. . . . .	245
8.19.2	Reclaiming a tape cartridge and changing its assignment . . . . .	245
8.20	Offline tape cartridges . . . . .	245
8.20.1	Do not modify the files of offline tape cartridges . . . . .	245
8.20.2	Solving problems . . . . .	245
8.21	Scheduling reconciliation and reclamation . . . . .	245
8.22	License Expiration Handling . . . . .	246
8.23	Disaster recovery . . . . .	247
8.23.1	Tiers of disaster recovery . . . . .	247
8.23.2	Preparing IBM Spectrum Archive EE for a tier 1 disaster recovery strategy (offsite vaulting) . . . . .	248
8.23.3	IBM Spectrum Archive EE tier 1 DR procedure . . . . .	250
8.24	IBM Spectrum Archive EE problem determination . . . . .	252
8.25	Collecting IBM Spectrum Archive EE logs for support . . . . .	253
8.26	Backing up file systems that are not managed by IBM Spectrum Archive EE . . . . .	255
8.26.1	Considerations . . . . .	255
8.26.2	Backing up a GPFS or IBM Spectrum Scale environment . . . . .	256
8.26.3	Backing up non-IBM Spectrum Scale file systems . . . . .	256
8.26.4	IBM TS4500 Automated Media Verification with IBM Spectrum Archive EE . . . . .	258
<b>Chapter 9. Use cases</b>		<b>265</b>
9.1	Use cases overview . . . . .	266
9.1.1	Use case for archive . . . . .	266
9.1.2	Use case for tiered and scalable storage . . . . .	267
9.1.3	Use case data exchange . . . . .	268
9.2	AFM use cases . . . . .	269
9.2.1	Centralized Archive Repository . . . . .	269
9.2.2	Asynchronous Archive Replication . . . . .	270
<b>Chapter 10. Troubleshooting</b>		<b>273</b>
10.1	Overview . . . . .	274
10.1.1	Quick health check . . . . .	274
10.2	Hardware . . . . .	274
10.2.1	Tape library . . . . .	275
10.2.2	Tape drives . . . . .	275
10.2.3	Tape cartridge . . . . .	279
10.3	Recovering data from a write failure tape . . . . .	284
10.4	Recovering data from a read failure tape . . . . .	286
10.5	Handling export errors . . . . .	287
10.6	Software . . . . .	288
10.6.1	Linux . . . . .	289
10.6.2	IBM Spectrum Scale . . . . .	291
10.6.3	IBM Spectrum Archive LE+ component . . . . .	293
10.6.4	Hierarchical storage management . . . . .	294
10.6.5	IBM Spectrum Archive EE logs . . . . .	295
10.7	Recovering from system failures . . . . .	297
10.7.1	Power failure . . . . .	297
10.7.2	Mechanical failure . . . . .	297

10.7.3 Inventory failure . . . . .	298
10.7.4 Abnormal termination . . . . .	298
<b>Chapter 11. Reference . . . . .</b>	<b>299</b>
11.1 Command-line reference. . . . .	300
11.1.1 IBM Spectrum Archive EE commands . . . . .	300
11.1.2 IBM Spectrum Scale commands. . . . .	324
11.1.3 Tivoli Storage Manager for Space Management commands . . . . .	324
11.2 Formats for IBM Spectrum Scale to IBM Spectrum Archive EE migration . . . . .	326
11.3 System calls and IBM tools. . . . .	328
11.3.1 Downloading the IBM Tape Diagnostic Tool. . . . .	328
11.3.2 Using the IBM LTFS Format Verifier. . . . .	329
11.4 IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products . . .	331
<b>Appendix A. Additional material . . . . .</b>	<b>333</b>
Locating the web material . . . . .	333
Using the web material. . . . .	333
Downloading and extracting the web material . . . . .	333
<b>Related publications . . . . .</b>	<b>335</b>
IBM Redbooks . . . . .	335
Other publications . . . . .	335
Online resources . . . . .	335
Help from IBM . . . . .	336

# Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

*IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US*

INTERNATIONAL BUSINESS MACHINES CORPORATION 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. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

## COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

# Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

GPFST™	IBM Spectrum Scale™	Storwize®
IBM®	IBM Spectrum Storage™	System Storage®
IBM Spectrum™	Linear Tape File System™	Tivoli®
IBM Spectrum Archive™	Redbooks®	
IBM Spectrum Protect™	Redbooks (logo)  ®	

The following terms are trademarks of other companies:

Intel, Intel Xeon, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Linear Tape-Open, LTO, Ultrium, the LTO Logo and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

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

# Preface

This IBM® Redbooks® publication helps you with the planning, installation, and configuration of the new IBM Spectrum™ Archive V1.2.6 for the IBM TS3310, IBM TS3500, IBM TS4300, and IBM TS4500 tape libraries. IBM Spectrum Archive™ EE enables the use of the LTFS for the policy management of tape as a storage tier in an IBM Spectrum Scale™ based environment. It helps encourage the use of tape as a critical tier in the storage environment. This is the sixth edition of IBM Spectrum Archive Installation and Configuration Guide.

IBM Spectrum Archive EE can run any application that is designed for disk files on a physical tape media. IBM Spectrum Archive EE supports the IBM Linear Tape-Open (LTO) Ultrium 8, 7, 6, and 5 tape drives in IBM TS3310, TS3500, TS4300, and TS4500 tape libraries. In addition, IBM TS1155, TS1150, and TS1140 tape drives are supported in TS3500 and TS4500 tape library configurations.

IBM Spectrum Archive EE can play a major role in reducing the cost of storage for data that does not need the access performance of primary disk. The use of IBM Spectrum Archive EE to replace disks with physical tape in tier 2 and tier 3 storage can improve data access over other storage solutions because it improves efficiency and streamlines management for files on tape. IBM Spectrum Archive EE simplifies the use of tape by making it transparent to the user and manageable by the administrator under a single infrastructure.

This publication is intended for anyone who wants to understand more about IBM Spectrum Archive EE planning and implementation. This book is suitable for IBM clients, IBM Business Partners, IBM specialist sales representatives, and technical specialists.

## Authors

This book was produced by a team working at the IBM Tokyo and IBM Tucson development labs.



**Larry Coyne** is a Project Leader at the International Technical Support Organization, Tucson Arizona center. He has 35 years of IBM experience with 23 in IBM storage software management. He holds degrees in Software Engineering from the University of Texas at El Paso and Project Management from George Washington University. His areas of expertise include client relationship management, quality assurance, development management, and support management for IBM Storage Software.



**Khanh Ngo** is an IBM Senior Technical Staff Member (STSM) in Tucson, Arizona serving as the Tape Storage Test Architect. He sets test strategies, designs tests, and tests solutions for the tape drive, tape library, and IBM Spectrum Archive product lines. He joined IBM in 2000 with a Bachelor of Science degree in Electrical Engineering and a Bachelor of Science in Computer Science. Later, he received a Master of Science degree in Engineering Management. He is also an IBM Master Inventor and on his IBM 5th patent plateau. Because of his work with various IBM Spectrum Archive Enterprise Edition (EE) customers across multiple industries, Khanh is often sought out for his expertise to lead, execute, and successfully complete proof of concepts and custom engineering solutions to integrate IBM Spectrum Archive EE into customers' production environments.



**Illarion Borisevich** is an IBM software developer in Tucson, Arizona as an FVT engineer for both IBM Spectrum Archive and the IBM next generation of storage management system, the IBM TS4500 tape library. He assists customers with setting up the IBM TS4500 and is a customer advocate for IBM Spectrum Archive. He has a bachelor's degree in computer science from the University of Arizona. Illarion works on the TS3500 CLI and develops scripts to stress both IBM Spectrum Archive EE and the IBM TS4500 tape library solution systems.



**Markus Schäfer** is an IBM Senior Accredited IT Specialist at the ESCC Tape Solution Center in Mainz, Germany. He has 17 years of experience in tape storage systems and tape solutions as a product field engineer. In his current job role as OEM Product Application Engineer (PAE), he is responsible for all EMEA OEM storage clients in terms of technical solutions and sales opportunities.

Thanks to the following people for their contributions and support for this project:

Atsushi Abe  
Said Ahmad  
Takashi Ashida  
Hiroshi Itagaki  
Takeshi Ishimoto  
Hiroyuki Miyoshi  
Hironobu Nagura  
**IBM Systems**

Thanks to the authors of the previous editions of the IBM Spectrum Archive EE Redbooks:

Illarion Borisevich, Larry Coyne, Chris Hoffmann, Stefan Neff, Khanh Ngo, Wei Zheng Ong

## Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at this website:

<http://www.ibm.com/redbooks/residencies.html>

## Comments welcome

Your comments are important to us!

We want our books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

- ▶ Use the online **Contact us** review Redbooks form found at:

<http://www.ibm.com/redbooks>

- ▶ Send your comments in an email to:

[redbooks@us.ibm.com](mailto:redbooks@us.ibm.com)

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization  
Dept. HYTD Mail Station P099  
2455 South Road  
Poughkeepsie, NY 12601-5400

## Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- ▶ Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>



# Summary of changes

This section describes the technical changes that are made in this edition of the book and in previous editions. This edition might also include minor corrections and editorial changes that are not identified.

Summary of Changes  
for SG24-8333-05  
for IBM Spectrum Archive Enterprise Edition V1.2.6: Installation and Configuration Guide  
as created or updated on June 22, 2018.

## June 2018, Sixth Edition Version 1.2.6

### **New information Version 1.2.6**

- ▶ Added support for IBM Power Server in Little Endian Mode (Power8 8 or later)
- ▶ Library replacement procedure phase two
- ▶ Tape intermix in pool for technology upgrade
- ▶ Datamigrate command for technology upgrade

### **Changed Information**

- ▶ Updated RHEL supported levels and IBM Spectrum Scale levels
- ▶ Updated various examples to include new information

## January 2018, Fifth Edition Version 1.2.5.1

### **New information Version 1.2.5.1**

- ▶ Provides new tape media support for LTO 8 tape drives with LTO 8 Type M cartridge (M8). The LTO program introduced a new capability with LTO 8 tape drives: the ability to write 9 TB (native) on a brand new LTO Ultrium 7 cartridge instead of 6 TB (native) as specified by the LTO 7 format.
- ▶ Added support for Red Hat Enterprise Linux Server 7.4.

### **New information Version 1.2.5**

- ▶ Provides support for the new 12 TB LTO 8 tape drive and TS1155 FC tape drive in the TS3500 tape library
- ▶ A library replacement procedure has been provided to allow the replacing of an old tape library (for example, TS3500 tape library) with a new tape library (for example, TS4500 tape library)

### **Changed Information**

- ▶ Upgraded the HSM component to version 8.1.2

## August 2017, Fourth Edition Version 1.2.4

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized below.

### New information

- ▶ Added support for Active File Management (AFM) Independent Writer (IW) mode starting with IBM Spectrum Archive V1.2.3.0 and later
- ▶ Added support for a RESTful API
- ▶ Added high availability features:
  - Control node failover
  - Monitoring daemon
  - New start/stop
- ▶ Added a GUI dashboard for data monitoring
- ▶ Added low pool threshold attribute for pools
- ▶ Added support for 15 TB tape support with TS1155 tape drive
- ▶ Added new `ltfsee node show` command
- ▶ Added new `ltfsee failover` command
- ▶ Added new IBM Spectrum Archive EE database backup
- ▶ Added IBM Swift HLM support

### Changed Information

- ▶ Added new traps to SNMP
- ▶ Updated `ltfsee info nodes` command

## February 2017, Third Edition Version 1.2.2 minor update

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized below.

- ▶ Added the following link for the new Performance white paper information in section 3.4.3, “Performance” on page 44.

**Note:** For additional migration performance information see the *IBM Spectrum Archive Enterprise Edition V1.2.2 Performance White Paper* at:

<https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102688>

- ▶ Updated the following examples: Example 7-57 on page 162, Example 7-59 on page 163, Example 7-60 on page 165, Example 7-61 on page 168, Example 7-63 on page 170, Example 8-7 on page 236, Example 8-8 on page 236, Example 8-9 on page 237, Example 8-21 on page 249.

## January 2017, Third Edition Version 1.2.2

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized below.

### New information

The following topics are new in Version 1.2.2:

- ▶ Added new write failure state, Write Fenced
- ▶ Added new pool remove option, -E for removing tapes with no file references
- ▶ Increased stability and performance
- ▶ Improved export/import
- ▶ Improved reconcile
- ▶ Automated the recover process of write failure tapes
- ▶ Added improved method for recovering read failure tapes

### Changed information

The following information was changed from the previous edition:

- ▶ Added performance section
- ▶ Renamed “Recovering data from a write-failure tape” to 7.16, “Handling of write-failure tapes” on page 189
- ▶ Updated the `ltfsee recover` command in 11.1, “Command-line reference” on page 300
- ▶ Added section about recovering data from write failure tapes
- ▶ Added section about recovering data from read failure tapes
- ▶ Added section for handling Normal Export errors
- ▶ Added “boost-filesystem” to system requirements
- ▶ Added section about memory considerations on the IBM Spectrum Scale file system for increased performance
- ▶ Added new information about how migrations are handled
- ▶ Added section about handling read failure tapes
- ▶ Added new rule for adding tapes into pools
- ▶ Added table about valid commands on different tape status

## August 2016, Second Edition Version 1.2.1

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized below.

### New information

What’s New in Version 1.2.1:

- ▶ Added procedures for upgrading from IBM Spectrum Archive Enterprise Edition version 1.1.x
- ▶ Added support for Red Hat Enterprise Linux (RHEL) 6 and Linux Enterprise Server 11
- ▶ Added information about avoiding errors during an upgrade on multiple nodes in 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60

## Changed information

- ▶ Removed commands for creating a list of files to use with migration, premigration, and save functions. V1.2.1.x and later now only uses the IBM GPFS™ scan result file.
- ▶ Updated the options for the `1tfsee_install` command.

## June 2016, First Edition, Version 1.2

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized below.

## New information

What's New in Version 1.2:

- ▶ Multiple tape library attachment (up to two) support to a single IBM Spectrum Scale cluster:
  - Data replication to the pools in separate libraries for additional data resiliency.
  - You can use the read replica policy to specify file recall behavior if there is a failure.
  - Total capacity expansion beyond a single library limit.
  - IBM Spectrum Scale cluster in single site or across metro distance locations through IBM Spectrum Scale synchronous mirroring.
  - Mixed type library support with different drive types.
  - Pools can be configured with a subset of tapes in one tape library as an LTO tape pool, 3592 tape pool, or 3592 write-once, read-many (WORM) tape pool.
- ▶ Data recording on WORM tape cartridges:
  - WORM cartridges are great for long-term records retention applications.
  - Support for 3592 WORM cartridges (3592-JY and 3592-JZ tapes).
  - Selection of rewritable or WORM cartridges per tape pool.
  - Files can be migrated (or premigrated) to the specified WORM pool as the destination in the IBM Spectrum Scale policy file or CLI option.
  - Files on WORM tapes will not be overwritten and are not erasable:
    - Reclamation, reconciliation, and reformat operations are disabled.
    - As a preferred practice, set the immutable flag on IBM Spectrum Scale disk to prevent the deletion of stub file for additional data protection.
- ▶ Expand storage capacity with LTO7 support:
  - You have 2.4 - 4 times the capacity in the same foot print compared to LTO6 and LTO5 technology.
  - Support for migration of files larger than 2.x TB (up to 6.x TB).
  - Intermixing of LTO generation in single library is supported (the pool must be homogeneous).
- ▶ Performance improvement for large-scale systems:
  - Optimization of file migration operations for many small files and for multi-node configuration.
  - Reduced IBM Spectrum Scale file system size requirements for IBM Spectrum Archive EE metadata.

- Collocation of files within a pool for speedier file recall.
- ▶ Flexibility in pool-based data management:
  - Improved automatic recovery process on failure by switching to replica pools.
  - Improved support of tape drive intermix.

## January 2015, Third Edition Version 1.1.1.2

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized below.

### **New information**

- ▶ Added support for IBM GPFS V4.1
- ▶ Added support for IBM TS1150 tape drive and media
- ▶ V1.1.1.2 (PGA2.1) updates

### **Changed information**

Added the `-u` option to the `reconcile` command to skip the pretest that checks for the necessity to reconcile before mounting the tapes.

## November 2014, Second Edition Version 1.1.1.1

This revision reflects the addition, deletion, or modification of new and changed information, which is summarized below:

### **New information**

- ▶ Premigration commands
- ▶ Added support for IBM TS4500 and TS3310 tape libraries
- ▶ Additional OS platform support
- ▶ Preserving file system objects on tape
- ▶ Recall commands
- ▶ Version 1.1.1.0 (PGA1) and Version 1.1.1.1 (PGA2) updates

### **Changed information**

- ▶ Improved small file migration performance
- ▶ Improved data resiliency (All copies are now referenced.)





# IBM Spectrum Archive Enterprise Edition

This chapter introduces the IBM Spectrum Archive Enterprise Edition (formerly IBM Linear Tape File System™ Enterprise Edition (LTFS EE)) and describes its business benefits, general use cases, technology, components, and functions.

This chapter includes the following topics for IBM Spectrum Archive Enterprise Edition (EE):

- ▶ Introduction
- ▶ IBM Spectrum Archive EE functions
- ▶ IBM Spectrum Archive EE components
- ▶ IBM Spectrum Archive EE cluster configuration introduction

# 1.1 Introduction

IBM Spectrum Archive, a member of the IBM Spectrum Storage™ family, enables direct, intuitive, and graphical access to data stored in IBM tape drives and libraries by incorporating the LTFS format standard for reading, writing, and exchanging descriptive metadata on formatted tape cartridges. IBM Spectrum Archive eliminates the need for additional tape management and software to access data.

IBM Spectrum Archive offers three software solutions for managing your digital files with the LTFS format: Single Drive Edition (SDE), Library Edition (LE), and Enterprise Edition (EE). This book focuses on the IBM Spectrum Archive EE.

IBM Spectrum Archive EE provides seamless integration of LTFS with IBM Spectrum Scale, which is another member of the IBM Spectrum Storage family, by creating a tape-based storage tier. You can run any application that is designed for disk files on tape by using IBM Spectrum Archive EE because it is fully transparent and integrates in the IBM Spectrum Scale file system. IBM Spectrum Archive EE can play a major role in reducing the cost of storage for data that does not need the access performance of primary disk.

With IBM Spectrum Archive EE, you can enable the use of LTFS for the policy management of tape as a storage tier in an IBM Spectrum Scale environment and use tape as a critical tier in the storage environment.

The use of IBM Spectrum Archive EE to replace online disk storage with tape in tier 2 and tier 3 storage can improve data access over other storage solutions because it improves efficiency and streamlines management for files on tape. IBM Spectrum Archive EE simplifies the use of physical tape by making it not apparent to the user and manageable by the administrator under a single infrastructure.

Figure 1-1 shows the integration of an IBM Spectrum Archive EE archive solution.

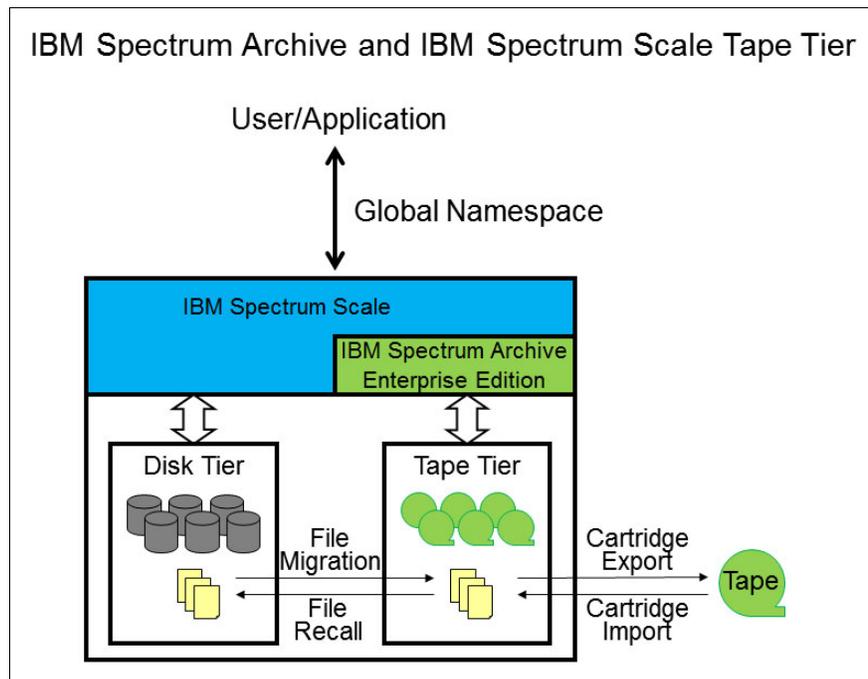


Figure 1-1 High-level overview of an IBM Spectrum Archive EE archive solution

IBM Spectrum Archive EE uses an enhanced version of the IBM Spectrum Archive LE, which is referred to as *the IBM Spectrum Archive LE+ component*, for the movement of files to and from tape devices. The scale-out architecture of IBM Spectrum Archive EE can add nodes and tape devices as needed to satisfy bandwidth requirements between IBM Spectrum Scale and the IBM Spectrum Archive EE tape tier.

Low-cost storage tier, data migration, and archive needs that are described in the following use cases can benefit from IBM Spectrum Archive EE:

- ▶ Operational storage
  - Provides a low-cost, scalable tape storage tier.
- ▶ Active archive
  - A local or remote IBM Spectrum Archive EE node serves as a migration target for IBM Spectrum Scale that transparently archives data to tape that is based on policies set by the user.

The following IBM Spectrum Archive EE characteristics cover a broad base of integrated storage management software with leading tape technology and the highly scalable IBM tape libraries:

- ▶ Integrates with IBM Spectrum Scale by supporting file-level migration and recall with an innovative database-less storage of metadata.
- ▶ Provides a scale-out architecture that supports multiple IBM Spectrum Archive EE nodes that share tape inventory with load balancing over multiple tape drives and nodes.
- ▶ Enables tape cartridge pooling and data exchange for IBM Spectrum Archive EE tape tier management:
  - Tape cartridge pooling allows the user to group data on sets of tape cartridges.
  - Multiple copies of files can be written on different tape cartridge pools, including different tape libraries in different locations.
  - Supports tape cartridge export with and without the removal of file metadata from IBM Spectrum Scale.
  - Supports tape cartridge import with pre-population of file metadata in IBM Spectrum Scale.

Furthermore, IBM Spectrum Archive EE provides the following key benefits:

- ▶ A low-cost storage tier in an IBM Spectrum Scale environment.
- ▶ An active archive or big data repository for long-term storage of data that requires file system access to that content.
- ▶ File-based storage in the LTFS tape format that is open, self-describing, portable, and interchangeable across platforms.
- ▶ Lowers capital expenditure and operational expenditure costs by using cost-effective and energy-efficient tape media without dependencies on external server hardware or software.
- ▶ Allows the retention of data on tape media for long-term preservation (10+ years).
- ▶ Provides the portability of large amounts of data by bulk transfer of tape cartridges between sites for disaster recovery and the initial synchronization of two IBM Spectrum Scale sites by using open-format, portable, self-describing tapes.
- ▶ Migration of data to newer tape or newer technology that is managed by IBM Spectrum Scale.

- ▶ Provides ease of management for operational and active archive storage.
- ▶ Expand archive capacity simply by adding and provisioning media without affecting the availability of data already in the pool.

**Tip:** For a no-cost trial version of the IBM Spectrum Archive EE, contact your local IBM sales representative.

### 1.1.1 Operational storage

This section describes how IBM Spectrum Archive EE is used as a storage tier in an IBM Spectrum Scale environment.

Using an IBM Spectrum Archive tape tier as operational storage is useful when a significant portion of files on a disk storage system infrastructure is static, meaning the data is not changing.

In this case, as shown in Figure 1-2, it is optimal to move the content to a lower-cost storage tier, in this case a physical tape. The files that are migrated to the IBM Spectrum Archive EE tape tier remain online, meaning they are accessible from the IBM Spectrum Scale file system under the IBM Spectrum Scale namespace at any time. Tape cartridge pools within IBM Spectrum Archive EE can also be used for backup.

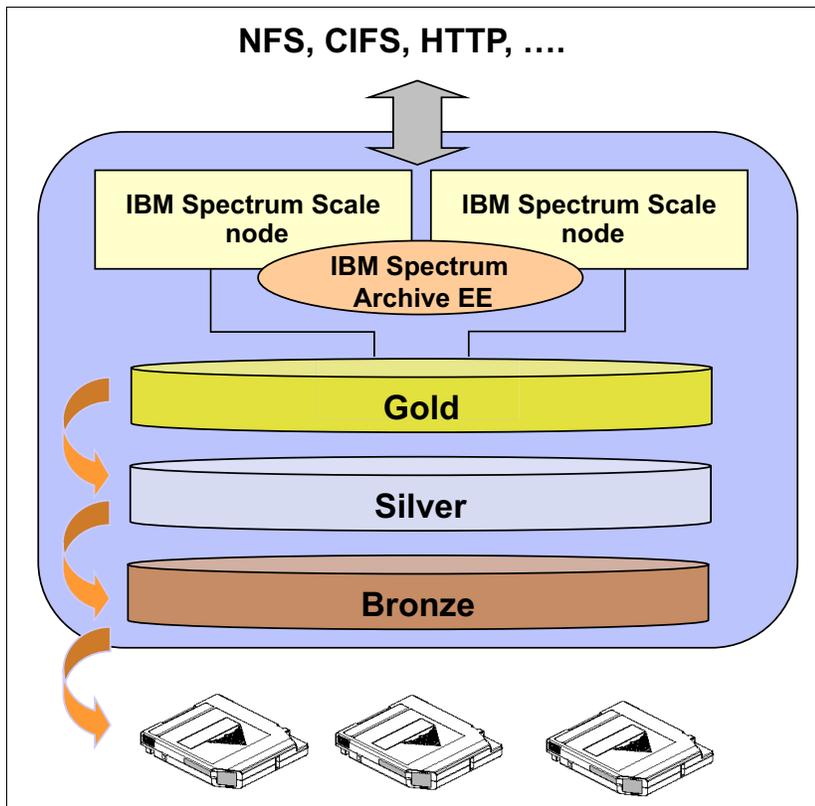


Figure 1-2 Tiered operational storage with IBM Spectrum Archive EE managing the tape tier

With IBM Spectrum Archive EE, the user specifies files to be migrated to the IBM Spectrum Archive tape tier by using standard IBM Spectrum Scale scan policies. IBM Spectrum Archive EE then manages the movement of IBM Spectrum Scale file data to the IBM Spectrum Archive tape cartridges. It also edits the metadata of the IBM Spectrum Scale files to point to the content on the IBM Spectrum Archive tape tier.

Access to the migrated files through the IBM Spectrum Scale file system remains unchanged, with the file data provided at the data rate and access times of the underlying tape technology. The IBM Spectrum Scale namespace is unchanged after migration, making the placement of files in the IBM Spectrum Archive tape tier not apparent to users and applications. See 8.10.1, “Creating a traditional archive system policy” on page 235.

### 1.1.2 Active archive

This section describes how IBM Spectrum Archive EE is used as an active archive in an IBM Spectrum Scale environment.

The use of an LTF5 tape tier as an active archive is useful when you need a low-cost, long-term archive for data that is maintained and accessed for reference. IBM Spectrum Archive satisfies the needs of this type of archiving by using open-format, portable, and self-describing tapes based on the LTF5 standard.

In an active archive, the IBM Spectrum Archive file system is the main store for the data while the IBM Spectrum Scale file system, with its limited disk capacity, is used as a staging area, or cache, in front of IBM Spectrum Archive EE. IBM Spectrum Scale policies are used to stage and de-stage data from the IBM Spectrum Scale disks to the IBM Spectrum Archive EE tape cartridge.

Figure 1-3 shows the archive storage management with the IBM Spectrum Archive tape tier in the IBM Spectrum Scale file system, the disk that is used for caching, and the namespace that is mapped to the tape cartridge pool.

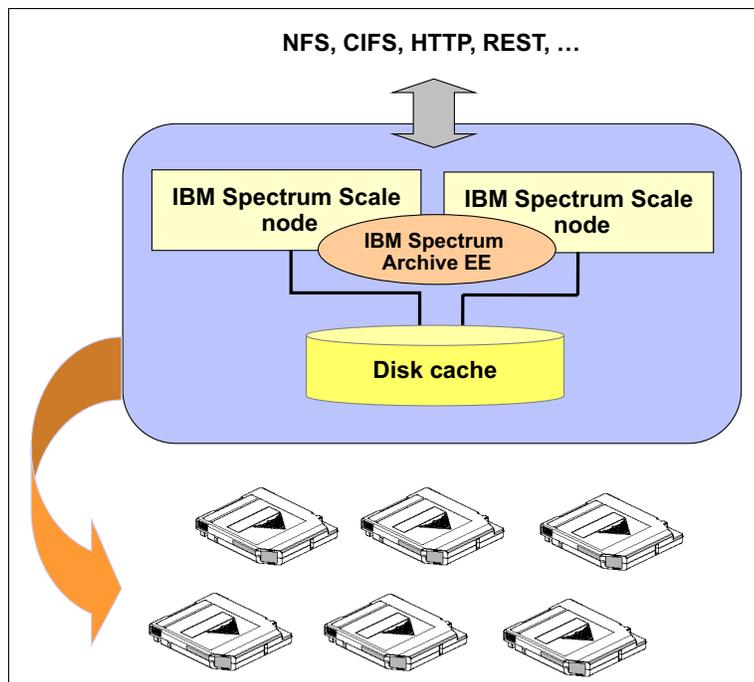


Figure 1-3 Archive storage management with IBM Spectrum Archive EE

The tapes from the archive can be exported for vaulting or for moving data to another location. Because the exported data is in the LTFS format, it can be read on any LTFS-compatible system.

For more information see 8.10.2, “Creating active archive system policies” on page 236.

## 1.2 IBM Spectrum Archive EE functions

This section describes the main functions that are found within IBM Spectrum Archive EE. Figure 1-4 shows where IBM Spectrum Archive EE fits within the solution architecture that integrates with IBM Spectrum Archive LE and IBM Spectrum Scale. This integration enables the functions of IBM Spectrum Archive to represent the external tape cartridge pool to IBM Spectrum Scale and file migration based on IBM Spectrum Scale policies. IBM Spectrum Archive EE can be configured on multiple nodes with those instances of IBM Spectrum Archive EE sharing a physical tape library.

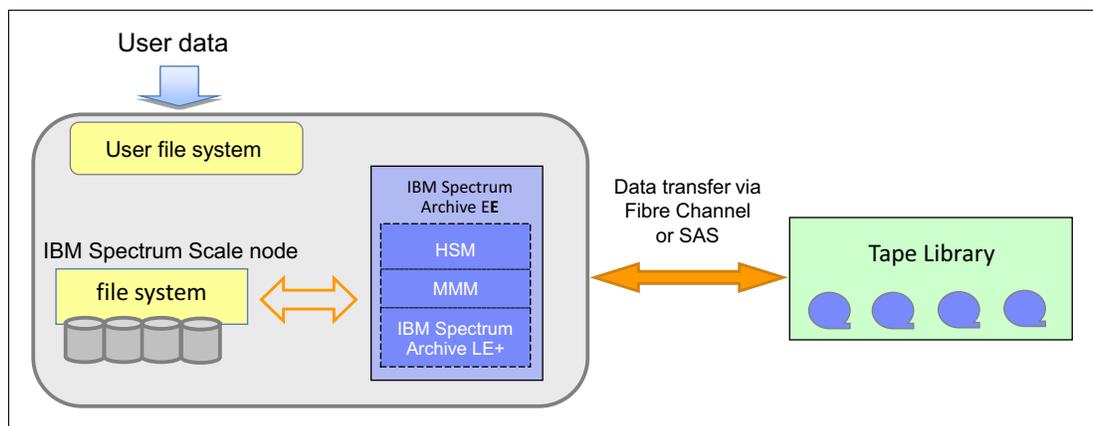


Figure 1-4 IBM Spectrum Archive EE integration with IBM Spectrum Scale and IBM Spectrum Archive LE

With IBM Spectrum Archive EE, you can perform the following management tasks on your system:

- ▶ Create and define tape cartridge pools for file migrations.
- ▶ Migrate files in the IBM Spectrum Scale namespace to the IBM Spectrum Archive tape tier.
- ▶ Recall files that were migrated to the IBM Spectrum Archive tape tier back into IBM Spectrum Scale.
- ▶ Reconcile file inconsistencies between files in IBM Spectrum Scale and their equivalents in IBM Spectrum Archive.
- ▶ Reclaim tape space that is occupied by non-referenced files and non-referenced content that is present on the physical tapes.
- ▶ Export tape cartridges to remove them from your IBM Spectrum Archive EE system.
- ▶ Import tape cartridges to add them to your IBM Spectrum Archive EE system.
- ▶ Add tape cartridges to your IBM Spectrum Archive EE system to expand the tape cartridge pool with no disruption to your system.
- ▶ Obtain inventory, job, and scan status of your IBM Spectrum Archive EE solution.

## 1.3 IBM Spectrum Archive EE components

This section describes the components that make up IBM Spectrum Archive EE:

- ▶ EE component (multi-tape management module (MMM))
- ▶ LE+ component
- ▶ HSM component

IBM Spectrum Scale is a required component for the IBM Spectrum Archive solution.

IBM Spectrum Archive EE is composed of multiple components that enable an IBM Spectrum Archive tape tier to be used for migration and recall with the IBM Spectrum Scale. Files are migrated to, and recalled from, the IBM Spectrum Archive tape tier by using the IBM Spectrum Archive EE components that are shown in Figure 1-5 on page 8 and Figure 1-6 on page 9.

### 1.3.1 IBM Spectrum Archive EE terms

This list highlights the components of an IBM Spectrum Archive EE solution:

<b>IBM Spectrum Archive EE node</b>	An x86_64 IBM Spectrum Scale server that is running on IBM Spectrum Archive EE. Each EE node must be connected to a set of tape drives in a tape library, through an FC connection. One EE node cannot be connected to more than one logical library.
<b>IBM Spectrum Archive EE Cluster</b>	<p>A set of EE nodes that are connected to a single IBM Spectrum Scale cluster. All nodes in a cluster can see the files on the IBM Spectrum Scale file system with same inode number.</p> <p>IBM Spectrum Scale clusters that are connected by active file management (AFM) are considered as two separate IBM Spectrum Scale clusters by IBM Spectrum Archive EE.</p>
<b>IBM Spectrum Scale Cluster</b>	IBM Spectrum Scale servers (non-EE nodes) and EE nodes.
<b>IBM Spectrum Scale only Node</b>	An IBM Spectrum Scale server that is running on a supported platform, such as Linux or Microsoft Windows, without IBM Spectrum Archive EE.
<b>Tape Pool</b>	<p>A set of tape cartridges of the same type (either Write Once Read Many (WORM) or Non-WORM, and either LTO or 3592) that are in one logical tape library. A tape pool uses the same generation of tapes within the pool.</p> <p>A tape pool does not span across multiple tape libraries.</p> <p>A tape pool is assigned to only one node group.</p>

**Node Group**

Nodes that are connected to the same tape library. Normally there is a one-to-one relationship between a node group and a tape library, so a dual-library EE cluster has two node groups, at minimum. In theory, you can divide one tape library into multiple node groups, just like partitioning.

Tape pools are assigned to only one node group, but a node group can access multiple tape pools.

**Control Node**

An EE node that is running in an MMM. IBM Spectrum Archive EE V1R2 and subsequent releases require you to configure one control node per tape library. The control node manages all the requests for access to its associated tape library. The control node redirects requests for access to other tape libraries to the control nodes of the other tape libraries.

Figure 1-5 shows the components that make up IBM Spectrum Archive EE. The components are shown with IBM Spectrum Scale configured on separate nodes for maximum scalability.

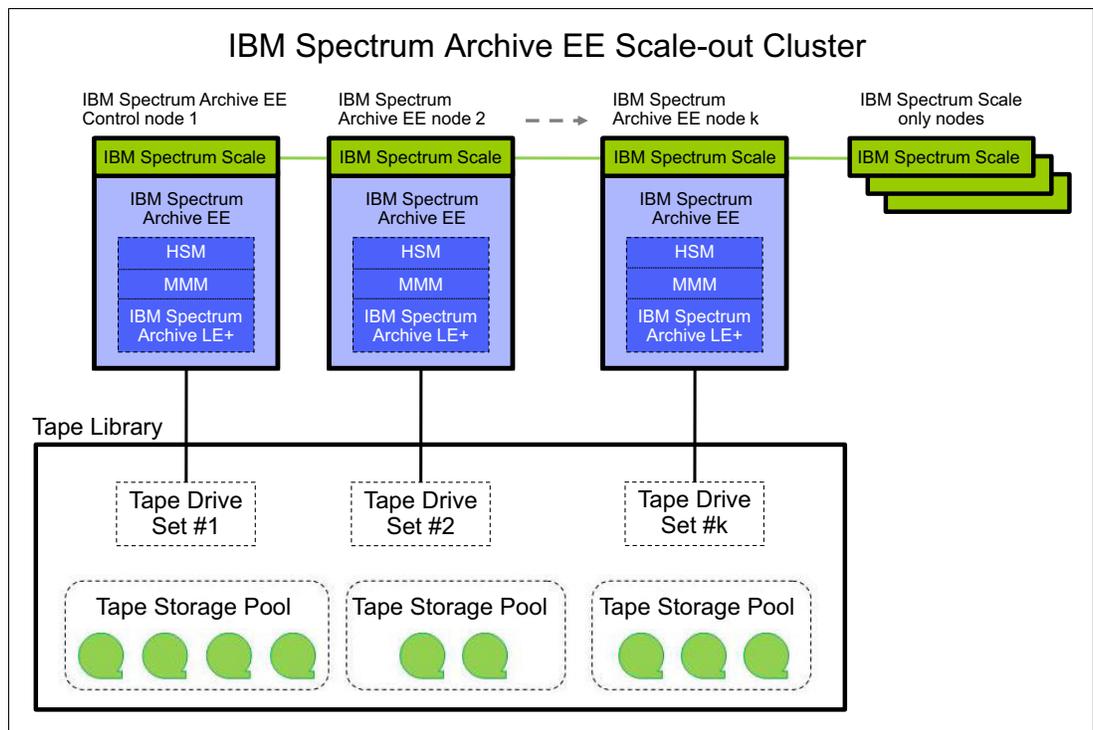


Figure 1-5 Components of IBM Spectrum Archive EE with separate IBM Spectrum Scale nodes

In Figure 1-6, the components that make up IBM Spectrum Archive EE are shown with no separate IBM Spectrum Scale nodes. This diagram shows how IBM Spectrum Scale can be configured to run on the same nodes as the IBM Spectrum Archive EE nodes.

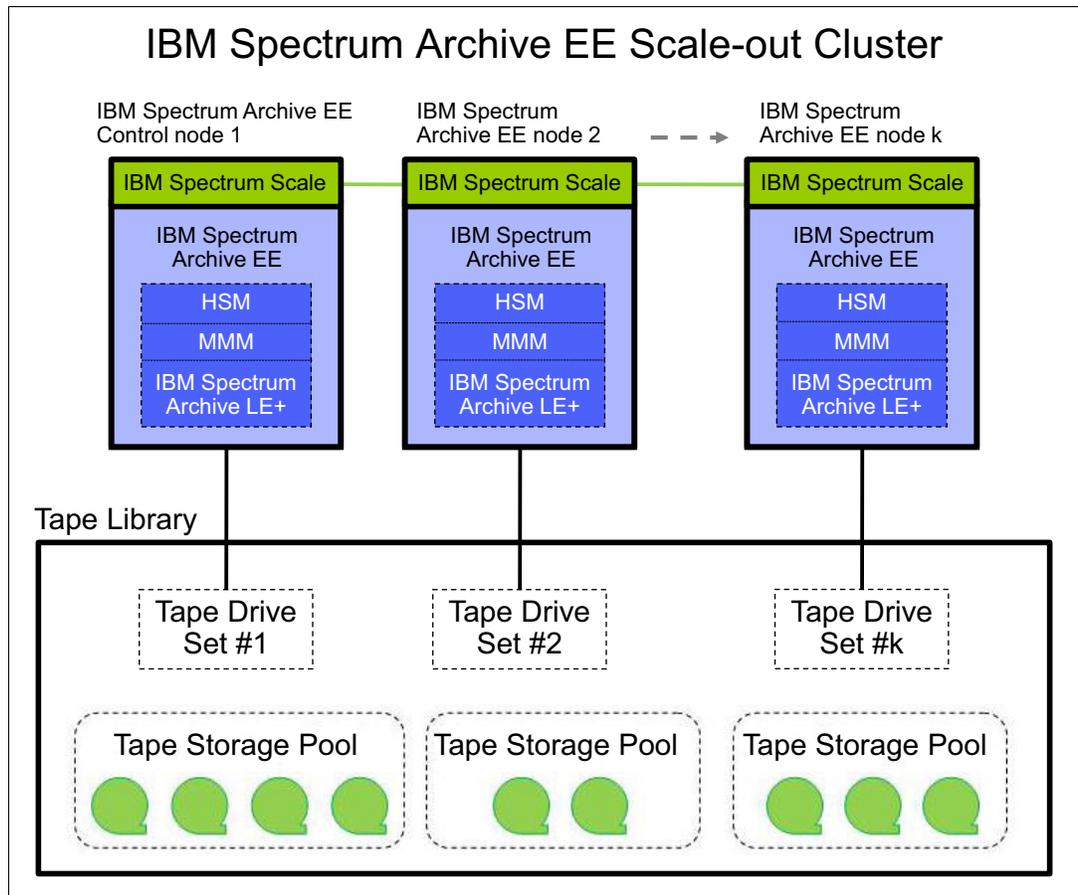


Figure 1-6 Components of IBM Spectrum Archive EE with no separate IBM Spectrum nodes

A second tape library can be added to the configuration, expanding the storage capacity and offering the opportunity to add even more nodes and tape devices. Availability can be improved by storing redundant copies on different tape libraries.

With multiple tape library attachments, the tape libraries can be connected to an IBM Spectrum Scale cluster in a single site or can be placed in the metro distance locations through IBM Spectrum Scale synchronous mirroring (stretched cluster).

The following distances are supported for stretched cluster with synchronous mirroring using block-level replication:

<http://www.ibm.com/support/knowledgecenter/STXKQY/gpfsc1ustersfaq.html>

- ▶ For V4.2.2.3 and later, distances greater than 300 km
- ▶ For V4.2.1 and later, distances up to 300 km
- ▶ For V4.2 and previous, distances less than 100 km

It is important to remember that this is still a single IBM Spectrum Scale cluster. In a configuration using IBM Spectrum Scale replication, a single IBM Spectrum Scale cluster is defined over two geographically separated sites consisting of two active production sites and by using tiebreaker disks.

One or more file systems are created, mounted, and accessed concurrently from the two active production sites. The data and metadata replication features of IBM Spectrum Scale are used to maintain a secondary copy of each file system block, relying on the concept of disk failure groups to control the physical placement of the individual copies:

1. Separate the set of available disk volumes into two failure groups. Define one failure group at each of the active production sites.
2. Create a replicated file system. Specify a replication factor of 2 for both data and metadata.

When allocating new file system blocks, IBM Spectrum Scale always assigns replicas of the same block to distinct failure groups. This feature provides a sufficient level of redundancy, allowing each site to continue operating independently should the other site fail.

For more information about synchronous mirroring that uses IBM Spectrum Scale replication, see the following website:

[https://www.ibm.com/support/knowledgecenter/en/STXKQY\\_4.2.3/com.ibm.spectrum.scale.v4r23.doc/blladv\\_continuous\\_replication\\_SSdata.htm](https://www.ibm.com/support/knowledgecenter/en/STXKQY_4.2.3/com.ibm.spectrum.scale.v4r23.doc/blladv_continuous_replication_SSdata.htm)

**Important:** Stretched cluster is available for distances as listed below. For longer distances, use the AFM feature of IBM Spectrum Scale with IBM Spectrum Archive. With the release of IBM Spectrum Archive V1.2.3.0, limited support is provided for IBM Spectrum Scale AFM. The only supported AFM mode of IBM Spectrum Scale is independent-writer (IW). The use of AFM will be with two different IBM Spectrum Scale clusters with one instance of IBM Spectrum Archive at each site. For more details about IBM Spectrum Scale AFM, see 2.3.5, “Active File Management” on page 29.

The following distances are supported for stretched cluster with synchronous mirroring using block-level replication:

<http://www.ibm.com/support/knowledgecenter/STXKQY/gpfclustersfaq.html>

- ▶ For V4.2.2.3 and later, distances greater than 300 km
- ▶ For V4.2.1 and later, with distances up to 300 km
- ▶ For V4.2 and previous, with distances less than 100 km

Figure 1-7 shows a fully configured IBM Spectrum Archive EE.

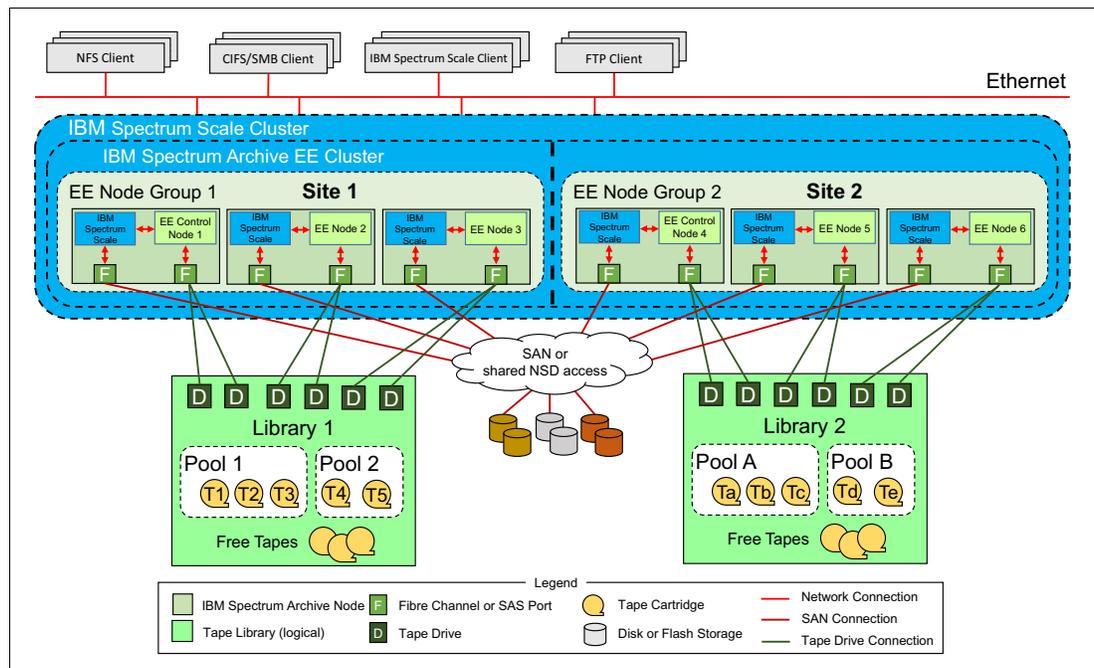


Figure 1-7 IBM Spectrum Archive EE shown with scaled-out capacity, multiple libraries, and redundant copies

### 1.3.2 Hierarchical Storage Manager

A Hierarchical Storage Manager (HSM) solution typically moves the file's data to back-end storage (in most cases, physical tape media) and leaves a small stub file in the local storage file system. The stub file uses minimal space, but leaves all metadata information about the local storage in such a way that for a user or a program the file looks like a normal, local stored file. When the user or a program accesses the file, the HSM solution automatically recalls (moves back) the file's data from the back-end storage and gives the reading application access to the file when all the data is back retrieved and available online again.

### 1.3.3 Multi-Tape Management Module

This component is a service of IBM Spectrum Archive EE control node. The MMM service implements policy-based tape cartridge selection and maintains the state of all of the resources that are available in the system. There is one control node per tape library.

The scheduler component of the control node uses policy-based cartridge selection to schedule and process job requests, such as migration and recall requests, which are fulfilled by using available system nodes and tape resources. The following tasks are done by the scheduler:

- ▶ Choosing a job from the job queue
- ▶ Choosing an appropriate tape cartridge and tape drive to handle the work
- ▶ Starting the job

The control node also manages the creation of replicas across multiple tape libraries. For example, when the `ltfsee migrate` command specifies making replicas of files in multiple tape libraries, the command accesses the control node that manages the tape pool for the primary copy.

The control node puts the copy job in the job queue for the primary tape library, then passes the secondary copy job to the control node for the second tape library. The second control node puts the copy job in the job queue for the second tape library.

When the scheduler component of the control node selects a tape cartridge and tape drive for a migration job, it manages the following conditions:

- ▶ If the migration is to a tape cartridge pool, the tape drive must belong to the node group that owns the tape cartridge pool.
- ▶ If a format generation property is defined for a tape cartridge pool, the tape cartridge must be formatted as that generation, and the tape drive must support that format.
- ▶ The number of tape drives that are being used for migration to a tape cartridge pool at one time must not exceed the defined mount limit.
- ▶ If there are multiple candidate tapes available for selection, the scheduler tries to choose a tape cartridge that is already mounted on an available tape drive.

When the scheduler selects a tape cartridge and tape drive for jobs other than migration, it makes the following choices:

- ▶ Choosing an available tape drive in the node group that owns the tape cartridge and tape cartridge pool
- ▶ Choosing the tape drive that has the tape drive attribute for the job

When the control node scheduler selects a tape cartridge for transparent recalls such as double-clicks or application reads, it manages the following conditions:

- ▶ If the file has a replica, the scheduler always chooses the primary copy. The first tape cartridge pool that is used by the migration process contains the primary copy.
- ▶ If the primary copy cannot be accessed, the scheduler automatically retries the recall job by using the other replicas if available.

Other functions that are provided by the control node include the following functions:

- ▶ Maintains a catalog of all known drives that are assigned to each IBM Spectrum Archive node in the system
- ▶ Maintains a catalog of tape cartridges in the tape library/libraries
- ▶ Maintains an estimate of the free space on each tape cartridge
- ▶ Allocates space on tape cartridges for new data

The MMM service is started when IBM Spectrum Archive EE is started by running the **ltfsee start** command. The MMM service runs on only one IBM Spectrum Archive EE control node at a time. Several operations, including migration and recall, fail if the MMM service stops. If SNMP traps are enabled, a notification is sent when the MMM service starts or stops.

For more information, see 7.4, “Starting and stopping IBM Spectrum Archive EE” on page 143, and 7.27, “Monitoring the system with SNMP” on page 212.

**Important:** If the **ltfsee start** command does not return after several minutes, it might be because the firewall is running. The firewall service must be disabled on the IBM Spectrum Archive EE nodes. For more information, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60.

The **ltfsee start** command also does the unmount of the tape drives, so the process might take a long time if there are many mounted tape drives.

### 1.3.4 IBM Spectrum Archive Library Edition Plus component

The IBM Spectrum Archive Library Edition Plus (LE+) component is the IBM Spectrum Archive tape tier of IBM Spectrum Archive EE. It is an enhanced version of the IBM Spectrum Archive LE that is designed to work with the EE.

The LE+ component is installed on all of the IBM Spectrum Scale nodes that are connected to the IBM Spectrum Archive EE library. It is the migration target for IBM Spectrum Scale. The LE+ component accesses the recording space on the physical tape cartridges through its file system interface and handles the user data as file objects and associated metadata in its namespace.

With IBM Spectrum Archive EE V1.2.4.0 and later, IBM Spectrum Archive LE+ is started automatically when running the `ltfsee start` command. If errors occur during start of the IBM Spectrum Archive EE system, run the `ltfsee info nodes` command to display which component failed to start. For more information about the updated `ltfsee info nodes` command, see 7.6, “IBM Spectrum Archive EE automatic node failover” on page 145.

## 1.4 IBM Spectrum Archive EE cluster configuration introduction

This section describes a cluster configuration for IBM Spectrum Archive EE. This configuration is for single-library, multiple-node access.

Single-library, multiple-node access enables access to the same set of IBM Spectrum Archive EE tape cartridges from more than one IBM Spectrum Archive EE node. The purpose of enabling this capability is to improve data storage and retrieval performance by assigning fewer tape drives to each node.

When this cluster configuration is used, each IBM Spectrum Archive EE node must have its own set of drives that is not shared with any other node. In addition, each IBM Spectrum Archive EE node must have at least one control path drive that is designated as a control path by an operator of the attached IBM tape library.

IBM Spectrum Archive EE uses the drive that is designated as a control path to communicate with the tape library. This type of control path is also known as a *media changer device*. IBM Spectrum Archive EE is scalable so you can start out with a single node and add more nodes later.

**Important:** As part of your planning, work with your IBM tape library administrator to ensure that each IBM Spectrum Archive EE node in your configuration has its own media changer device (control path) defined in its logical library.

Figure 1-8 shows the typical setup for an IBM Spectrum Archive EE single-library, multiple-node access.

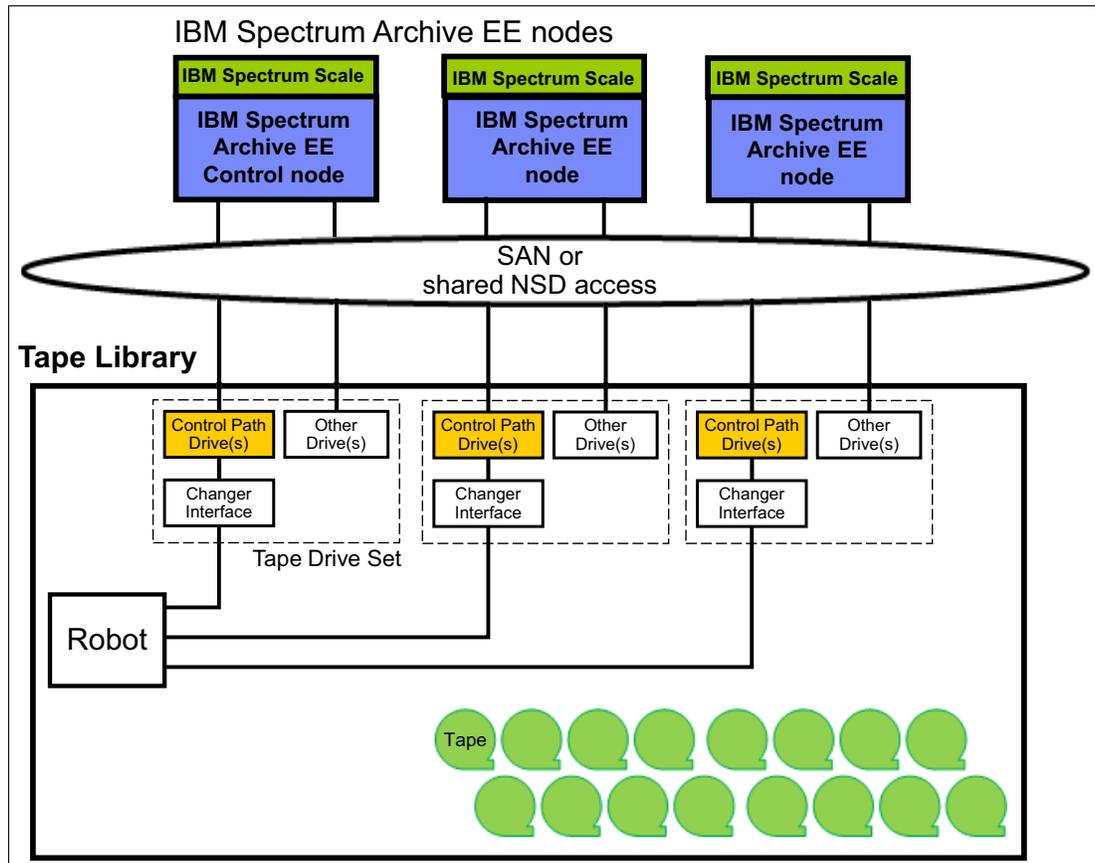


Figure 1-8 Single-library multiple-node access setup

IBM Spectrum Archive EE manages all aspects of the single-library, multiple-node access, which includes the management of the following areas:

- ▶ Multiple tenancy

The contents of the tape cartridge are managed automatically by the IBM Spectrum Archive EE system so that each IBM Spectrum Archive EE node does not have to be aware of any changes made on other IBM Spectrum Archive EE nodes. The index on each tape cartridge is updated when the tape is mounted and the index is read from this tape.

- ▶ Single node management of library inventory

The IBM Spectrum Archive EE system automatically keeps the library inventory up to date to manage the available drives and tape cartridges. The library inventory is kept on the node on which the MMM service runs.

- ▶ Space reclaim management

When data is moved from one tape cartridge to another to reclaim the space on the first tape cartridge, the IBM Spectrum Archive EE system ensures that the internal database reflects the change in the index of the IBM Spectrum Archive physical tape cartridge.



# IBM Spectrum Archive overview

This chapter provides an overview of the IBM Spectrum Archive product family and the individual components of the IBM Spectrum Archive Enterprise Edition (EE).

This chapter includes the following topics:

- ▶ Introduction to IBM Spectrum Archive and LTFS
- ▶ Comparison of the IBM Spectrum Archive products
- ▶ IBM Spectrum Scale
- ▶ OpenStack SwiftHLM
- ▶ IBM Spectrum Archive EE dashboard
- ▶ IBM Spectrum Archive EE REST API
- ▶ Types of archiving

## 2.1 Introduction to IBM Spectrum Archive and LTFS

LTFS is the first file system that works with LTO tape technology and IBM Enterprise tape drives, providing ease of use and portability for open systems tape storage.

**Note:** Throughout this publication, the terms *supported tape libraries*, *supported tape drives*, and *supported tape media* are used to represent the following tape libraries, tape drives, and tape media. Unless otherwise noted, as of the date of publication, IBM Spectrum Archive EE supports these libraries:

- ▶ IBM TS4500 tape library<sup>a</sup>
- ▶ IBM TS4300 tape library
- ▶ IBM TS3500 tape library<sup>a</sup>
- ▶ IBM TS3310 tape library
- ▶ IBM LTO Ultrium 8, 7, 6, or 5 tape drives, and IBM TS1155, TS1150, or TS1140 tape drives
- ▶ LTO Ultrium 8, M8<sup>b</sup>, 7, 6, and 5, and 3592 JC, JB, JD, JK, and JL tape media

Refer to the following website to check for the latest system requirements:

[https://www.ibm.com/support/knowledgecenter/en/ST9MBR\\_1.2.6/ltfs\\_ee\\_system\\_reqs.html](https://www.ibm.com/support/knowledgecenter/en/ST9MBR_1.2.6/ltfs_ee_system_reqs.html)

For the latest IBM Spectrum Archive Library Edition Support Matrix (supported tape library and tape drive firmware levels, check Fix Central:

<http://www.ibm.com/support/fixcentral>

- a. IBM TS1155, TS1150, and IBM TS1140 support on IBM TS4500 and IBM TS3500 tape libraries only
- b. Uninitialized M8 media (MTM 3589-452) is only supported on the following tape libraries in IBM Spectrum Archive EE: TS4500, TS4300, and TS3310. TS3500 will only support pre-initialized media.

With this application, accessing data that is stored on an IBM tape cartridge is as easy and intuitive as the use of a USB flash drive. Tapes are self-describing, and you can quickly recall any file from a tape cartridge without having to read the whole tape cartridge from beginning to end. Furthermore, any LTFS-capable system can read a tape cartridge that is created by any other LTFS-capable system (regardless of the operating system). Any LTFS-capable system can identify and retrieve the files that are stored on it. LTFS-capable systems have the following characteristics:

- ▶ Files and directories are shown to you as a directory tree listing.
- ▶ More intuitive searches of tape cartridges and library content are now possible because of the addition of file tagging.
- ▶ Files can be moved to and from LTFS tape cartridges by using the familiar drag method that is common to many operating systems.
- ▶ Many applications that were written to use files on disk can now use files on tape cartridges without any modification.
- ▶ All standard File Open, Write, Read, Append, Delete, and Close functions are supported.
- ▶ No need for an additional, external tape management system or database tracking the content of each tape.

Archival data storage requirements are growing at over 60% annually. The LTFS format is an ideal option for long-term archiving of large files that must be easily shared with others. This option is especially important because the tape media that it uses (LTO and 3592) are designed to have a 15 - 30 year lifespan (depending on the number of read/write passes).

Industries that most benefit from this tape file system are the banking, digital media, medical, geophysical, and entertainment industries. Many users in these industries use Linux or Macintosh systems, which are fully compatible with LTFS.

LTO Ultrium tape cartridges from earlier LTO generations (that is, LTO-1 through LTO-4) *cannot* be partitioned and be used by LTFS/IBM Spectrum Archive. Also, if LTO Ultrium 4 tape cartridges are used in an LTO Ultrium 5 tape drive to write data, the LTO-4 tape cartridge is treated as an unpartitioned LTO-5 tape cartridge. Even if an application can manage partitions, it is not possible to partition the LTO-4 media that is mounted in an LTO Ultrium 5 drive.

Starting with the release of IBM Spectrum Archive EE V1R2, corresponding Write Once, Read Many (WORM) tape cartridges are supported in an IBM Spectrum Archive EE solution operating supported IBM Enterprise tape drives. With the same release, tape drives in mixed configurations are supported. For more information, see 7.23, "Tape drive intermix support" on page 204.

Although LTFS presents the tape cartridge as a disk drive, the underlying hardware is still a tape cartridge and is therefore sequential in nature. Tape does not allow random access. Data is always appended to the tape, and there is no overwriting of files. File deletions do not erase the data from tape, but instead erase the pointers to the data. So, although with LTFS you can simultaneously copy two (or more) files to an LTFS tape cartridge, you get better performance if you copy files sequentially.

To operate the tape file system, the following components are needed:

- ▶ Software in the form of an open source LTFS package
- ▶ Data structures that are created by LTFS on tape

Together, these components can manage a file system on the tape media as though it is a disk file system for accessing tape files, including the tape directory tree structures. The metadata of each tape cartridge, after it is mounted, is cached in server memory. Therefore, metadata operations, such as browsing the directory or searching for a file name, do not require any tape movement and are quick.

## 2.1.1 Tape media capacity with IBM Spectrum Archive

Table 2-1 lists the tape drives and media that are supported by LTFS. The table also gives the native capacity of supported media, and raw capacity of the LTFS data partition on the media.

Table 2-1 Tape media capacity with IBM Spectrum Archive

Tape drive	Tape media <sup>a</sup>	Native capacity <sup>b, c</sup>	LTFS data partition size <sup>cd</sup>
IBM TS1155 tape drive <sup>e, f</sup>	Advanced type C data (JC)	7000 GB (6519 GiB)	6757 GB (6293 GiB)
	Advanced type C WORM data (JY)	7000 GB (6519 GiB)	6757 GB (6293 GiB)
	Extended type D data (JD)	15000 GB (13969 GiB)	14562 GB (13562 GiB)
	Extended type D WORM data (JZ)	15000 GB (14969 GiB)	14562 GB (13562 GiB)
	Advanced type C economy data (JK)	900 GB (838 GiB)	869 GB (809 GiB)
	Advanced type D economy data (JL)	3000 GB (1862 GiB)	2912 GB (1804 GiB)
IBM TS1150 tape drive <sup>g, h</sup>	Advanced type C data (JC)	7000 GB (6519 GiB)	6757 GB (6293 GiB)
	Advanced type C WORM data (JY)	7000 GB (6519 GiB)	6757 GB (6293 GiB)
	Extended type D data (JD)	10,000 GB (9313 GiB)	9687 GB (9022 GiB)
	Extended type D WORM data (JZ)	10,000 GB (9313 GiB)	9687 GB (9022 GiB)
	Advanced type C economy data (JK)	900 GB (838 GiB)	869 GB (809 GiB)
	Advanced type D economy data (JL)	3000 GB (2794 GiB)	2912 GB (2712 GiB)
IBM TS1140 tape drive <sup>g, h</sup>	Advanced data type C (JC)	4000 GB (3725 GiB)	3650 GB (3399 GiB)
	Advanced data type C WORM (JY)	4000 GB (3725 GiB)	3650 GB (3399 GiB)
	Extended data (JB)	1600 GB (1490 GiB)	1457 GB (1357 GiB)
	Advanced type C economy data (JK)	500 GB (465 GiB)	456 GB (425 GiB)
IBM LTO Ultrium 8 tape drive	LTO 8	12000 GB (11175 GiB)	11711 GB (10907 GiB)
	LTO 8 (M8)	9000 GB (8382 Gib)	8731 GB (8132 GiB)
IBM LTO Ultrium 7 tape drive	LTO 7	6000 GB (5588 GiB)	5731 GB (5338 GiB)

Tape drive	Tape media <sup>a</sup>	Native capacity <sup>b, c</sup>	LTFS data partition size <sup>cd</sup>
IBM LTO Ultrium 6 tape drive	LTO 6	2500 GB (2328 GiB)	2408 GB (2242 GiB)
IBM LTO Ultrium 5 tape drive	LTO 5	1500 GB (1396 GiB)	1425 GB (1327 GiB)

- a. WORM media are not supported by IBM Spectrum Archive SDE and IBM Spectrum Archive LE, only with EE.
- b. The actual usable capacity is greater when compression is used.
- c. See the topic *Data storage values*, found at [https://www.ibm.com/support/knowledgecenter/en/STQNYL\\_2.4.0/lufs\\_data\\_storage\\_values.html](https://www.ibm.com/support/knowledgecenter/en/STQNYL_2.4.0/lufs_data_storage_values.html).
- d. Values that are given are the default size of the LTFS data partition, unless otherwise indicated.
- e. TS1155, TS1150, and TS1140 tape drives support enhanced partitioning for cartridges.
- f. Media that are formatted on a 3592 drive must be read on the same generation of drive. For example, a JC cartridge that was formatted by a TS1150 tape drive cannot be read on a TS1140 tape drive.
- g. TS1155, TS1150, and TS1140 tape drives support enhanced partitioning for cartridges.
- h. Media that are formatted on a 3592 drive must be read on the same generation of drive. For example, a JC cartridge that was formatted by a TS1150 tape drive cannot be read on a TS1140 tape drive.

## 2.2 Comparison of the IBM Spectrum Archive products

The following sections give a brief overview of the IBM Spectrum Archive software products that are available at the time of writing. Their main features are summarized in Table 2-2.

**Note:** IBM LTFS Storage Manager (LTFS SM) was discontinued from marketing effective 12/14/2015. IBM support for the LTFS SM still continues.

Table 2-2 LTFS product comparison

Name	License required	Market	Tape Drive Support	Tape Library support	Internal database	Integrates with IBM Spectrum Scale
IBM Spectrum Archive Single Drive Edition (SDE)	No	Entry - Midrange	Yes	No	No	No
IBM Spectrum Archive Library Edition (LE)	No	Midrange - Enterprise	Yes	Yes	No	No
IBM Spectrum Archive Enterprise Edition (EE)	Yes	Enterprise	Yes	Yes	No	Yes

## 2.2.1 IBM Spectrum Archive Single Drive Edition

The IBM Spectrum Archive SDE provides direct, intuitive, and graphical access to data that is stored with the supported IBM tape drives and libraries that use the supported Linear Tape-Open (LTO) Ultrium tape cartridges and IBM Enterprise tape cartridges. It eliminates the need for more tape management and software to access data. The LTFS format is the first file system that works with tape technology that provides ease of use and portability for open systems tape storage. With this system, accessing data that is stored on an IBM tape cartridge is as easy and intuitive as using a USB flash drive.

**Note:** For in-depth information about IBM Spectrum Archive Single Drive Edition, see *IBM Spectrum Archive Single Drive Edition and Library Edition: Installation and Configuration Guide*, SG24-8090.

Figure 2-1 shows the IBM Spectrum Archive SDE user view, which resembles standard file folders.

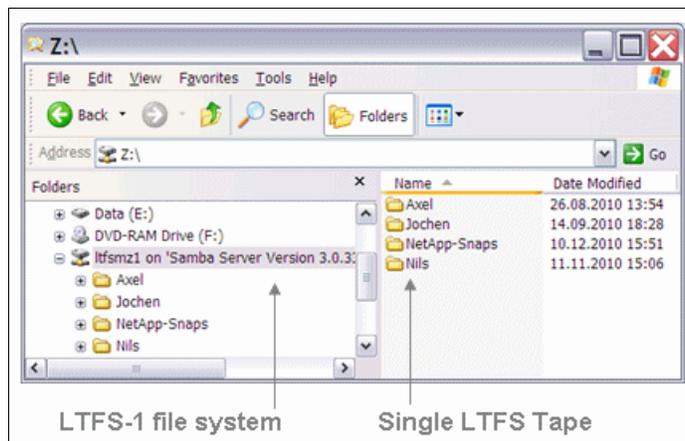


Figure 2-1 IBM Spectrum Archive SDE user view showing the file folders from the single LTFS tape cartridge

It runs on Linux, Windows, and MacOS, and with the operating system's graphical File Manager, reading data on a tape cartridge is as easy as dragging file data sets. Users can run any application that is designed for disk files against tape data without concern for the fact that the data is physically stored on tape. IBM Spectrum Archive SDE allows access to all of the data in a tape cartridge that is loaded on a single drive as though it were on an attached disk drive.

It supports existing stand-alone versions of LTFS, such as those running on IBM, HP, Quantum, FOR-A, 1 Beyond, and other platforms.

### IBM Spectrum Archive SDE software, systems, tape drives and media requirements

The most current software, systems, tape drives and media requirements can be found at the IBM Spectrum Archive Single Drive Edition IBM Knowledge Center website:

<https://www.ibm.com/support/knowledgecenter/en/STQNYL>

Select the most current IBM Spectrum Archive SDE version and then select **Planning**. The Supported tape drives and media, system requirements, and required software topics are displayed.

IBM Spectrum Archive SDE supports the use of multiple tape drives at one time. The method for using multiple tape drives depends on the operating system being used.

For Linux and Mac OS X users, it is possible to use multiple tape drives by starting multiple instances of the LTFS software, each with a different target tape device name in the `-o devname` parameter. For more information, see the topic *Mounting media by using the ltfs command*, found at:

[https://www.ibm.com/support/knowledgecenter/en/STQNYL\\_2.4.0/ltfs\\_managing\\_command\\_line\\_ltfs.html](https://www.ibm.com/support/knowledgecenter/en/STQNYL_2.4.0/ltfs_managing_command_line_ltfs.html)

For Windows users, the LTFS software detects each of the installed tape drives, and it is possible to assign a different drive letter to each drive by using the configuration window. For more information, see the topic *Assigning a drive letter to a tape drive*, found at:

[https://www.ibm.com/support/knowledgecenter/en/STQNYL\\_2.4.0/ltfs\\_assigning\\_drive\\_letter.html](https://www.ibm.com/support/knowledgecenter/en/STQNYL_2.4.0/ltfs_assigning_drive_letter.html)

**Note:** A certain level of tape drive firmware is required to fully use IBM Spectrum Archive SDE functions. To find the supported firmware version and for more information about connectivity and configurations, see the IBM System Storage® Interoperation Center (SSIC) at the following website:

<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

## Migration path to IBM Spectrum Archive EE

There is no direct migration path from IBM Spectrum Archive SDE to IBM Spectrum Archive EE software. Any existing IBM Spectrum Archive SDE software should be uninstalled before IBM Spectrum Archive EE is installed. Follow the uninstallation procedure that is documented in the *IBM Spectrum Archive Single Drive Edition and Library Edition: Installation and Configuration Guide*, SG24-8090.

Data tapes that are used by IBM Spectrum Archive SDE version 1.3.0 or later can be imported into IBM Spectrum Archive EE. For more information about this procedure, see 7.21.1, “Importing” on page 196. Tapes that were formatted in LTFS 1.0 format by older versions of IBM Spectrum Archive are automatically upgraded to LTFS 2.0 format on first write.

## 2.2.2 IBM Spectrum Archive Library Edition

IBM Spectrum Archive LE uses the open, non-proprietary LTFS format that allows any application to write files into a large archive. It provides direct, intuitive, and graphical access to data that is stored on tape cartridges within the supported IBM tape libraries that use either LTO IBM Enterprise supported tape drives.

**Note:** For in-depth information about IBM Spectrum Archive Library Edition see the *IBM Spectrum Archive Single Drive Edition and Library Edition: Installation and Configuration Guide*, SG24-8090.

Figure 2-2 shows the user view of multiple IBM Spectrum Archive tapes appearing as different library folders.

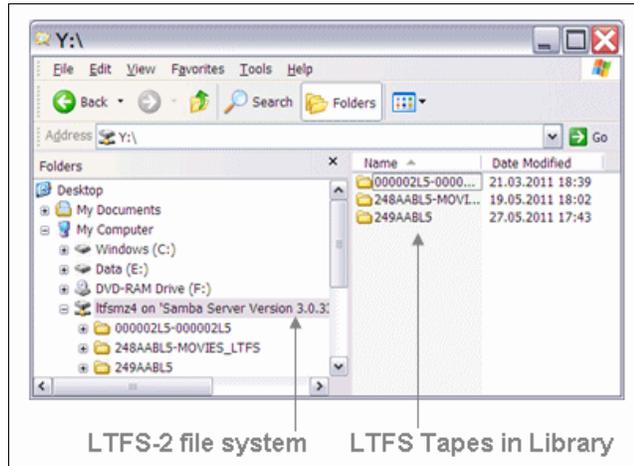


Figure 2-2 IBM Spectrum Archive LE User view of multiple LTFS tape cartridges

In addition, IBM Spectrum Archive LE enables users to create a single file system mount point for a logical library that is managed by a single instance of IBM Spectrum Archive, which runs on a single computer system.

The LTFS metadata of each tape cartridge, after it is mounted, is cached in server memory. So, even after the tape cartridge is ejected, the tape cartridge metadata information remains viewable and searchable, with no remounting required. Every tape cartridge and file is accessible through the operating system file system commands, from any application. This improvement in search efficiency can be substantial, considering the need to search hundreds or thousands of tape cartridges that are typically found in tape libraries.

## IBM Spectrum Archive LE software, systems, tape drives and media requirements

The most current software, systems, tape drives and media requirements can be found at the IBM Spectrum Archive Library Edition IBM Knowledge Center website:

<https://www.ibm.com/support/knowledgecenter/en/STZMZN>

Select the most current IBM Spectrum Archive LE version and then select **Planning**. The Supported tape drives and media, system requirements, and required software topics will be displayed.

For more information about connectivity and configurations, see the SSIC website:

<http://www.ibm.com/systems/support/storage/ssic/interoperability.wss>

## Migration path to IBM Spectrum Archive EE

There is no direct migration path from IBM Spectrum Archive LE to IBM Spectrum Archive EE software. Any existing IBM Spectrum Archive LE software should be uninstalled before IBM Spectrum Archive EE is installed. Follow the uninstall procedure that is documented in *IBM Spectrum Archive Single Drive Edition and Library Edition: Installation and Configuration Guide*, SG24-8090.

Data tapes that were created and used by IBM Spectrum Archive LE Version 2.1.2 or later can be imported into IBM Spectrum Archive EE. For more information about this procedure, see 7.21.1, “Importing” on page 196.

### 2.2.3 IBM Spectrum Archive Enterprise Edition

As enterprise-scale data storage, archiving, and backup expands, there is a need to lower storage costs and improve manageability. IBM Spectrum Archive EE provides such a solution that offers IBM Spectrum Scale users a new low-cost, scalable storage tier.

IBM Spectrum Archive EE provides seamless integration of LTFS with IBM Spectrum Scale by providing an IBM Spectrum Archive tape tier under IBM Spectrum Scale. IBM Spectrum Scale policies are used to move files between online disks storage and IBM Spectrum Archive tape tiers without affecting the IBM Spectrum Scale namespace.

IBM Spectrum Archive EE uses IBM Spectrum Archive LE for the movement of files to and from the physical tape devices and cartridges. IBM Spectrum Archive EE can manage multiple IBM Spectrum Archive LE nodes in parallel, so bandwidth requirements between IBM Spectrum Scale and the tape tier can be satisfied by adding nodes and tape devices as needed.

Figure 2-3 shows the IBM Spectrum Archive EE system view with IBM Spectrum Scale providing the global namespace and IBM Spectrum Archive EE installed on two IBM Spectrum Scale nodes. IBM Spectrum Archive EE can be installed on one or more IBM Spectrum Scale nodes. Each IBM Spectrum Archive EE instance has dedicated tape drives that are attached in the same tape library partition. IBM Spectrum Archive EE instances share tape cartridges and LTFS index. The workload is distributed over all IBM Spectrum Archive EE nodes and their attached tape drives.

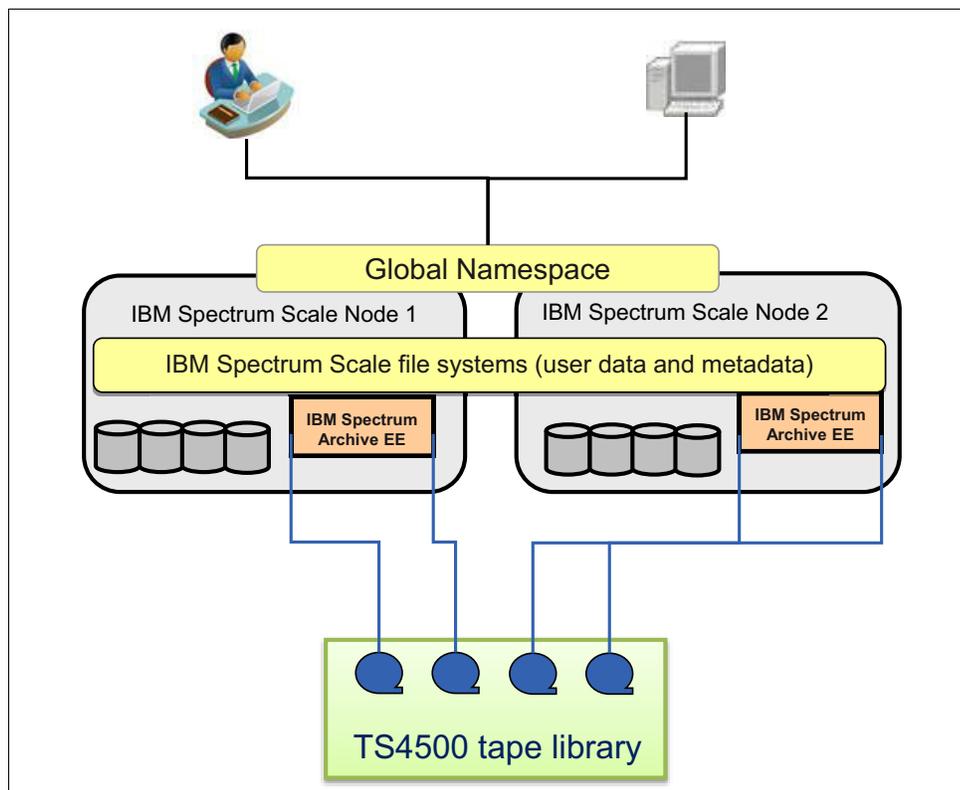


Figure 2-3 IBM Spectrum Archive EE system view

A local or remote IBM Spectrum Archive LE node serves as a migration target for IBM Spectrum Scale, which transparently archives data to tape based on policies set by the user.

IBM Spectrum Archive EE provides the following benefits:

- ▶ A low-cost storage tier in an IBM Spectrum Scale environment.
- ▶ An active archive or big data repository for long-term storage of data that requires file system access to that content.
- ▶ File-based storage in the LTFS tape format that is open, self-describing, portable, and interchangeable across platforms.
- ▶ Lowers capital expenditure and operational expenditure costs by using cost-effective and energy-efficient tape media without dependencies on external server hardware or software.
- ▶ Provides unlimited capacity scalability for the IBM supported tape libraries and keeping offline tape cartridges on shelves.
- ▶ Allows the retention of data on tape media for long-term preservation (10+ years).
- ▶ Provides the portability of large amounts of data by bulk transfer of tape cartridges between sites for disaster recovery and the initial synchronization of two IBM Spectrum Scale sites by using open-format, portable, self-describing tapes.
- ▶ Provides ease of management for operational and active archive storage.

Figure 2-4 provides a conceptual overview of processes and data flow in IBM Spectrum Archive EE.

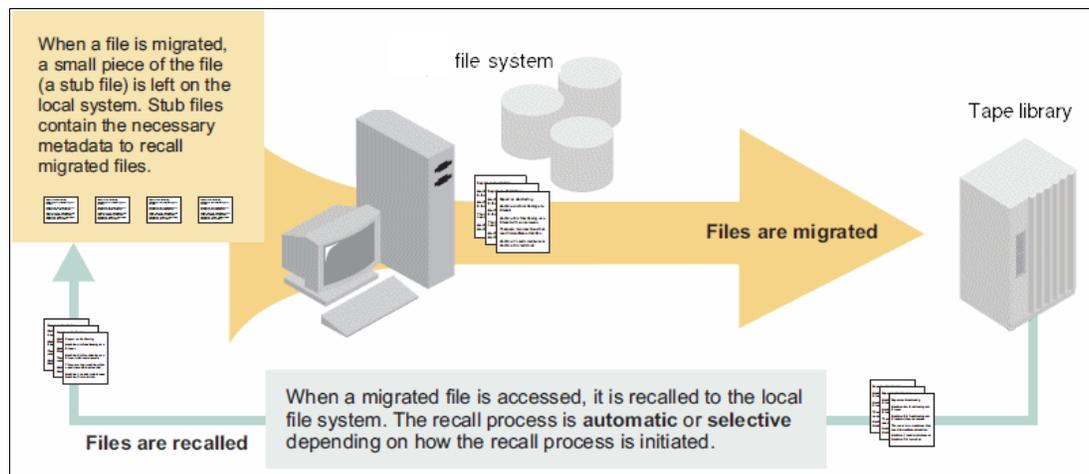


Figure 2-4 IBM Spectrum Archive EE data flow

IBM Spectrum Archive EE can be used for a low-cost storage tier, data migration, and archive needs as described in the following use cases.

### Operational storage

The use of an IBM Spectrum Archive tape tier as operational storage is useful when a significant portion of files on an online disk storage system is static, meaning the data does not change. In this case, it is more efficient to move the content to a lower-cost storage tier, for example, to a physical tape cartridge. The files that are migrated to the IBM Spectrum Archive tape tier remain online, meaning they are accessible at any time from IBM Spectrum Scale under the IBM Spectrum Scale namespace.

With IBM Spectrum Archive EE, the user specifies files to be migrated to the IBM Spectrum Archive tape tier by using standard IBM Spectrum Scale scan policies. IBM Spectrum Archive EE then manages the movement of IBM Spectrum Scale file data to IBM Spectrum Archive tape cartridges. It also edits the metadata of the IBM Spectrum Scale files to point to the content on the IBM Spectrum Archive tape tier.

Access to the migrated files through the IBM Spectrum Scale file system remains unchanged with the file data provided at the data rate and access times of the underlying tape technology. The IBM Spectrum Scale namespace is unchanged after migration, which makes the placement of files in the IBM Spectrum Archive tape tier not apparent to users and applications.

### **Active archive**

The use of an IBM Spectrum Archive tape tier as an active archive is useful when there is a need for a low-cost, long-term archive for data that is maintained and accessed for reference. IBM Spectrum Archive satisfies the needs of this type of archiving by using open-format, portable, self-describing tapes. In an active archive, the LTFS-file system is the main store for the data. The IBM Spectrum Scale file system, with its limited disk capacity, is used as a staging area, or cache, in front of IBM Spectrum Archive.

IBM Spectrum Scale policies are used to stage and destage data from the IBM Spectrum Scale disks space to the IBM Spectrum Archive tape cartridges. The tape cartridges from the archive can be exported for vaulting or for moving data to another location. Because the exported data is in the LTFS format, it can be read on any LTFS-compatible system.

### **History of the IBM Spectrum Archive EE releases**

The following list shows you the evolution and releases of the IBM Spectrum Archive EE software since going to market in the year 2013:

- ▶ Version 1.R1.0.0 (GA in 6/2013)
- ▶ Version 1.R1.1.0 (1/2014) as pGA 1
- ▶ Version 1.R1.1.1 (6/2014) as pGA 2
- ▶ Version 1.R1.1.2 (12/2014) as pGA 2.1
- ▶ Version 1.R1.1.3 (3/2015) as pGA 2.2
- ▶ Version 1.R1.2.0 (5/2015) for the IBM Spectrum Archive renaming
- ▶ Version 1.R2.0.0 (12/2015)
- ▶ Version 1.R2.1.0 (6/2016)
- ▶ Version 1.R2.2.0 (12/2016)
- ▶ Version 1.R2.3.0 (3/2017)
- ▶ Version 1.R2.4.0 (6/2017)
- ▶ Version 1.R2.5.0 (11/2017)
- ▶ Version 1.R2.6.0 (3/2018)

**Note:** pGA stands for post-GA releases of IBM software or products. The pGA releases usually contain a set of selected fixes. In this case, there are enhancements and some new product function and features.

## **2.3 IBM Spectrum Scale**

IBM Spectrum Scale is a cluster file system solution, which means that it provides concurrent access to one or more file systems from multiple nodes. These nodes can all be SAN-attached, network-attached, or both, which enables high-performance access to this common set of data to support a scale-out solution or provide a high availability platform.

The entire file system is striped across all storage devices, typically disk and flash storage subsystems.

**Note:** IBM Spectrum Scale is a member of the IBM Spectrum product family. The most recent version at the time of writing is Version 4.2.3. Version 4.1.1 was the first version with the name IBM Spectrum Scale. The prior versions, Versions 4.1.0 and 3.5.x, are still called GPFS. Both IBM Spectrum Scale and GPFS are used interchangeably in this book.

You can find the current documentation and publications for GPFS version at the following website:

[https://www.ibm.com/support/knowledgecenter/SSFKCN/gpfs\\_welcome.html](https://www.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html)

You can find the current documentation and publications for IBM Spectrum Scale at the following website:

[https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale\\_welcome.html](https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale_welcome.html)

## 2.3.1 Overview

IBM Spectrum Scale can help you achieve Information Lifecycle Management (ILM) efficiencies through powerful policy-driven automated tiered storage management. The IBM Spectrum Scale ILM toolkit helps you manage sets of files, pools of storage, and automate the management of file data. By using these tools, IBM Spectrum Scale can automatically determine where to physically store your data regardless of its placement in the logical directory structure. Storage pools, file sets, and user-defined policies can match the cost of your storage resources to the value of your data.

You can use IBM Spectrum Scale policy-based ILM tools to perform the following tasks:

- ▶ Create storage pools to provide a way to partition a file system's storage into collections of disks or a redundant array of independent disks (RAID) with similar properties that are managed together as a group.

IBM Spectrum Scale has the following types of storage pools:

- A required system storage pool that you create and manage through IBM Spectrum Scale.
- Optional user storage pools that you create and manage through IBM Spectrum Scale.
- Optional external storage pools that you define with IBM Spectrum Scale policy rules and manage through an external application, such as IBM Spectrum Archive EE.
- ▶ Create file sets to provide a way to partition the file system namespace to allow administrative operations at a finer granularity than that of the entire file system.
- ▶ Create policy rules that are based on data attributes to determine initial file data placement and manage file data placement throughout the life of the file.

## 2.3.2 Storage pools

Physically, a *storage pool* is a collection of disks or RAID arrays. You can use storage pools to group multiple storage systems within a file system. By using storage pools, you can create tiers of storage by grouping storage devices based on performance, locality, or reliability characteristics. For example, one pool can be an enterprise class storage system that hosts high-performance FC disks and another pool might consist of numerous disk controllers that host a large set of economical SATA disks.

There are two types of storage pools in an IBM Spectrum Scale environment: Internal storage pools and external storage pools. Internal storage pools are managed within IBM Spectrum Scale. External storage pools are managed by an external application, such as IBM Spectrum Archive EE. For external storage pools, IBM Spectrum Scale provides tools that you can use to define an interface that IBM Spectrum Archive EE uses to access your data.

IBM Spectrum Scale does not manage the data that is placed in external storage pools. Instead, it manages the movement of data to and from external storage pools. You can use storage pools to perform complex operations such as moving, mirroring, or deleting files across multiple storage devices, which provide storage virtualization and a single management context.

Internal IBM Spectrum Scale storage pools are meant for managing online storage resources. External storage pools are intended for use as near-line storage and for archival and backup operations. However, both types of storage pools provide you with a method to partition file system storage for the following considerations:

- ▶ Improved price-performance by matching the cost of storage to the value of the data
- ▶ Improved performance by:
  - Reducing the contention for premium storage
  - Reducing the impact of slower devices
  - Allowing you to retrieve archived data when needed
- ▶ Improved reliability by providing for:
  - Replication based on need
  - Better failure containment
  - Creation of storage pools as needed

### 2.3.3 Policies and policy rules

IBM Spectrum Scale provides a means to automate the management of files by using policies and rules. If you correctly manage your files, you can use and balance efficiently your premium and less expensive storage resources. IBM Spectrum Scale supports the following policies:

- ▶ File placement policies are used to place automatically newly created files in a specific storage pool.
- ▶ File management policies are used to manage files during their lifecycle by moving them to another storage pool, moving them to near-line storage, copying them to archival storage, changing their replication status, or deleting them.

A policy is a set of rules that describes the lifecycle of user data that is based on the file's attributes. Each rule defines an operation or definition, such as migrate to a pool and replicate the file. The rules are applied for the following uses:

- ▶ Initial file placement
- ▶ File management
- ▶ Restoring file data

When a file is created or restored, the placement policy determines the location of the file's data and assigns the file to a storage pool. All data that is written to that file is placed in the assigned storage pool. The placement policy that is defining the initial placement of newly created files and the rules for placement of restored data must be installed into IBM Spectrum Scale by running the `mmchpolicy` command. If an IBM Spectrum Scale file system does not have a placement policy that is installed, all the data is stored into the system storage pool. Only one placement policy can be installed at a time.

If you switch from one placement policy to another, or change a placement policy, that action has no effect on existing files. However, newly created files are always placed according to the currently installed placement policy.

The management policy determines file management operations, such as migration and deletion. To migrate or delete data, you must run the `mmapplypolicy` command. You can define the file management rules and install them in the file system together with the placement rules. As an alternative, you can define these rules in a separate file and explicitly provide them to `mmapplypolicy` by using the `-P` option. In either case, policy rules for placement or migration can be intermixed. Over the life of the file, data can be migrated to a different storage pool any number of times, and files can be deleted or restored.

With Version 3.1, IBM Spectrum Scale introduced the policy-based data management that automates the management of storage resources and the data that is stored on those resources. Policy-based data management is based on the storage pool concept. A storage pool is a collection of disks or RAID configurations with similar properties that are managed together as a group. The group under which the storage pools are managed together is the file system.

IBM Spectrum Scale provides a single name space across all pools. Files in the same directory can be in different pools. Files are placed in storage pools at creation time by using placement policies. Files can be moved between pools based on migration policies and files can be removed based on specific policies.

For more information about the SQL-like policy rule language, see *IBM Spectrum Scale: Administration Guide*, which is available at this website:

<https://www.ibm.com/support/knowledgecenter/STXKQY>

IBM Spectrum Scale V3.2 introduced external storage pools. You can set up external storage pools and GPFS policies that allow the GPFS policy manager to coordinate file migrations from a native IBM Spectrum Scale online pool to external pools in IBM Spectrum Archive EE. The GPFS policy manager starts the migration through the HSM client command-line interface embedded in the IBM Spectrum Archive EE solution.

For more information about GPFS policies, see 7.10, “Migration” on page 159.

## 2.3.4 Migration or premigration

The migration or premigration candidate selection is identical to the IBM Spectrum Scale native pool-to-pool migration/premigration rule. The Policy Engine uses the `ltfsee migrate` or `ltfsee premigrate` command for the migration or premigration of files from a native storage pool to an IBM Spectrum Archive EE tape cartridge pool.

There are two different approaches that can be used to drive an IBM Spectrum Archive EE migration through GPFS policies: Manual and automated. These approaches are only different in how the `mmapplypolicy` command (which performs the policy scan) is started.

### Manual

The manual IBM Spectrum Scale driven migration is performed when the user or a UNIX cron job runs the `mmapplypolicy` command with a predefined migration or premigration policy. The rule covers the migration or premigration of files from the system pool to the external IBM Spectrum Scale pool, which means that the data is physically moved to the external tape pool, which must be defined in IBM Spectrum Archive EE.

## Automated

The GPFS threshold migration is performed when the user specifies a threshold policy and the GPFS policy daemon is enabled to monitor the storage pools in the file system for that threshold. If a predefined high threshold is reached (which means the filling level of the storage pool reached the predefined high water mark), the monitor daemon automatically starts the `mmapplypolicy` command to perform an inode scan.

For more information about migration, see 7.10, “Migration” on page 159.

## 2.3.5 Active File Management

IBM Spectrum Scale Active File Management (AFM) is a scalable, high-performance file-system caching layer that is integrated with the IBM Spectrum Scale cluster file system. AFM is based on a home-cache model. A single home provides the primary file storage that is exported. One or more caches provide a view into the exported home file system without storing the file data locally. Upon file access in the cache, the data is fetched from home and stored in cache.

Another way to get files transferred from home to cache is through prefetching. Prefetching can use the IBM Spectrum Scale policy engine to quickly identify files that match certain criteria.

When files are created or changed in cache, they can be replicated back to home. A file that was replicated back to home can be evicted in cache. In this case, the user still sees the file in cache (the file is uncached), but the actual file content is stored in home. Eviction is triggered by the quota that is set on the AFM file set and can evict files based on size or last recent used criteria.

Cache must be an IBM Spectrum Scale independent file set. Home can be an IBM Spectrum Scale file system, a Network File System (NFS) export from any other file system, or a file server (except for the disaster-recovery use case). The caching relationship between home and cache can be based on the NFS or native IBM Spectrum Scale protocol. In the latter case, home must be an IBM Spectrum Scale file system in a different IBM Spectrum Scale cluster. The examples in this Redpaper that feature AFM use an NFS protocol at the home cluster.

The AFM relation is typically configured on the cache file set in one specific mode. The AFM mode determines where files can be processed (created, updated, and deleted) and how files are managed by AFM according to the file state. See Figure 2-5 on page 30 for AFM file states.

**Important:** IBM Spectrum Archive EE V1.2.3.0 is the first release that started supporting IBM Spectrum Scale AFM with IBM Spectrum Scale V4.2.2.3. AFM has multiple cache modes that can be created. However, IBM Spectrum Archive EE only supports the independent-writer (IW) cache mode. More cache modes might be supported soon.

### Independent-writer

AFM's IW cache mode makes the AFM target home for one or more caches. All changes in the caches are replicated to home asynchronously. Changes to the same data are applied in the home file set so that the changes are replicated from the caches. There is no cross-cache locking. Potential conflicts must be resolved at the respective cache site.

A file in the AFM cache can have different states as shown in Figure 2-5. File states can be different depending on the AFM modes.

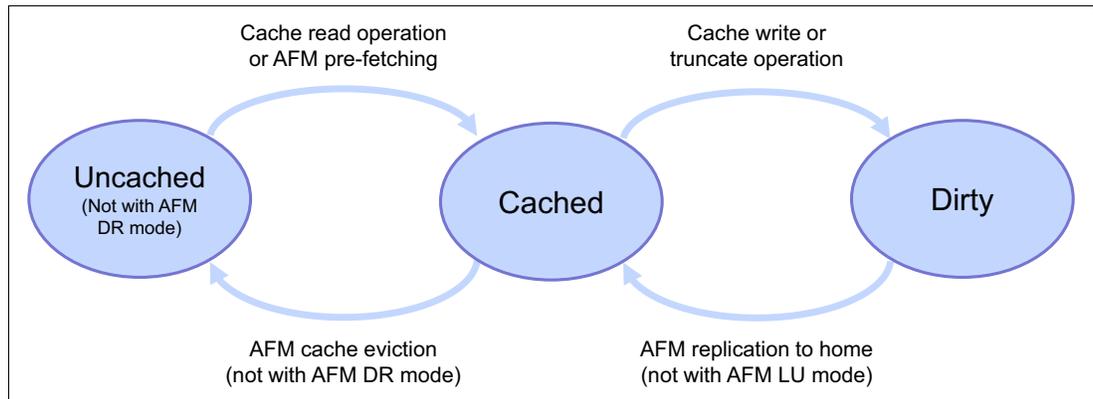


Figure 2-5 AFM file state transitions

### Uncached

When an AFM relation is created between cache and home, and files are available in home, these files can be seen in cache without being present. This state means that the file metadata is present in the cache, but the file content is still on home. Such files are in status uncached. In addition, the uncached status is achieved by evicting files from the cache.

### Cached

When an uncached file is accessed in cache for a read or write operation, the file is fetched from home. Fetching is the process of copying a file from home to cache. Files fetched from home to cache are in cached state. Another way to fetch files from home to cache is by using the AFM prefetch command (`mmafmctl prefetch`). This command can use the policy engine to identify files quickly, according to certain criteria.

### Dirty

When a cached file in the AFM cache is modified, that file is marked as dirty, indicating that it is a candidate for replication back to home. The dirty status of the file is reset to cached if the file has been replicated to home. When a file is deleted in cache, this delete operation is also done on home.

For information about how to configure AFM with IBM Spectrum Archive EE, see 6.4, “Configuring IBM Spectrum Archive EE with IBM Spectrum Scale AFM” on page 118. For use cases, see 9.2, “AFM use cases” on page 269.

## 2.4 OpenStack SwiftHLM

The Swift High Latency Media (SwiftHLM) project seeks to create a high-latency storage back end that makes it easier for users to perform bulk operations of data tiering within a Swift data ring. SwiftHLM enables IBM Spectrum Scale, IBM Spectrum Archive, and IBM Spectrum Protect as the key products for this software-defined hybrid storage with object interface to tape technology. Data is produced at significantly higher rates than a decade ago.

The storage and data management solutions of the past can no longer keep up with the data demands of today. The policies and structures that decide and execute how that data is used, discarded, or retained determines how efficiently the data is used. The need for intelligent data management and storage is more critical now than ever before.

Traditional management approaches hide cost-effective, high-latency media (HLM) storage, such as tape or optical disk archive back ends, underneath a traditional file system. The lack of HLM-aware file system interfaces and software makes it difficult for users to understand and control data access on HLM storage. Coupled with data-access latency, this lack of understanding results in slow responses and potential timeouts that affect the user experience.

The Swift HLM project addresses this challenge. Running OpenStack Swift on top of HLM storage allows you to cheaply store and efficiently access large amounts of infrequently used object data. Data that is stored on tape storage can be easily adopted to an Object Storage data interface. SwiftHLM can be added to OpenStack Swift (without modifying Swift) to extend Swift's interface.

This ability allows users to explicitly control and query the state (on disk or on HLM) of Swift object data, including efficient pre-fetch of bulk objects from HLM to disk when those objects must be accessed. This function, previously missing in Swift, provides similar functions as Amazon Glacier does through the Glacier API or the Amazon S3 Lifecycle Management API.

BDT Tape Library Connector (open source) and IBM Spectrum Archive or IBM Spectrum Protect™ are examples of HLM back ends that provide important and complex functions to manage HLM resources (tape mounts and unmounts to drives, serialization of requests for tape media, and tape drive resources). They can use SwiftHLM functions for a proper integration with Swift.

Although access to data that is stored on HLM can be done transparently without the use of SwiftHLM, this process does not work well in practice for many important use cases and other reasons. SwiftHLM function can be orthogonal and complementary to Swift (ring to ring) tiering (source). The high-level architecture of the low cost, high-latency media storage solution is shown in Figure 2-6 on page 32.

For more information, see *Implementing OpenStack SwiftHLM with IBM Spectrum Archive EE or IBM Spectrum Protect for Space Management*, REDP-5430:

<http://www.redbooks.ibm.com/abstracts/redp5430.html>

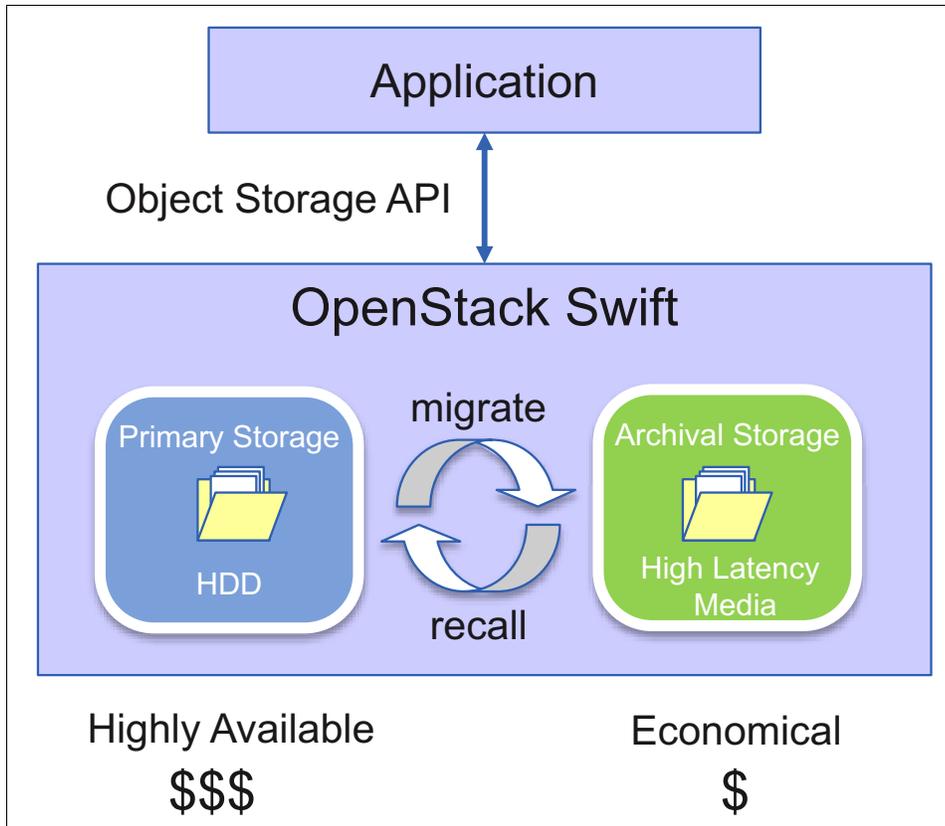


Figure 2-6 High-level SwiftHLM architecture

## 2.5 IBM Spectrum Archive EE dashboard

IBM Spectrum Archive EE V1.2.4.0 introduces dashboard capabilities that allow customers to visualize their data through a graphical user interface (GUI). By using the dashboard, you can see the following things without logging in to a system and typing commands by using a web-browser:

- ▶ See whether a system is running without error. If there is an error, see what kind of error is detected.
- ▶ See basic tape-related configurations like how many pools and how much space is available.
- ▶ See time-scaled storage consumption for each tape pool.
- ▶ See the throughput for each drive for migration and recall.

This monitoring feature consists of multiple components that are installed in the IBM Spectrum Archive EE nodes as well as dedicating an external node for displaying the dashboard.

The dashboard consists of the following components:

- ▶ Logstash
- ▶ Elasticsearch
- ▶ Grafana

Logstash is used for data collection, and should be installed on all IBM Spectrum Archive EE nodes. The data that is collected by Logstash is then sent to Elasticsearch on the external monitoring node where it can query data quickly and send it to the Grafana component for visualization. Figure 2-7 shows the IBM Spectrum Archive EE Dashboard architecture.

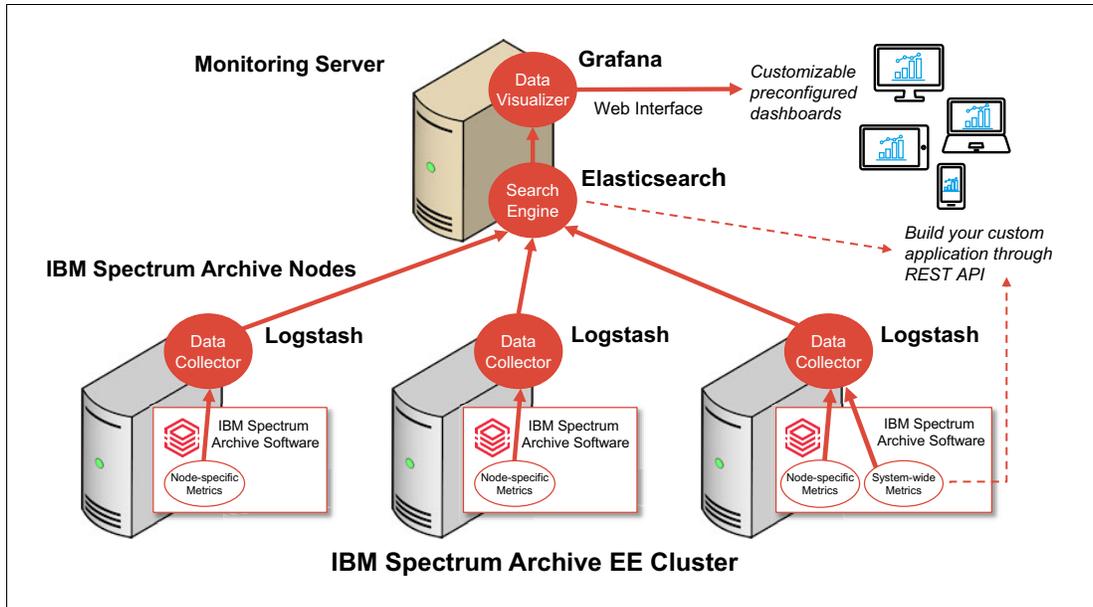


Figure 2-7 IBM Spectrum Archive Dashboard architecture

The Dashboard views are System Health, Storage, Activity, and Config. Figure 2-8 shows an example of the IBM Spectrum Archive EE Dashboard Activity view.

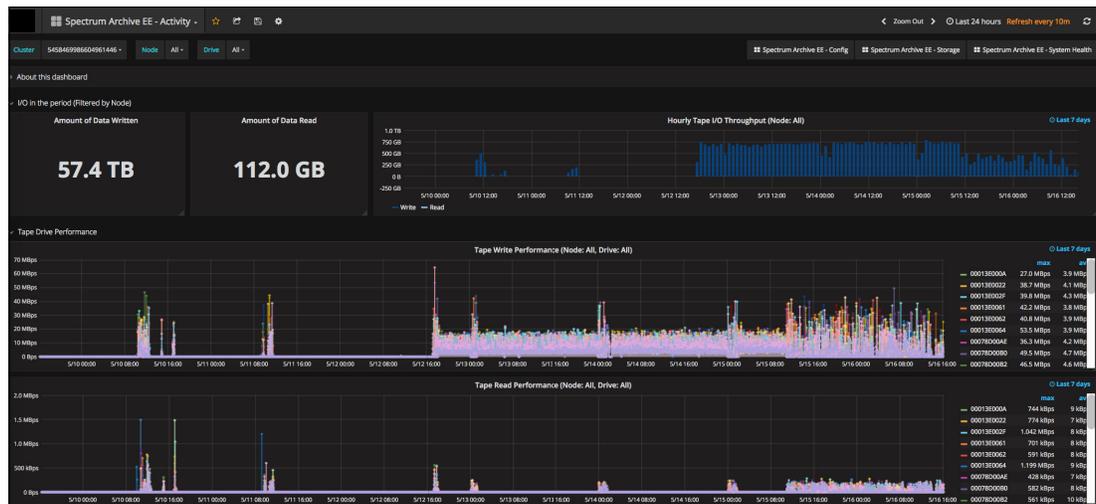


Figure 2-8 IBM Spectrum Archive EE Dashboard activity view

For more information on configuring the Dashboard within your environment, see the *IBM Spectrum Archive Enterprise Edition Dashboard Deployment Guide*:

<https://www.ibm.com/developerworks/community/wikis/form/anonymous/api/wiki/fa32927c-e904-49cc-a4cc-870bcc8e307c/page/89889f48-40ee-41b6-9510-5b6de6fc8409/attachment/779a3b43-cd79-42a7-919a-bf7e2bb03c06/media/SpectrumArchiveDashboard-DeploymentGuide.pdf>

## 2.6 IBM Spectrum Archive EE REST API

The Representational State Transfer (REST) API for IBM Spectrum Archive Enterprise Edition can be used to access data on the IBM Spectrum Archive Enterprise Edition system. Starting with Version 1.2.4, IBM Spectrum Archive EE provides the configuration information through its REST API. The GET operation returns the array of configured resources similar to CLI commands, but in well-defined JSON format.

They are equivalent to what the `ltfsee info XXX` and `ltfsee XXX show` commands display. With the REST API, you can automate these queries and integrate the information into your applications including the web/cloud.

For installation instructions, see 4.4, “Installing a RESTful server” on page 66.

For usage examples including commonly used parameters, see 7.29, “IBM Spectrum Archive REST API” on page 215.

## 2.7 Types of archiving

It is important to differentiate between archiving and the HSM process that is used by IBM Spectrum Archive EE. When a file is migrated by IBM Spectrum Archive EE from your local system to tape storage, a placeholder or stub file is created in place of the original file. Stub files contain the necessary information to recall your migrated files and remain on your local file system so that the files appear to be local. This process contrasts with archiving, where you often delete files from your local file system after archiving them.

The following types of archiving are used:

- ▶ Archive with no file deletion
- ▶ Archive with deletion
- ▶ Archive with stub file creation (HSM)
- ▶ Compliant archiving

Archiving with no file deletion is the typical process that is used by many backup and archive software products. In the case of IBM Spectrum Protect, an archive creates a copy of one or more files in IBM Spectrum Protect with a set retention period. It is often used to create a point-in-time copy of the state of a server's file system and this copy is kept for an extended period. After the archive finishes, the files are still on the server's file system.

Contrast this with archiving with file deletion where after the archive finishes the files that form part of the archive are deleted from the file system. This is a feature that is offered by the IBM Spectrum Protect archive process. Rather than a point-in-time copy, it can be thought of as a point-in-time move as the files are moved from the servers' file system into IBM Spectrum Protect storage.

If the files are needed, they must be manually retrieved back to the file system. A variation of this is *active archiving*, which is a mechanism for moving data between different tiers of storage depending on its retention requirements. For example, data that is in constant use is kept on high-performance disk drives, and data that is rarely referenced or is required for long-term retention is moved to lower performance disk or tape drives.

IBM Spectrum Archive EE uses the third option, which instead of deleting the archived files, it creates a stub file in its place. If the files are needed, they are automatically retrieved back to the file system by using the information that is stored in the stub file when they are accessed.

The final type of archiving is compliant archiving, which is a legislative requirement of various countries and companies data retention laws, such as Sarbanes-Oxley in the US. These laws require a business to retain key business information. Failure to comply with these laws can result in fines and sanctions. Essentially, this type of archiving results in data being stored by the backup software without the possibility of it being deleted before a defined period elapses. In certain cases, it can never be deleted.

**Important:** IBM Spectrum Archive EE is not a compliant archive solution.





# Planning

This chapter provides planning information that is related to the IBM Spectrum Archive Enterprise Edition (EE). Review the Planning section within the IBM Spectrum Archive EE IBM Knowledge Center website:

[https://www.ibm.com/support/knowledgecenter/en/ST9MBR\\_1.2.6/tfs\\_ee\\_planning.html](https://www.ibm.com/support/knowledgecenter/en/ST9MBR_1.2.6/tfs_ee_planning.html)

The most current information for IBM Spectrum Archive EE hardware and software configurations, notices, and limitations can always be found in the readme file of the software package.

This chapter includes the following topics:

- ▶ System requirements
- ▶ Required software
- ▶ Hardware and software setup
- ▶ Sizing and settings

## 3.1 System requirements

IBM Spectrum Archive EE supports the Linux operating systems and hardware platforms that are shown in Table 3-1. For the latest tape device drivers, check Fix Central:

<http://www.ibm.com/support/fixcentral>

Table 3-1 Linux system requirements

<b>Linux computers</b>	
Supported operating systems (x86_64)	Red Hat Enterprise Linux Server 7.1 Red Hat Enterprise Linux Server 7.2 Red Hat Enterprise Linux Server 7.3 Red Hat Enterprise Linux Server 7.4
Supported operating systems (ppc64le)	Red Hat Enterprise Linux Server 7.4 or later
Supported tape libraries	A single logical library within a supported tape library for each Spectrum Archive EE Control Node.
Supported tape drives	IBM TS1140 tape drive.
	IBM TS1150 tape drive.
	IBM TS1155 tape drive.
	LTO-5 full-high tape drive.
	LTO-6 full-high tape drive.
	LTO-7 full-high tape drive.
	LTO-8 full-high tape drive.
Supported tape media	TS1140 media: JB, JC, JK, and JY. TS1155/TS1150 media: JC, JD, JK, JL, JY, and JZ.
	LTO media: LTO 8, LTO 7 Type M8, LTO 7, LTO 6, and LTO 5.
<b>Server</b>	
Processor	Minimum: A x86_64 processor. Preferred: Dual socket server with the latest chipset.
Memory	Minimum: (d) x (f) + 1 GB of RAM available for the IBM Spectrum Archive EE program: <ul style="list-style-type: none"> <li>▶ d: Number of tape drives</li> <li>▶ f: Number of millions of files/directories on the tape cartridges</li> </ul> In addition, IBM Spectrum Scale must be configured with adequate RAM. Example: There are six tape drives in the system and three million files are stored on the tape cartridges. The required RAM is 19 GB (6 x 3 + 1 = 19). Preferred: 64 GB RAM and greater.
HBA	Minimum: Fibre Channel Host Bus Adapter supported by TS1155, TS1150, TS1140, LTO-8, LTO-7, LTO-6, and LTO-5 tape drives. Preferred: 8 Gbps/ 16 Gbps Dual port or Quad port Fibre Channel Host Bus Adapter.
Network	TCP/IP based protocol network.

Linux computers	
Disk device for LTFS EE tape file system metadata	For more information, see 3.4.1, “IBM Spectrum Archive EE metadata file system” on page 43.
One or more disk devices for the GPFS file system	The amount of disk space that is required depends on the IBM Spectrum Scale settings that are used.

### 3.1.1 Limitations

There are some limitations to the IBM Spectrum Archive EE software relating to IBM Spectrum Protect™, Scale-out Backup and Recovery (SOBAR), Advanced File Management (AFM), and Transparent Cloud Tiering (TCT).

If IBM Spectrum Protect clients are being used for managing the files in the same IBM Spectrum Scale cluster, the following limitations apply:

- ▶ If the IBM Spectrum Protect backup client needs to back up the files of the file system that is managed by IBM Spectrum Archive EE, you must schedule the migration process to start after the files are backed up.
- ▶ If an IBM Spectrum Scale cluster is used with IBM Spectrum Archive EE, any file system that is in the same IBM Spectrum Scale cluster cannot be used with IBM Spectrum Protect for Space Management for migrating the data to IBM Spectrum Protect servers.

The file system that is managed by IBM Spectrum Archive EE cannot be used with the following functions or features of IBM Spectrum Scale:

- ▶ Scale-out Backup and Recovery (SOBAR)
- ▶ Advanced File Management (AFM), in a mode other than Independent Writer mode
- ▶ Transparent Cloud Tiering (TCT)

## 3.2 Required software

This section describes the required software for IBM Spectrum Archive EE on Red Hat systems. The following RPM Package Manager software (RPMS) must be installed and be at latest levels for a Red Hat Enterprise Linux system before installing IBM Spectrum Archive EE V1.2.6.

x86\_64 operating system prerequisites:

- ▶ IBM Spectrum Scale version 4.2.2, 4.2.3, 5.0.0.x, or 5.0.1.x (Standard Edition, Advanced Edition, or Data Management Edition).

**Note:** IBM Spectrum Scale version 4.2.3 requires a patch to the installation toolkit when IBM Spectrum Archive Enterprise Edition is installed.

- ▶ pyOpenSSL
- ▶ fuse
- ▶ fuse-libs
- ▶ libxml2
- ▶ libuuid
- ▶ libicu
- ▶ rpcbind
- ▶ python-2.7.5-34 or later, but earlier than 3.0

- ▶ IBM tape device driver for Linux (lin\_tape)
- ▶ Net-snmp
- ▶ attr
- ▶ boost-thread
- ▶ boost-filesystem
- ▶ gperftools-libs
- ▶ lsof

**Note:** These required RPMS were gathered from an RHEL 7.2 machine. Your RPMS versions might differ depending on the version of your RHEL OS.

IBM Power server platform (ppc64le) operating system prerequisites:

- ▶ IBM Spectrum Scale version 4.2.3.5 or subsequent 4.2.3.x fix pack levels.
- ▶ IBM Spectrum Scale version 5.0.0.x or subsequent 5.0.1.x fix pack levels.
- ▶
- ▶ pyOpenSSL
- ▶ fuse
- ▶ fuse-libs
- ▶ libxml2
- ▶ libuuid
- ▶ libicu
- ▶ rpcbind
- ▶ python-2.7.5-34 or later, but earlier than 3.0
- ▶ IBM tape device driver for Linux (lin\_tape)
- ▶ Net-snmp
- ▶ attr
- ▶ boost-thread
- ▶ boost-filesystem
- ▶ gperftools-libs
- ▶ lsof

Java virtual machine (JVM) 1.7 or later must be installed before you install IBM Spectrum Archive Enterprise Edition.

The most current version of the IBM tape device driver for Linux (lin\_tape).

### 3.3 Hardware and software setup

Valid combinations of IBM Spectrum Archive EE components in an IBM Spectrum Scale cluster are shown in Table 3-2.

Table 3-2 Valid combinations for types of nodes in the IBM Spectrum Scale cluster

Node type	IBM Spectrum Scale	IBM Spectrum Archive internal Hierarchical Storage Manager (HSM)	IBM Spectrum Archive LE+	Multi-tape management module (MMM)
IBM Spectrum Scale only node	Yes	No	No	No
IBM Spectrum Archive EE node	Yes	Yes	Yes	Yes

All other combinations are invalid as an IBM Spectrum Archive EE system. IBM Spectrum Archive EE nodes have connections to the IBM tape libraries and drives.

Multiple IBM Spectrum Archive EE nodes enable access to the same set of IBM Spectrum Archive EE tapes. The purpose of enabling this capability is to increase the performance of the host migrations and recalls by assigning fewer tape drives to each IBM Spectrum Archive EE node. The number of drives per node depends on the HBA/switch/host combination. The idea is to have the maximum number of drives on the node such that all drives on the node can be writing or reading at their maximum speeds.

The following hardware/software/configuration setup must be prepared before IBM Spectrum Archive EE is installed:

- ▶ IBM Spectrum Scale is installed on each of the IBM Spectrum Archive EE nodes.
- ▶ The IBM Spectrum Scale cluster is created and all of the IBM Spectrum Archive EE nodes belong to the cluster.

A single NUMA node is preferable for better performance. For servers that contain multiple CPUs, the key is to remove memory from the other CPUs and group them in a single CPU to create a single NUMA node. This configuration allows all the CPUs to access the shared memory, resulting in higher read/write performances between the disk storage and tape storage.

FC switches can be added between the host and tape drives as well as between the host and the disk storage to create a storage area network (SAN) to further expand storage needs as required.

If you plan to configure your own server, an example of a server configuration is listed below for a Lenovo x3650 M5 rack server: Server System x3650 M5 - 8871KXU:

- ▶ Intel Xeon Processor E5-2640 v4 10C 2.4 GHz 25 MB Cache 2133 MHz 90 W (Standard)
- ▶ 16 GB TruDDR4 Memory (2Rx4, 1.2 V) PC4-19200 CL17 2400 MHz LP RDIMM (Standard) x 4
- ▶ ServeRAID M5210 SAS/SATA Controller (Standard)
- ▶ System x Enterprise Slides Kit (Standard)
- ▶ System x 900 W High Efficiency Platinum AC Power Supply (Standard)
- ▶ System x 900 W High Efficiency Platinum AC Power Supply (Standard)

- ▶ System x3650 M5 PCIe Riser 1 (1 x16 FH/FL + 1 x8 FH/HL Slots):
  - QLogic 8 Gb FC Dual-port HBA (16 Gb preferred)
  - QLogic 8 Gb FC Dual-port HBA (16 Gb preferred)
- ▶ System x3650 M5 PCIe Riser 2 (1 x16 FH/FL + 1 x8 FH/HL Slots):
  - QLogic 8 Gb FC Dual-port HBA (16 Gb preferred)
  - QLogic 8 Gb FC Dual-port HBA (16 Gb preferred)
- ▶ x3650 M5 PCIe x8:
  - QLogic 8 Gb FC Dual-port HBA (16 Gb preferred)
- ▶ x3650 M5 PCIe x8:
  - QLogic 8 Gb FC Dual-port HBA (16 Gb preferred)

**Note:** Servers with multiple sockets can have the sockets physically removed.

Figure 3-1 provides an example of an IBM Spectrum Archive EE node hardware setup with a x3650 M5 rack server.

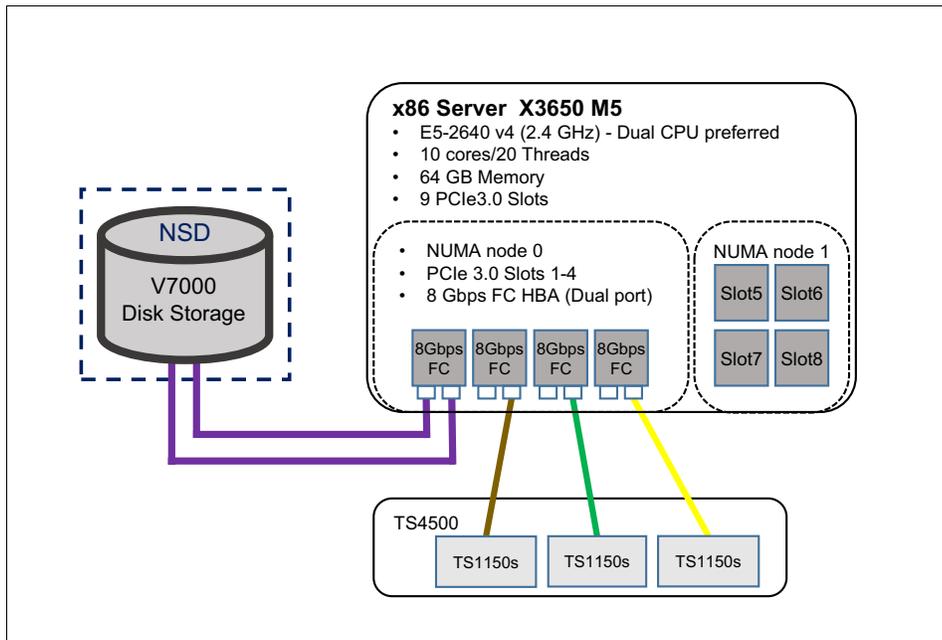


Figure 3-1 IBM Spectrum Archive EE node hardware setup with a x3650 M5 rack server

**Note:** The example uses a x3650 M5 rack server because all standard models come by default with an Intel Xeon E5-2600 v5 single socket processor with 4 PCIe slots dedicated to the socket. For more information about the x3650 m5 server, see this website:

<http://shop.lenovo.com/us/en/systems/servers/racks/x3650-m5>

## 3.4 Sizing and settings

There are many considerations that are required when you are planning for an IBM Spectrum Scale file system, including the IBM Spectrum Archive EE HSM-managed file system and the IBM Spectrum Archive EE metadata file system. Thus, this section describes the IBM Spectrum Scale file system aspects to help avoid the need to make changes later.

### 3.4.1 IBM Spectrum Archive EE metadata file system

IBM Spectrum Archive EE requires space for the file metadata that is stored on an IBM Spectrum Scale file system. If this metadata file system is separate from the IBM Spectrum Scale space-managed file systems, you must ensure that the size and number of inodes of the metadata file system is large enough to handle the number of migrated files.

The IBM Spectrum Archive EE metadata directory can be stored in its own IBM Spectrum Scale file system or it can share the IBM Spectrum Scale file system that is being space-managed.

When the IBM Spectrum Archive EE metadata file system is using the same IBM Spectrum Scale file system to be space-managed, it has the advantage of being flexible by sharing the resources. Space-managed and IBM Spectrum Archive EE metadata can accommodate each other by growing and shrinking as needed. Thus, it is suggested that you have a single file system. For metadata optimization, it is preferable to put the GPFS metadata and the IBM Spectrum Archive metadata on SSDs or flash.

**Note:** It is suggested that you just have a single file system.

The size requirements of the IBM Spectrum Scale file system that is used to store the IBM Spectrum Archive EE metadata directory depends on the block size and the number of files that are migrated to IBM Spectrum Archive EE.

The following calculation produces an estimate of the minimum number of inodes that the IBM Spectrum Scale file system must have available. The calculation depends on the number of cartridges:

Number of inodes =  $500 + (15 \times c)$  (Where  $c$  is the number of cartridges.)

**Important:** If there is more than one tape library, the number of cartridges in your calculation must be the total number of cartridges in all libraries.

The following calculation produces an estimate of the size of the metadata that the IBM Spectrum Scale file system must have available:

Number of GBs =  $10 + (3 \times F \times N)$

where:

- ▶  $F$  is the number of files, in millions, to migrate.
- ▶  $N$  is the number of replicas to create.

For example, to migrate 50 million files to two tape storage pools, 310 GB of metadata is required:

$10 + (3 \times 50 \times 2) = 310$  GB

## 3.4.2 Redundant copies

The purpose of redundant copies is to enable the creation of multiple LTFS copies of each GPFS file during migration. One copy is considered to be the primary, and the other copies are considered the redundant copies. The redundant copies can be created only in pools that are different from the pool of the primary copy and different from the pools of other redundant copies. The maximum number of redundant copies is two. The primary copy and redundant copies can be in a single tape library or spread across two tape libraries.

Thus, to ensure that file migration can occur on redundant copy tapes, the number of tapes in the redundant copy pools must be the same or greater than the number of tapes in the primary copy pool. For example, if the primary copy pool has 10 tapes, the redundant copy pools also should have at least 10 tapes. For more information about redundant copies, see 7.10.4, “Replicas and redundant copies” on page 170.

**Note:** The most common setup is to have two copies, that is, one primary and one copy pool with the same number of tapes.

## 3.4.3 Performance

Performance planning is an important aspect of any IBM Spectrum Archive EE implementation, specifically migration performance. The migration performance is the rate at which IBM Spectrum Archive EE can move data from disk to tape, freeing up space on disk. The number of tape drives (including tape drive generation) and servers that are required for the configuration can be determined based on the amount of data that needs to be moved per day and an estimate of the average file size of that data.

**Note:** Several components of the reference architecture affect the overall migration performance of the solution, including backend disk speeds, SAN connectivity, NUMA node configuration, and amount of memory. Thus, this migration performance data should be used as a guideline only and any final migration performance measurements should be done on the actual customer hardware.

For more migration performance information, see the *IBM Spectrum Archive Enterprise Edition V1.2.2 Performance* white paper, which is available at:

<https://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/WP102688>

The configuration shown in Figure 3-2 on page 45 was used to run the lab performance test in this section.

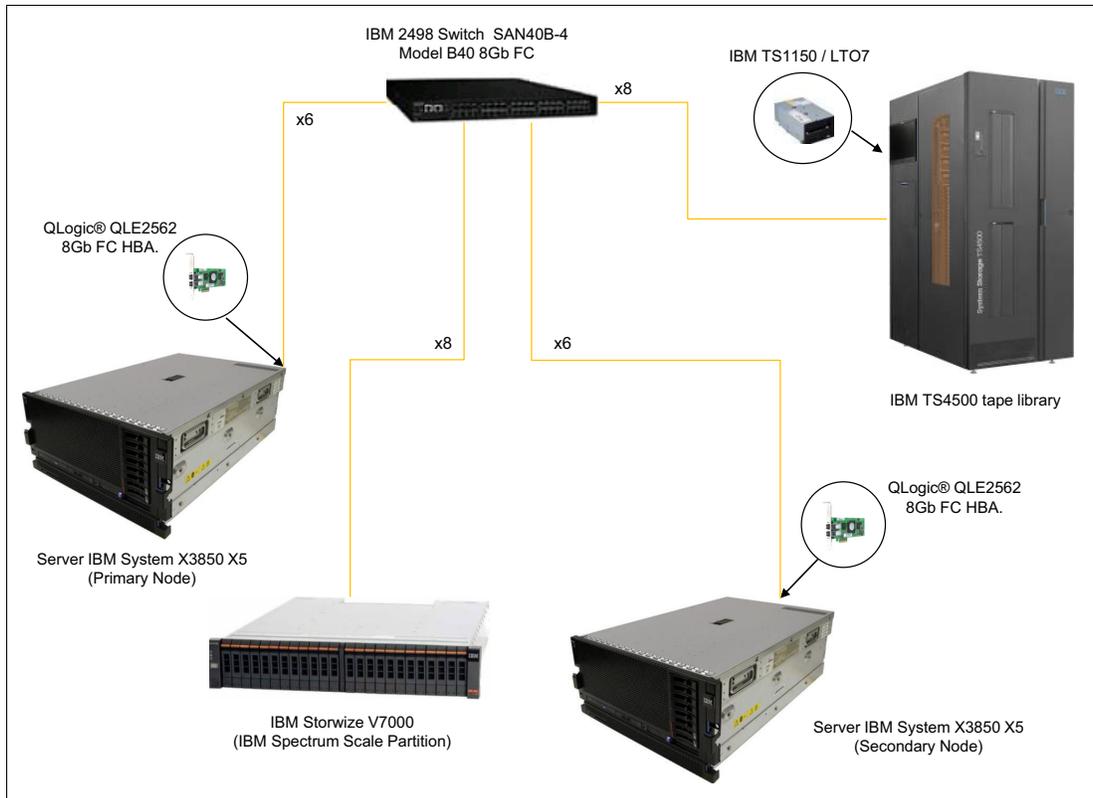


Figure 3-2 IBM Spectrum Archive EE configuration used in lab performance

The performance data shown in this section was derived by using two x3850 X5 servers consisting of multiple QLogic QLE2562 8 Gb FC HBA cards. The servers were also modified by moving all the RAM memory to a single CPU socket to create a multi-CPU, single NUMA node. The HBAs were relocated so they are on one NUMA node.

This modification was made so that memory can be shared locally to improve performance. The switch used in the lab was an IBM 2498 model B40 8 Gb FC switch that was zoned out so that it had a zone for each HBA in each node. An IBM Storwize® V7000 disk storage unit was used for the disk space and either TS1150 or TS1070 drives were used in a TS4500 tape library.

Figure 3-2 shows the example configuration. This is one of many configurations. Yours might be different.

Figure 3-3 on page 46 shows a basic internal configuration of Figure 3-2 in more detail. The two x3850 X5 servers are using a single NUMA and have all the HBAs performing out of it. This configuration allows all CPUs within the servers to work more efficiently. The 8 GB Fibre Channel switch is broken up so that each zone handles a single HBA on each server.

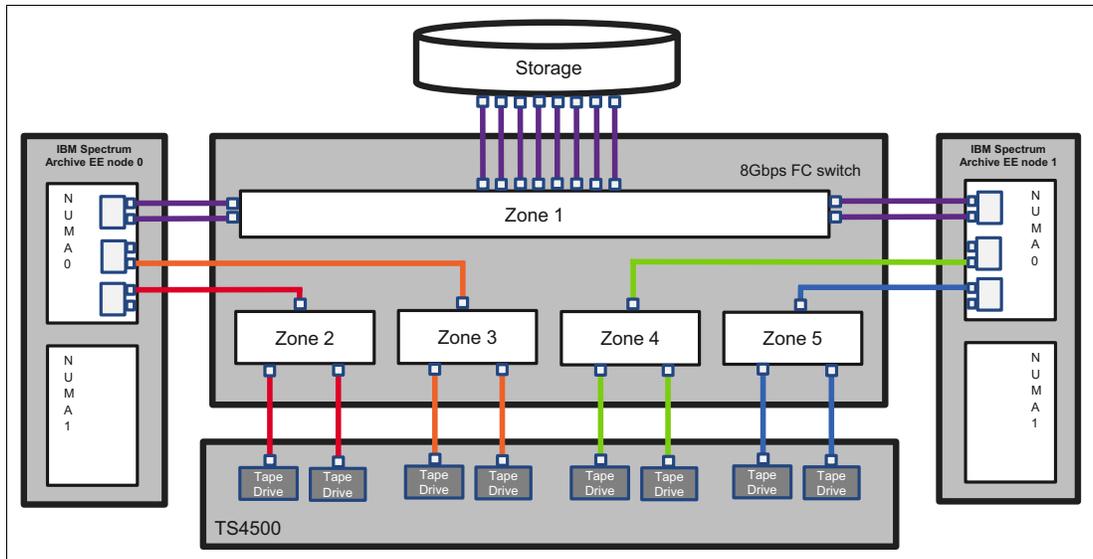


Figure 3-3 Internal diagram of two nodes running off NUMA 0 connected to an 8 Gb switch with four drives per node and one external storage unit

**Note:** Figure 3-3 shows a single Fibre Channel cable going to the drive zones. However, generally you should have a second Fibre Channel cable for failover scenarios. This is the same reason why the zone that goes to the external storage unit has two Fibre Channel cables from 1 HBA on each server.

This diagram should be used as a guide on how you would set up your own environment for best performances. This is one of many configurations. For instance, if you have more drives, you can add more tape drives per zone, or add more HBAs on each server and create a new zone.

Table 3-3 and Table 3-4 show the raw data of total combined transfer rate in MiB/s on multiple node configurations with various files sizes (the combined transfer rate of all drives). In these tables, N represents nodes, D represents number of drives per node, and T represents the total number of drives for the configuration.

With a TS1150 configuration of 1N4D4T, you can expect to see a combined total transfer rate of 1244.9 MBps for 10 GiB files. If that configuration is doubled to 2N4D8T, the total combined transfer rate is nearly doubled to 2315.3 MiB/s for 10 GiB files. With this information, you can estimate the total combined transfer rate for your configuration.

*Table 3-3 TS1150 raw performance data for multiple node/drive configurations with 5 MiB, 10 MiB, 1 GiB, and 10 GiB files*

Node/Drive Configuration	File Size				
	5 MiB	10 Mib	100 MiB	1 Gib	10 Gib
8 Drives (2N4D8T)	369.0	577.3	1473.4	2016.4	2315.3
6 Drives (2N3D6T)	290.6	463.2	1229.5	1656.2	1835.5
4 Drives (1N4D4T)	211.3	339.3	889.2	1148.4	1244.9
3 Drives (1N3D3T)	165.3	267.1	701.8	870.6	931.0
2 Drives (1N2D2T)	114.0	186.3	465.8	583.6	624.3

Table 3-4 shows the raw performance for the LTO7 drives.

*Table 3-4 LTO7 raw performance data for multiple node/drive configurations with 5 MiB, 10 MiB, 1 GiB, and 10 GiB files*

Node/Drive Configuration	File Size				
	5 MiB	10 Mib	100 MiB	1 Gib	10 Gib
8 Drives (2N4D8T)	365.4	561.8	1287.8	1731.8	1921.7
6 Drives (2N3D6T)	286.3	446.5	1057.9	1309.4	1501.7
4 Drives (1N4D4T)	208.5	328.1	776.7	885.6	985.1
3 Drives (1N3D3T)	162.9	254.4	605.5	668.3	749.1
2 Drives (1N2D2T)	111.0	178.4	406.7	439.2	493.7

Figure 3-4 shows a comparison line graph of the raw performance data obtain for TS1150 drives and LTO7 drives that use the same configurations. It is clear that using small files, the comparison between the two types of drives is minimal. However, when migrating file sizes of 1 GiB and greater, there is a noticeable difference. Comparing the biggest configuration of 2N4D8T, the LTO peaks at a total combined transfer rate of 1921.7 MiB/s. With the same configuration but with TS1150 drives, it peaks at a total combined transfer rate of 2315.3 MiB/s.

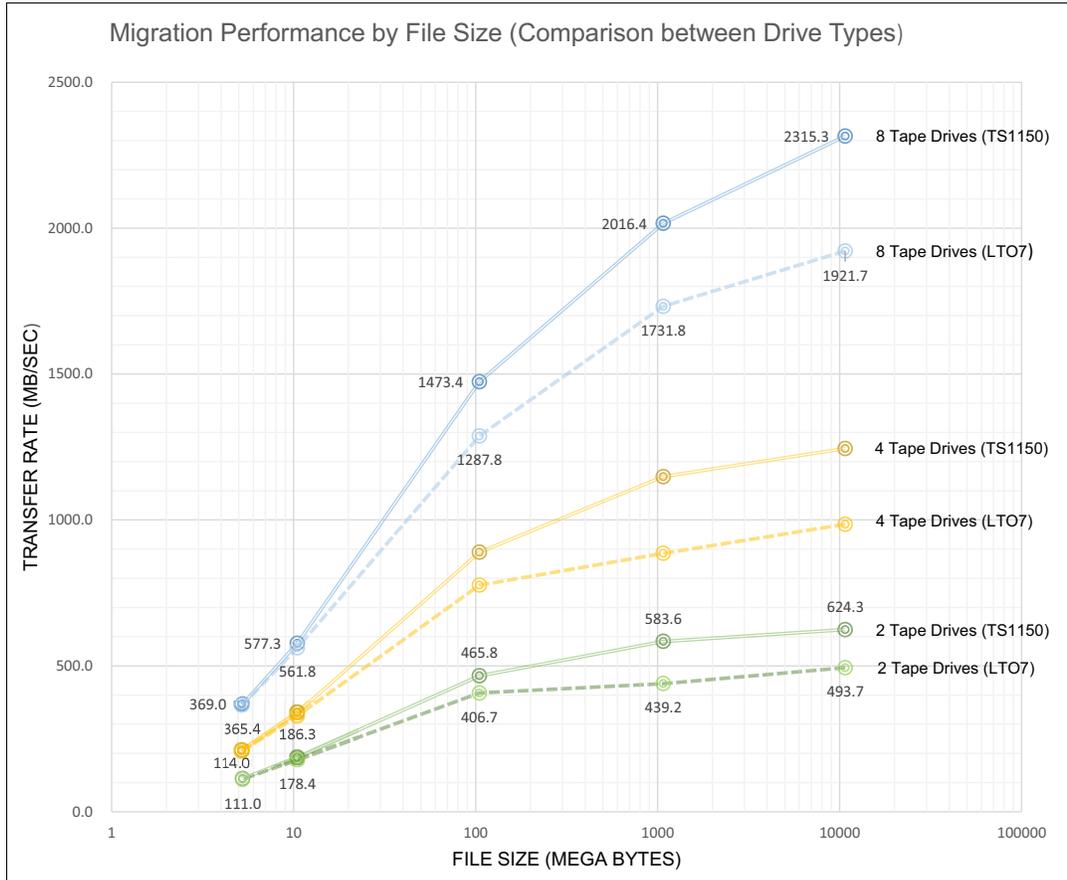


Figure 3-4 Comparison between TS1150 and LTO-7 drives using multiple node/drive configurations



# Installation

This chapter provides information about the distribution and installation of IBM Spectrum Archive Enterprise Edition (EE). It describes the following main aspects:

- ▶ Installing IBM Spectrum Archive EE on a Linux system

This section describes how to install the IBM Spectrum Archive EE program on a Linux system (in our example, we use a Red Hat-based Linux server system). It describes the installation routine step-by-step and reviews the prerequisites, such as verification of the Linux preinstallation packages and the installation of the IBM tape device driver under Linux.

- ▶ Quick installation guide for IBM Spectrum Archive EE

This optional section provides some background information about how to upgrade the tape library or tape drive firmware for use with the IBM Spectrum Archive EE.

This chapter includes the following topics:

- ▶ Installing IBM Spectrum Archive EE on a Linux system
- ▶ Installation prerequisites for IBM Spectrum Archive EE
- ▶ Installing IBM Spectrum Archive EE
- ▶ Installing a RESTful server
- ▶ Quick installation guide for IBM Spectrum Archive EE
- ▶ Library replacement

## 4.1 Installing IBM Spectrum Archive EE on a Linux system

The first part of this chapter describes how to install the IBM Spectrum Archive EE program on a Linux server system. In our lab setup for writing this book, we used a Red Hat-based Linux system. All the examples in this chapter are based on that release.

Although IBM Spectrum Archive EE is based on the IBM Linear Tape File System standard components as IBM Spectrum Archive and the IBM Spectrum Archive Library Edition (LE), these components are all included with the IBM Spectrum Archive EE installation package.

Before you can start with the installation routines, you must verify the following installation prerequisites:

- ▶ Installation prerequisites (see 4.2, “Installation prerequisites for IBM Spectrum Archive EE” on page 50)

This section describes the tasks that must be completed before installing IBM Spectrum Archive EE.

- ▶ Installing IBM Spectrum Archive EE on a Linux server (see 4.3, “Installing IBM Spectrum Archive EE” on page 57)

This section describes how to install the IBM Spectrum Archive EE package on a Linux server.

## 4.2 Installation prerequisites for IBM Spectrum Archive EE

This section describes the tasks that must be completed before installing IBM Spectrum Archive EE.

Ensure that the following prerequisites are met before IBM Spectrum Archive EE is installed. For more information, see the other topics in this section if needed.

- ▶ Verify that your computer meets the minimum hardware and software requirements for installing the product. For more information, see Chapter 3, “Planning” on page 37.
- ▶ Verify that your user ID meets the requirements for installing the product (such as you are working with the root user ID or have the root administration permissions).
- ▶ Ensure that you reviewed all of the planning information that is described in Chapter 3, “Planning” on page 37.
- ▶ If the standard IBM Spectrum Archive LE is already installed, it must be uninstalled before IBM Spectrum Archive EE is installed. For more information, see 11.4, “IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products” on page 331.
- ▶ Ensure that all prerequisite software is installed, as described in 4.2.1, “Installing the host bus adapter and device driver” on page 51.
- ▶ Ensure that the host bus adapter (HBA) and device driver are installed, as described in 4.2.1, “Installing the host bus adapter and device driver” on page 51.
- ▶ Ensure that the IBM tape device driver is installed, as described in 4.2.2, “Installing the IBM tape device driver” on page 51.

- ▶ Determine the distribution package for IBM Spectrum Archive EE that is required for your system. For more information, see 4.2.2, “Installing the IBM tape device driver” on page 51.
- ▶ For IBM Spectrum Scale prerequisites, see 3.3, “Hardware and software setup” on page 41.

## 4.2.1 Installing the host bus adapter and device driver

This section describes how to install the HBA and its device driver for use with IBM Spectrum Archive EE.

To install the HBA and its device driver, see the documentation that is provided by the HBA manufacturer.

If the HBA attached to the tape library is an Emulex adapter, add the following line to the `/etc/modprobe.d/lpfc.conf` file:

```
options lpfc lpfc_sg_seg_cnt=256
```

Then, restart the server system for the change to take effect.

For more information about fixes and updates for your system’s software, hardware, and operating system, see this website:

<http://www.ibm.com/support/fixcentral>

## 4.2.2 Installing the IBM tape device driver

This section describes how to install the IBM tape device driver for use with the integrated LTFS function and later on for the actual IBM Spectrum Archive EE. In our examples, we refer to a Red Hat-based Linux server system that we used for writing this book and creating the examples. Always check the IBM Fix Central website for the current version of the IBM tape device driver, as shown in the following steps. The drivers that were used during the time of the development of the book might have newer updates available.

Before you install the IBM Linux tape device driver, ensure that the following tasks are completed:

- ▶ The HBA and its device driver are installed.
- ▶ The tape drive and HBA are connected.
- ▶ The tape library is powered on.
- ▶ The tape drives in the library are powered on.
- ▶ The following tape library and tape drives are configured correctly:
  - A logical library was created and configured.
  - Tape drive and control path are assigned to the logical library.
  - Physical media is assigned to the logical library.
- ▶ The Linux server for this installation is powered on.

Complete the following steps to download the packages and documentation to install the IBM tape device driver:

**Important:** The IBM Fix Central portal is the preferred location for IBM tape device driver downloads and should replace any FTP sites that are used for the same purpose.

1. Download the most recent version of the Linux tape device driver from the IBM Fix Central portal by completing the following steps:
  - a. Open the following URL in your web browser:  
<http://www.ibm.com/support/fixcentral>
  - b. Click **Product Group** → **System Storage**.
  - c. Click **Product Family** → **Tape systems**.
  - d. Click **Product Type** → **Tape drivers and software**.
  - e. Click **Product** → **Tape device drivers**.
  - f. Select your operating system from the **Platform** menu.
  - g. Click **Continue**.
  - h. (Optional) Narrow down the search of available downloads according to your criteria.
  - i. Select the version of the IBM tape device driver that you want to download (ensure that you select the correct operating system, the correct release and fix level, and the appropriate support for 32- or 64-bit).
  - j. Follow the instructions on the Fix Central download page to download the new version.

**Important:** See the readme files that are provided within the IBM tape device driver package you downloaded to check for other required software and recently added information that might apply to this specific release.

2. Complete the following steps to download the .pdf document *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130, from the IBM Support portal:
  - a. Open the following URL in your web browser:  
<http://www.ibm.com/support/us/en/>
  - b. Enter the guide name in the **Search support** field and then click the blue button.
  - c. Click **IBM Tape Device Drivers Installation and User's Guide**.
  - d. Click **Device\_Driver\_IUG.pdf** to download the user's guide.
  - e. Follow the procedures in the guide to install the tape device driver.

The following steps show you how to install the IBM Linux tape device driver `lin_tape` (64-bit, for Red Hat Linux Release 7) during our setup of the equipment when writing this book:

1. Download the appropriate level of the `lin_tape` source rpm package to a directory of your choice on the Linux server system for which you want to install it (for more information, see 4.2.2, "Installing the IBM tape device driver" on page 51).
2. Run the `rpmbuild --rebuild <filename>` command where `<filename>` is the name of the source rpm file. This command creates an installable binary rpm package for your specific kernel from the source rpm package, as shown in the following example:

```
rpmbuild --rebuild lin_tape-3.0.3-1.src.rpm
```



```

LD [M] /root/rpmbuild/BUILD/lin_tape-3.0.3/lin_tape.o
Building modules, stage 2.
MODPOST 1 modules
CC      /root/rpmbuild/BUILD/lin_tape-3.0.3/lin_tape.mod.o
LD [M] /root/rpmbuild/BUILD/lin_tape-3.0.3/lin_tape.ko
make[2]: Leaving directory ~/usr/src/kernels/3.10.0-229.e17.x86_64'
make[1]: Leaving directory ~/root/rpmbuild/BUILD/lin_tape-3.0.3'
mv lin_tape.ko bldtmp/lin_tape-3.10.0-229.e17.x86_64.ko
+ exit 0
Executing(%install): /bin/sh -e /var/tmp/rpm-tmp.jejgYx
+ umask 022
+ cd /root/rpmbuild/BUILD
+ '[' /root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64 '!=' / ']'
+ rm -rf /root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64
++ dirname /root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64
+ mkdir -p /root/rpmbuild/BUILDROOT
+ mkdir /root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64
+ cd lin_tape-3.0.3
+ rm -rf /root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64
+ install -D -m 644 bldtmp/lin_tape-3.10.0-229.e17.x86_64.ko
/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64/lib/modules/3.10.0-229.e17.x86_64/kernel/drivers/scsi/lin_tape.ko
++ uname -m
+ PROC=x86_64
+ p=x86_64
+ '[' x86_64 == i386 ']'
+ '[' x86_64 == i586 ']'
+ '[' x86_64 == i686 ']'
+ install -D -m 700 lin_tape
/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64/etc/init.d/lin_tape
+ cp IBM_tape.h /usr/include/sys
+ cd /root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64
+ ln -sf /etc/init.d/lin_tape /usr/sbin/rc_lin_tape
+ cd -
/root/rpmbuild/BUILD/lin_tape-3.0.3
+ /usr/lib/rpm/check-buildroot
+ /usr/lib/rpm/redhat/brp-compress
+ /usr/lib/rpm/redhat/brp-strip /usr/bin/strip
+ /usr/lib/rpm/redhat/brp-strip-comment-note /usr/bin/strip /usr/bin/objdump
+ /usr/lib/rpm/redhat/brp-strip-static-archive /usr/bin/strip
+ /usr/lib/rpm/brp-python-bytecompile /usr/bin/python 1
+ /usr/lib/rpm/redhat/brp-python-hardlink
+ /usr/lib/rpm/redhat/brp-java-repack-jars
Processing files: lin_tape-3.0.3-1.x86_64
Executing(%doc): /bin/sh -e /var/tmp/rpm-tmp.A16QJg
+ umask 022
+ cd /root/rpmbuild/BUILD
+ cd lin_tape-3.0.3
+
DOCDIR=/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64/usr/share/doc/lin_tape-3.0.3
+ export DOCDIR
+ /usr/bin/mkdir -p
/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64/usr/share/doc/lin_tape-3.0.3

```

```

+ cp -pr lin_tape.ReadMe
/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64/usr/share/doc/lin_tape-3.0.3
+ cp -pr lin_tape_daemon.ReadMe
/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64/usr/share/doc/lin_tape-3.0.3
+ cp -pr COPYING
/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64/usr/share/doc/lin_tape-3.0.3
+ cp -pr COPYING.LIB
/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64/usr/share/doc/lin_tape-3.0.3
+ exit 0
Provides: lin_tape = 3.0.3-1 lin_tape(x86-64) = 3.0.3-1
Requires(interp): /bin/sh /bin/sh /bin/sh /bin/sh
Requires(rpmlib): rpmlib(FileDigests) <= 4.6.0-1 rpmlib(PayloadFilesHavePrefix)
<= 4.0-1 rpmlib(CompressedFileNames) <= 3.0.4-1
Requires(pre): /bin/sh
Requires(post): /bin/sh
Requires(preun): /bin/sh
Requires(postun): /bin/sh
Conflicts: IBMtape lin_tape < 3.0.3
Obsoletes: IBMtape
Checking for unpackaged file(s): /usr/lib/rpm/check-files
/root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64
Wrote: /root/rpmbuild/RPMS/x86_64/lin_tape-3.0.3-1.x86_64.rpm
Executing(%clean): /bin/sh -e /var/tmp/rpm-tmp.WpXVHJ
+ umask 022
+ cd /root/rpmbuild/BUILD
+ cd lin_tape-3.0.3
+ rm -rf /root/rpmbuild/BUILDROOT/lin_tape-3.0.3-1.x86_64
+ exit 0
Executing(--clean): /bin/sh -e /var/tmp/rpm-tmp.dJk1Lt
+ umask 022
+ cd /root/rpmbuild/BUILD
+ rm -rf lin_tape-3.0.3
+ exit 0

```

- The output from the binary build is shown. Near the end of the output, there is a line that indicates the file name and location of your binary rpm file on your Linux server (see Example 4-1 on page 53), as shown in the following example:

```
Wrote: /root/rpmbuild/RPMS/x86_64/lin_tape-3.0.3-1.x86_64.rpm
```

- To install the `lin_tape` driver from the binary package as the next step, run the `rpm -ivh <filename>` command against the `.rpm` file that was created in step 2 on page 52, as shown in the following example:

```
rpm -ivh /root/rpmbuild/RPMS/x86_64/lin_tape-3.0.3-1.x86_64.rpm
```

The command output is shown in Example 4-2.

*Example 4-2 Install the binary lin\_tape device driver*

```

# rpm -ivh /root/rpmbuild/RPMS/x86_64/lin_tape-3.0.3-1.x86_64.rpm
Preparing...          ##### [100%]
   1:lin_tape         ##### [100%]
Starting lin_tape:
lin_tape loaded

```

- The final package in this section is to install the `lin_taped` daemon. Download the binary rpm package to your Linux file system. This daemon comes directly as the binary rpm package and has no need for prior compilation with its source rpm package. Run the `rpm -ivh` command on the daemon rpm file, as shown in the following example:

```
rpm -ivh lin_taped-3.0.3-1-rhel7.x86_64.rpm
```

The output of the command is shown in Example 4-3.

*Example 4-3 Installing the binary lin\_taped daemon*

---

```
# rpm -ivh lin_taped-3.0.3-1-rhel7.x86_64.rpm
Preparing...          ##### [100%]
 1:lin_taped          ##### [100%]
Starting lin_tape:
lin_tape loaded
```

---

- Finally, you can confirm the successful installation and list the found and configured IBM tape libraries and tape device by running the following Linux command:

```
cat /proc/scsi/IBM*
```

Example 4-4 shows you an output of a Linux system that had many tape devices attached to it and the tape library was partitioned to several logical partitions.

*Example 4-4 Run the Linux cat command to display the attached tape devices and libraries*

---

```
# cat /proc/scsi/IBM*
lin_tape version: 3.0.3
lin_tape major number: 245
Attached Changer Devices:
Number  model      S/N          HBA          SCSI          FO Path
0       3573-TL    00L4U78A5815_LL0  qla2xxx      2:0:0:1       NA
1       3573-TL    00L4U78A5815_LL0  qla2xxx      2:0:1:1       NA
2       03584L22   0000078A42740401  qla2xxx      2:0:3:1       NA
3       3573-TL    00L4U78B1694_LL1  qla2xxx      2:0:4:1       NA
4       3573-TL    00L4U78B1694_LL0  qla2xxx      2:0:5:1       NA
5       03584L32   0000078A42740407  qla2xxx      2:0:6:1       NA
6       03584L22   0000078A42740401  qla2xxx      2:0:7:1       NA
7       03584L22   0000078A42740403  qla2xxx      2:0:8:1       NA
8       03584L22   0000078A42740403  qla2xxx      2:0:9:1       NA
9       03584L32   0000078A42740409  qla2xxx      2:0:11:1      NA
10      03584L32   0000078A42740407  qla2xxx      2:0:12:1     NA
11      03584L32   0000078A42740409  qla2xxx      2:0:13:1     NA
12      3573-TL    00L4U78C6680_LL0  qla2xxx      2:0:14:1     NA
13      3573-TL    00L4U78C6680_LL0  qla2xxx      2:0:16:1     NA
14      03584L32   0000078A42740405  qla2xxx      2:0:19:1     NA
15      03584L22   0000078A42740402  qla2xxx      2:0:20:1     NA
16      03584L32   0000078A42740405  qla2xxx      2:0:21:1     NA
17      03584L22   0000078A42740402  qla2xxx      2:0:22:1     NA
18      03584L22   0000078A42740401  qla2xxx      2:0:23:1     NA
19      03584L22   0000078A42740401  qla2xxx      2:0:24:1     NA
20      03584L22   0000078A42740401  qla2xxx      2:0:25:1     NA
lin_tape version: 3.0.3
lin_tape major number: 245
Attached Tape Devices:
Number  model      S/N          HBA          SCSI          FO Path
0       ULT3580-TD5 00078A218E    qla2xxx      2:0:0:0       NA
1       ULT3580-TD5 1168001144    qla2xxx      2:0:1:0       NA
```

2	ULT3580-TD6	1013000068	q1 a2xxx	2:0:2:0	NA
3	03592E07	0000013D0485	q1 a2xxx	2:0:3:0	NA
4	ULT3580-TD5	00078A1D8F	q1 a2xxx	2:0:4:0	NA
5	ULTRIUM-TD5	9A700L0077	q1 a2xxx	2:0:5:0	NA
6	ULT3580-TD6	00013B0037	q1 a2xxx	2:0:6:0	NA
7	03592E07	001013000652	q1 a2xxx	2:0:7:0	NA
8	03592E07	0000078DDAC3	q1 a2xxx	2:0:8:0	NA
9	03592E07	001013000255	q1 a2xxx	2:0:9:0	NA
10	03592E07	0000013D0734	q1 a2xxx	2:0:10:0	NA
11	ULT3580-TD5	1068000073	q1 a2xxx	2:0:11:0	NA
12	ULT3580-TD6	00013B0084	q1 a2xxx	2:0:12:0	NA
13	ULTRIUM-TD5	9A700M0029	q1 a2xxx	2:0:13:0	NA
14	ULT3580-TD5	1068000070	q1 a2xxx	2:0:14:0	NA
15	ULT3580-TD5	1068000016	q1 a2xxx	2:0:16:0	NA
16	ULT3580-TD5	00078AC0A5	q1 a2xxx	2:0:19:0	NA
17	03592E07	0000013D0733	q1 a2xxx	2:0:20:0	NA
18	ULT3580-TD5	00078AC08B	q1 a2xxx	2:0:21:0	NA
19	03592E07	0000078DDBF1	q1 a2xxx	2:0:22:0	NA
20	03592E07	0000013D0483	q1 a2xxx	2:0:23:0	NA
21	03592E07	0000013D0485	q1 a2xxx	2:0:24:0	NA
22	03592E07	0000078D13C1	q1 a2xxx	2:0:25:0	NA

**Linux cat command:** You will use the Linux `cat` command again in a later section to identify the tape devices and medium changer you use for the setup of the LTFs EE installation and configuration.

## 4.3 Installing IBM Spectrum Archive EE

This section describes the process that is used to install IBM Spectrum Archive EE package. This installation package is provided by IBM to you on a DVD.

Consider the following important notes:

- ▶ Information that is contained in the readme file and installation file that are provided with the IBM Spectrum Archive EE distribution package supersedes information that is presented in this book and the online IBM Knowledge Center.
- ▶ The IBM Spectrum Archive EE license does not entitle customers to use any other IBM Tivoli® Storage Manager components or products. All components that are needed to migrate data to the LTFs file space are integrated into IBM Spectrum Archive EE. They also are part of the provided installation package and IBM Spectrum Archive EE license and are to be used only in this context.
- ▶ If IBM Spectrum Archive LE is already installed, it must be uninstalled *before* IBM Spectrum Archive EE is installed.

If IBM Spectrum Archive EE is installed (such as an older version), it must be uninstalled *before* the latest version of IBM Spectrum Archive EE is installed. To update the EE package, an installation script (`ltfsee_install`) is provided that does the automatic uninstallation during the software update. The next sections show you how to use this IBM Spectrum Archive EE installation script for different purposes and maintenance.

It also is possible to install, upgrade, or uninstall IBM Spectrum Archive EE manually. For more information, see 11.4, “IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products” on page 331.

### 4.3.1 Extracting binary rpm files from an installation package

This first task lists the necessary steps to perform before binary rpm files are extracted. It also presents the available methods for extracting binary rpm files from an installation package for IBM Spectrum Archive EE on a Linux server system.

Interactive console mode is the method that is used for extracting binary rpm files from an installation package.

Before you use any of these methods to extract the IBM Spectrum Archive EE binary rpm files, you must confirm or set the run permission of the installation package.

**Important:** The `1tfsee-1.2.4.0-[revision].[architecture].bin` installation package includes rpm files for the revision and supported platforms.

Before the IBM Spectrum Archive EE binary rpm files are extracted from an installation package, complete the following steps:

1. Confirm the run permission of `1tfsee-1.2.4.0-[revision].[architecture].bin` by running the following command:  

```
# ls -l 1tfsee-1.2.4.0-[revision].[architecture].bin
```
2. If it is not already set, set the run permission by running the following command:  

```
# chmod +x 1tfsee-1.2.4.0-[revision].[architecture].bin
```
3. Proceed with the extraction of the binary IBM Spectrum Archive EE rpm files by selecting one of the procedures that are described next.

In the lab setup that was used for this book, we used the interactive console mode method, which is the option most users are likely to use.

#### Extracting binary rpm files in interactive console mode

This section describes how to extract binary rpm files from the IBM Spectrum Archive EE installation package by using the interactive console mode.

**Important:** The steps in this section extract binary rpm files to your local disk only. To complete the installation process, more steps are required. After you complete the extraction of the binary rpm files, see “Installing, upgrading, or uninstalling IBM Spectrum Archive EE automatically” on page 61 or 4.5, “Quick installation guide for IBM Spectrum Archive EE” on page 71 for more information.

To extract IBM Spectrum Archive EE binary rpm files in interactive console mode, complete the following steps:

1. Run IBM Spectrum Archive EE on the system by running the appropriate command for your environment:
  - If your operating system is running on the command-line interface (CLI), run the following command:  

```
#./1tfsee-1.2.4.0-[revision].[architecture].bin
```
  - If your operating system is running on the GUI (X Window System), run the following command:  

```
#./1tfsee-1.2.4.0-[revision].[architecture].bin -i console
```

The messages in Example 4-5 are displayed.

*Example 4-5 Extract binary rpm files in interactive console mode*

---

```
Preparing to install...
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...
```

```
Launching installer...
```

---

**Important:** You cannot select the installation and link folders with the console installer. They are created in the `~/LTFSEE/` directory, which is the default folder of the installer that extracts the required files.

The installation script `ltfsee_install` for the command-line installation is found under the `~/LTFSEE/rpm.[revision]` folder, for example, the `~/LTFSEE/rpm.1220_10326/` subfolder.

2. Read the International Program License Agreement. Enter 1 to accept the agreement and press **Enter** to continue, as shown in Example 4-6.

*Example 4-6 IBM Spectrum Archive EE International Program License Agreement*

---

```
=====
IBM Spectrum Archive Enterprise Edition (created with InstallAnywhere)
=====
```

```
Preparing CONSOLE Mode Installation...
```

```
=====
International License Agreement for Early Release of Programs
```

```
Part 1 - General Terms
```

```
BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON AN
"ACCEPT" BUTTON, OR OTHERWISE USING THE PROGRAM, LICENSEE AGREES TO
THE TERMS OF THIS AGREEMENT. IF YOU ARE ACCEPTING THESE TERMS ON
BEHALF OF LICENSEE, YOU REPRESENT AND WARRANT THAT YOU HAVE FULL
AUTHORITY TO BIND LICENSEE TO THESE TERMS. IF YOU DO NOT AGREE TO
THESE TERMS,
```

```
* DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, CLICK ON AN "ACCEPT" BUTTON,
OR USE THE PROGRAM; AND
```

```
* PROMPTLY RETURN THE UNUSED MEDIA, DOCUMENTATION, AND PROOF OF
ENTITLEMENT TO THE PARTY FROM WHOM IT WAS OBTAINED FOR A REFUND OF THE
AMOUNT PAID. IF THE PROGRAM WAS DOWNLOADED, DESTROY ALL COPIES OF THE
PROGRAM.
```

```
1. Definitions
```

```
Press Enter to continue viewing the license agreement, or enter "1" to
accept the agreement, "2" to decline it, "3" to print it, or "99" to go back
to the previous screen.:
```

---

3. An Installing... message displays while the files are extracted to the ~/LTFSEE/ installation folder, as shown in Example 4-7. You can monitor the progress by watching the text-animated progress bars.

*Example 4-7 IBM Spectrum Archive EE installation of the binary files*

```

=====
Installing...
-----

[-----|-----|-----|-----]
[-----|-----|-----|-----]

```

When the files are successfully extracted, the text-based installer completes.

**Important:** The following symbolic links are created in your home directory:

- ▶ A link to the rpm folder that keeps the extracted rpm files.
- ▶ A link to the “Change IBM Linear Tape File System Enterprise Edition Installation” executable file that uninstalls IBM Spectrum Archive EE.

4. Go to the rpm.[version]\_[revision] folder under the installation folder to locate the rpm folder and access the installation file. If you created symbolic links, click the rpm symbolic link or use the Linux operating system `cd ~/rpm` command to open the rpm folder.

**Important:** Two files, `INSTALL_EE.[revision]` and `README_EE.[revision]`, are placed in the rpm folder. Folders that correspond to the supported platforms are created in the rpm folder as well. The specific rpm files for the supported platform are in the platform subdirectory.

When you successfully finish, continue to “Installing, upgrading, or uninstalling IBM Spectrum Archive EE automatically” on page 61 to complete the installation. If you prefer to install manually, see 4.5, “Quick installation guide for IBM Spectrum Archive EE” on page 71.

### 4.3.2 Installing, upgrading, or uninstalling IBM Spectrum Archive EE

This section describes how to install, upgrade, or uninstall binary rpm files for IBM Spectrum Archive EE after extracting them from the installation package, as described in 4.3.1, “Extracting binary rpm files from an installation package” on page 58.

IBM Spectrum Archive EE can be automatically installed, upgraded, or uninstalled.

IBM Spectrum Archive EE nodes communicate by using several TCP/UDP ports. Because some ports are assigned dynamically within a wide range, you must disable any firewall program that runs on these nodes.

**Important:** On Red Hat Enterprise Linux 7, managing system services is done with the `systemctl` command.

To stop the firewall service, run the `systemctl stop firewalld` command.

To prevent the firewall service from being automatically started at start time, run the `systemctl disable firewalld` command.

In addition, you can mask the firewalld service to prevent it from being started manually or by another service by running the `systemctl mask firewalld` command.

These commands should be run on every power cycle of the server.

During this installation for the IBM Spectrum Archive EE rpm files, there is also an MIB file that is provided if you plan to use SNMP for monitoring of your IBM Spectrum Archive EE setup. SNMP monitoring software usually requires such an MIB file to manage the various SNMP traps sent to it. The IBM Spectrum Archive EE MIB file is in the `/opt/ibm/ltfsee/data/IBMSA-MIB.txt` directory.

## Installing, upgrading, or uninstalling IBM Spectrum Archive EE automatically

This section describes how to install, upgrade, or uninstall binary rpm files for IBM Spectrum Archive EE automatically after extracting them from the installation package. We used this method during our lab setup to write this book and document the examples.

The automated method is based on a utility (a shell script), which is provided by the IBM Spectrum Archive EE installation package. The script is named `ltfsee_install` and can be found after extracting the binary installation files in the `~/LTFSEE/rpm.[revision]` directory with the IBM Spectrum Archive EE rpm files (such as `/root/LTFSEE/rpm.1220_10326/`).

The following prerequisite conditions must be met before you proceed with this task:

- ▶ Required software:
  - For more information about installing the required software on a Linux system, see 4.2.1, “Installing the host bus adapter and device driver” on page 51.
  - For more information about installing the IBM tape device driver, see 4.2.2, “Installing the IBM tape device driver” on page 51.
- ▶ IBM Spectrum Scale settings:
  - The node must belong to an IBM Spectrum Scale cluster.
  - The GPFS daemons must be started.
  - For information about how to create an IBM Spectrum Scale cluster, see the IBM Knowledge Center for IBM Spectrum Scale at:

[https://www.ibm.com/support/knowledgecenter/en/STXKQY\\_4.2.3/com.ibm.spectrum.scale.v4r23.doc/b1lins\\_estart.htm](https://www.ibm.com/support/knowledgecenter/en/STXKQY_4.2.3/com.ibm.spectrum.scale.v4r23.doc/b1lins_estart.htm)

### ***ltfsee\_install*** utility

Use the `ltfsee_install` command-line utility to install rpm packages automatically to the system. You must have root user authority to use this command.

For more information, see “Installing, upgrading, or uninstalling IBM Spectrum Archive EE automatically” on page 61.

The `ltfsee_install <option>` command installs the following rpm packages to the system:

- ▶ IBM Spectrum Archive LE+ component
- ▶ Integrated customized Tivoli Storage Manager for Space Management with IBM Spectrum Archive EE
- ▶ IBM Spectrum Archive Migration Driver

The command includes the following options:

- ▶ **--install**  
Install rpm packages. If rpm packages are already installed, the installation is stopped.
- ▶ **--upgrade**  
Upgrade installed rpm packages.
- ▶ **--clean**  
Uninstall rpm packages.
- ▶ **--verify**  
Verify the prerequisite conditions and IBM Spectrum Archive Enterprise Edition package installation only. No installation will be performed.
- ▶ **--check**  
Check the prerequisite conditions only. No installation will be performed.

Verify that these conditions are met by logging on to the operating system as a root user and running the following command:

```
# ./ltfsee_install --check
```

If the conditions are met, the following message is shown as the last line of an output to the screen:

The prerequisites checking is completed successfully.

Example 4-8 shows the complete output.

*Example 4-8 Screen output for the `ltfsee_install --check` command*

---

```
./ltfsee_install --check  
Checking rpm installation and version.
```

The prerequisites checking is completed successfully.

---

The `ltfsee_install` file installs or upgrades all required rpm packages on the server node. It can also uninstall those rpm packages from the node if needed.

**Important:** The `ltfsee_install` command procedures in this topic automatically perform all operations from 4.5, “Quick installation guide for IBM Spectrum Archive EE” on page 71, except for installing optional TIVsm language packages in case they are needed.

Complete the following steps to automatically install, upgrade, or uninstall IBM Spectrum Archive EE by running the `ltfsee_install` command:

1. Log on to the operating system as a root user.
2. On each node in the cluster, complete the set of tasks for the action you want to take:
  - a. Installing IBM Spectrum Archive EE on the node:

- i. Run the following command:

```
# ./ltfsee_install --install
```

Example 4-9 shows you the complete output of the `ltfsee_install --install` command.

- ii. Verify that the command completed successfully. Check for the following success message in the command output:

All rpm packages are installed successfully.

*Example 4-9 Output for the ltfsee\_install --install command*

---

```
# ./ltfsee_install --install
Checking rpm installation and version.

The prerequisites checking completed successfully.

Preparing... #####
[100%]
 1:ltfsle #####
[100%]
/sbin/ldconfig: /lib/ld-linux.so.2 is not a symbolic link

Preparing... #####
[100%]
 1:ltfsle-library #####
[100%]
/sbin/ldconfig: /lib/ld-linux.so.2 is not a symbolic link

Preparing... #####
[100%]
 1:ltfsle-library-plus #####
[100%]
/sbin/ldconfig: /lib/ld-linux.so.2 is not a symbolic link

Preparing... #####
[100%]
 1:gskcrypt64 #####
Preparing... #####
[100%]
 1:gskssl64 #####
Preparing... #####
[100%]
 1:TIVsm-API64 #####
Preparing... #####
[100%]
```

```

1:TIVsm-BA #####
[100%]
Preparing... #####
[100%]
1:TIVsm-HSM #####
[100%]
Deactivating failover operations on the node.
-----
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
Client Version 7, Release 1, Level 6.3
Client date/time: 12/08/2016 19:47:28
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights
Reserved.

TSM The node failover option is now disabled.
-----

Stopping the space management daemon.
-----
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
Client Version 7, Release 1, Level 6.3
Client date/time: 12/08/2016 19:47:35
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights
Reserved.

-----

Stopping the HSM service.
-----
hsm: unrecognized service
-----

Terminated the space management daemon.
Preparing... #####
[100%]
1:ltfs-mig #####
[100%]
ldconfig: /lib/ld-linux.so.2 is not a symbolic link

Starting the HSM service.
-----
hsm: unrecognized service
-----

Starting the space management daemon.
-----
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
Client Version 7, Release 1, Level 6.3
Client date/time: 12/08/2016 19:47:48
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights
Reserved.

```

-----  
IBM Tivoli Storage Manager  
Command Line Space Management Client Interface  
Client Version 7, Release 1, Level 6.3  
Client date/time: 12/08/2016 19:47:49  
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

TSM The node failover option is now enabled in mode ENABLED.  
All rpm packages are installed successfully.  
Complete the configuration using `/opt/ibm/ltfsee/bin/ltfsee_config` command.

---

- iii. Complete the configuration by running the `/opt/ibm/ltfsee/bin/ltfsee_config` command, as described in 6.2, “Configuring IBM Spectrum Archive EE” on page 104.
- b. Upgrading the rpm files to the latest versions:

**Note:** To avoid any unexpected failures during a multi node system upgrade, run `ltfsee_install --upgrade` on one node at a time and wait until the process completes before running the upgrade on the next node.

- i. Run the `ltfsee stop` command.
- ii. Run the `pidof mmm` command and wait until there are no processes returned.
- iii. Run the `pidof ltfs` command on every EE node and wait until there are no processes returned.

To perform the upgrade process, run the command shown in Example 4-10.

*Example 4-10 Running the pidof ltfs command*

---

```
# ./ltfsee_install --upgrade  
[root@ltfsm11 rpm.1220_10326]# ./ltfsee_install --upgrade
```

Force installation option is selected.

This option uninstalls the following preinstalled packages before installation:

- LE+ component
- HSM component
- LTFS Migration Driver

These packages will be upgraded to the version that IBM Spectrum Archive Enterprise Edition provides.

Do you want to continue? (y/n) y

---

- iv. Verify that the command completed. Check for the following success message in the command output:  
All rpm packages are installed successfully.

- c. Uninstalling IBM Spectrum Archive EE from the node:
  - i. Run the following command:
 

```
# ltfsee_install --clean
```
  - ii. Verify that the command completed. Check for the following success message in the command output:
 

```
Uninstallation is completed.
```
3. Verify that the installation or uninstallation completed successfully by running the following command:
 

```
# ltfsee_install --verify
```

If the installation was successful, the following message is shown:

```
Module installation check is completed.
```

## 4.4 Installing a RESTful server

This section describes how to install the IBM Spectrum Archive EE REST API after extracting the installation package, as described in 4.3.1, “Extracting binary rpm files from an installation package” on page 58 and installing IBM Spectrum Archive EE as described in 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60. This section should be followed after IBM Spectrum Archive EE has been installed.

The rest service can be installed on any node within the cluster that has IBM Spectrum Archive EE installed. To start the installation, some software requirements must be met before the rpm can be installed.

The following are the required software applications:

- ▶ IBM Spectrum Archive V1.2.4 or later
- ▶ httpd
- ▶ mod\_ssl
- ▶ mod\_wsgi
- ▶ Python 2.4 or later, but earlier than 3.0
- ▶ Flask

Example 4-11 shows how to perform the installation of the required software to run REST.

*Example 4-11 Required software for rest installation*

---

```
[root@l1tfseesrv1 ~]# yum install -y httpd mod_ssl mod_wsgi
Loaded plugins: langpacks, product-id, rhnplugin, search-disabled-repos,
subscription-manager
This system is receiving updates from RHN Classic or Red Hat Satellite.
Resolving Dependencies
--> Running transaction check
----> Package httpd.x86_64 0:2.4.6-45.e17 will be installed
--> Processing Dependency: httpd-tools = 2.4.6-45.e17 for package:
httpd-2.4.6-45.e17.x86_64
--> Processing Dependency: /etc/mime.types for package: httpd-2.4.6-45.e17.x86_64
----> Package mod_ssl.x86_64 1:2.4.6-45.e17 will be installed
----> Package mod_wsgi.x86_64 0:3.4-12.e17_0 will be installed
--> Running transaction check
----> Package httpd-tools.x86_64 0:2.4.6-45.e17 will be installed
----> Package mailcap.noarch 0:2.1.41-2.e17 will be installed
```

--> Finished Dependency Resolution

Dependencies Resolved

```
=====
=====
Package           Arch          Size          Version
Repository
=====
Installing:
  httpd            x86_64        1.2 M         2.4.6-45.e17
rhel-x86_64-server-7
  mod_ssl          x86_64        105 k         1:2.4.6-45.e17
rhel-x86_64-server-7
  mod_wsgi         x86_64        76 k          3.4-12.e17_0
rhel-x86_64-server-7
Installing for dependencies:
  httpd-tools      x86_64        84 k          2.4.6-45.e17
rhel-x86_64-server-7
  mailcap          noarch        31 k          2.1.41-2.e17
rhel-x86_64-server-7
=====
```

Transaction Summary

```
=====
=====
Install 3 Packages (+2 Dependent packages)
=====
```

```
Total download size: 1.5 M
Installed size: 4.4 M
Downloading packages:
(1/5): httpd-2.4.6-45.e17.x86_64.rpm
| 1.2 MB 00:00:00
(2/5): httpd-tools-2.4.6-45.e17.x86_64.rpm
| 84 kB 00:00:00
(3/5): mailcap-2.1.41-2.e17.noarch.rpm
| 31 kB 00:00:00
(4/5): mod_ssl-2.4.6-45.e17.x86_64.rpm
| 105 kB 00:00:00
(5/5): mod_wsgi-3.4-12.e17_0.x86_64.rpm
| 76 kB 00:00:00
=====
```

```
-----
Total
467 kB/s | 1.5 MB 00:00:03
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
Warning: RPMDB altered outside of yum.
Installing : mailcap-2.1.41-2.e17.noarch
1/5
Installing : httpd-tools-2.4.6-45.e17.x86_64
2/5
```

```

Installing : httpd-2.4.6-45.e17.x86_64
3/5
Installing : mod_wsgi-3.4-12.e17_0.x86_64
4/5
Installing : 1:mod_ssl-2.4.6-45.e17.x86_64
5/5
Verifying : httpd-tools-2.4.6-45.e17.x86_64
1/5
Verifying : mod_wsgi-3.4-12.e17_0.x86_64
2/5
Verifying : mailcap-2.1.41-2.e17.noarch
3/5
Verifying : 1:mod_ssl-2.4.6-45.e17.x86_64
4/5
Verifying : httpd-2.4.6-45.e17.x86_64
5/5

```

Installed:

```

httpd.x86_64 0:2.4.6-45.e17          mod_ssl.x86_64 1:2.4.6-45.e17
mod_wsgi.x86_64 0:3.4-12.e17_0

```

Dependency Installed:

```

httpd-tools.x86_64 0:2.4.6-45.e17          mailcap.noarch
0:2.1.41-2.e17

```

Complete!

---

If pip is not installed on the designated node (pip is installed by default if the version of Python is 2.7.9 or greater), it can be installed by running the following commands:

```

curl "https://bootstrap.pypa.io/get-pip.py" -o "get-pip.py"
python get-pip.py

```

After pip is installed, run the following command to install flask:

```

pip install flask

```

Example 4-12 shows how to install flask.

*Example 4-12 Install flask V0.12*

---

```

[root@ltfseesrv1 ~]# pip install flask
Collecting flask
  Downloading Flask-0.12.2-py2.py3-none-any.whl (83kB)
    100% |#####| 92kB 513kB/s
Collecting click>=2.0 (from flask)
  Downloading click-6.7-py2.py3-none-any.whl (71kB)
    100% |#####| 71kB 1.4MB/s
Collecting Jinja2>=2.4 (from flask)
  Downloading Jinja2-2.9.6-py2.py3-none-any.whl (340kB)
    100% |#####| 348kB 1.0MB/s
Collecting Werkzeug>=0.7 (from flask)
  Downloading Werkzeug-0.12.2-py2.py3-none-any.whl (312kB)
    100% |#####| 317kB 1.3MB/s
Collecting itsdangerous>=0.21 (from flask)
  Downloading itsdangerous-0.24.tar.gz (46kB)
    100% |#####| 51kB 5.0MB/s

```

```

Collecting MarkupSafe>=0.23 (from Jinja2>=2.4->flask)
  Downloading MarkupSafe-1.0.tar.gz
Building wheels for collected packages: itsdangerous, MarkupSafe
  Running setup.py bdist_wheel for itsdangerous ... done
  Stored in directory:
  /root/.cache/pip/wheels/fc/a8/66/24d655233c757e178d45dea2de22a04c6d92766abfb741129
  a
  Running setup.py bdist_wheel for MarkupSafe ... done
  Stored in directory:
  /root/.cache/pip/wheels/88/a7/30/e39a54a87bcbe25308fa3ca64e8ddc75d9b3e5afa21ee32d5
  7
Successfully built itsdangerous MarkupSafe
Installing collected packages: click, MarkupSafe, Jinja2, Werkzeug, itsdangerous,
flask
  Found existing installation: MarkupSafe 0.11
  Uninstalling MarkupSafe-0.11:
    Successfully uninstalled MarkupSafe-0.11
Successfully installed Jinja2-2.9.6 MarkupSafe-1.0 Werkzeug-0.12.2 click-6.7
flask-0.12.2 itsdangerous-0.24

```

---

After all the required software has been installed, in the same directory that the IBM Spectrum Archive EE was extracted to, there is an RHEL7 directory that contains a file called `.ibmsa-rest-[build]-[revision].x86_64.rpm`. To install the restful service, run the **yum install** command on this file, as outlined in Example 4-13.

*Example 4-13 Installing IBM Spectrum Archive Rest service*

---

```

[root@l1fseesrv1 RHEL7]# yum install -y ibmsa-rest-1.2.4.0-12441.x86_64.rpm
Loaded plugins: langpacks, product-id, rhnplugin, subscription-manager
This system is receiving updates from RHN Classic or Red Hat Satellite.
Examining ibmsa-rest-1.2.4.0-12441.x86_64.rpm: ibmsa-rest-1.2.4.0-12441.x86_64
Marking ibmsa-rest-1.2.4.0-12441.x86_64.rpm to be installed
Resolving Dependencies
--> Running transaction check
---> Package ibmsa-rest.x86_64 0:1.2.4.0-12441 will be installed
--> Finished Dependency Resolution

```

Dependencies Resolved

```

=====
Package            Arch      Version      Repository
Size
=====
Installing:
  ibmsa-rest        x86_64    1.2.4.0-12441
/ibmsa-rest-1.2.4.0-12441.x86_64      51 k

```

Transaction Summary

```

=====
Install 1 Package

Total size: 51 k
Installed size: 51 k

```

```

Downloading packages:
Running transaction check
Running transaction test
Transaction test succeeded
Running transaction
  Installing : ibmsa-rest-1.2.4.0-12441.x86_64
1/1
#####
# ibmsa-rest have been installed successfully.
# Please restart or reload httpd to enable REST server.
#####
  Verifying : ibmsa-rest-1.2.4.0-12441.x86_64
1/1

Installed:
  ibmsa-rest.x86_64 0:1.2.4.0-12441

```

Complete!

---

At the bottom of a successful installation, it says the installation was successful and that a restart of the httpd service is required to enable the rest server. To restart the service, run the following command:

```
systemctl restart httpd
```

When this is all done, to quickly test that the rest service has been successfully installed, run the following command:

```
curl -i -X GET 'http://localhost:7100/ibmsa/v1'
```

Example 4-14 shows using a **test curl** command to see whether the installation was successful.

*Example 4-14 Test curl command*

---

```

[root@l1fseesrv1 ~]# curl -i -XGET 'http://localhost:7100/ibmsa/v1'
HTTP/1.1 200 OK
Date: Wed, 12 Jul 2017 22:03:00 GMT
Server: Apache/2.4.6 (Red Hat Enterprise Linux) OpenSSL/1.0.1e-fips mod_wsgi/3.4
Python/2.7.5
Content-Length: 83
Content-Type: application/json

{"message":"IBM Spectrum Archive REST API server is working.,"status_code":"200"}

```

---

The default port is on 7100 and the default protocol to use is http. If SSL is required, uncomment `SSLEngine`, `SSLCertificateFile`, `SSLCertificateKeyFile` and provide the direct path to both the certificate file and the certificate key file in the following file:

```
/etc/httpd/conf.d/ibmsa-rest-httpd.conf
```

For an overview of IBM Spectrum Archive EE Rest API and commands, see 7.29, “IBM Spectrum Archive REST API” on page 215.

## 4.5 Quick installation guide for IBM Spectrum Archive EE

This section summarizes the overall installation process of IBM Spectrum Archive EE:

1. Ensure that all prerequisite software (Linux packages) and HBA driver are installed.
2. Ensure that the IBM Spectrum Scale file system daemon is started by running the following command:

```
mmstartup -a
```

3. Install the IBM tape device driver `lin_tape`.

Download the appropriate level of the source rpm package for IBM `lin_tape` and the `lin_tape` daemon rpm package from this website:

<http://www.ibm.com/support/fixcentral>

To install the `lin_tape` device driver, run the following commands:

```
- rpmbuild --rebuild lin_tape-3.0.3-1.src.rpm
- rpm -ivh /root/rpmbuild/RPMS/x86_64/lin_tape-3.0.3-1.x86_64.rpm
- rpm -ivh lin_taped-lin_tape-3.0.3-1.rhel6.x86_64.rpm
```

4. Extract the binary rpm files from an IBM Spectrum Archive EE installation package by running the following command:

```
./ltfsee-1.2.2.2-10326.x86_64.bin
```

5. Install IBM Spectrum Archive EE automatically by using the `ltfsee_install` tool. Use `--check` for pre-installation check, `--install` for the installation, and `--verify` for a postinstallation verification, as shown in the following commands:

```
- ~/LTFSEE/rpm.[revision]/ltfsee_install --check
- ~/LTFSEE/rpm.[revision]/ltfsee_install --install
- ~/LTFSEE/rpm.[revision]/ltfsee_install --verify
```

## 4.6 Library replacement

Library replacement has been introduced to give users the ability to upgrade their older tape library to a newer tape library. There are two library replacement scenarios available for users to perform: a complete library replacement, and a pool relocation. Both are disruptive, but the pool relocation offers less down time and the ability to continue running IBM Spectrum Archive while the pool that is going to be relocated is disabled.

The library replacement procedure requires the user to halt their environment to perform, due to the relocation of all tape cartridges and possibly all tape drives to a newer library.

The pool relocation procedure requires a multi-library configuration so that a pool can be relocated from one library to the other. While the pool is being relocated operations are not allowed on that pool however the other pools are still available for operations.

In Addition to these new procedures, IBM Spectrum Archive EE now appends a new state for tape cartridges, appendable. This new state is determined by the pool settings and the state of the tape to determine if the tape is candidate for migration. An appendable tape cartridge within a pool allows data to be written to that tape. Those tapes can either be empty, or partially filled with data.

Non-appendable tape cartridges are tapes that do not allow writes to the tape. Non-appendable tape cartridges result from being completely full, erroneous tapes, or tapes within pools that do not match their media-restriction or format type. With the introduction of this new tape state, users have better control over the flow of their data from disk to tape and from tape to tape. Refer to 7.10.5, “Data Migration” on page 173.

**Note:** The following library replacement procedures are supported with IBM Spectrum Archive EE v1.2.6 and above.

## 4.6.1 Library replacement procedure

This procedure should be used if you are switching or upgrading your current tape library to another tape library and want to continue using the tapes and tape drives and cluster configuration within your IBM Spectrum Archive EE environment. For example, you might be replacing your existing TS3500 tape library with a new TS4500 tape library. To do so, complete these steps:

1. Identify the library name and serial number that you intend to replace, for use in the subsequent steps. Example 4-15 shows how to get a list of libraries and serial numbers by running the `ltfsee_config -m LIST_LIBRARIES` command.

*Example 4-15 ltfsee\_config -m LIST\_LIBRARIES*

---

```
[root@server1 ~]# ltfsee_config -m LIST_LIBRARIES
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
LIST_LIBRARIES
LIST_LIBRARIES mode starts .
```

```
## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file
ltfsee_config.filesystem exists.
```

Library Name	Serial Number	Control Node
lib1	0000013FA0520401	9.11.244.22
lib2	0000013FA052041F	9.11.244.23

---

2. If the tape drives from the old library will not be used in the new library, remove the tape drives by using the `ltfsee drive remove` command. Refer to “The ltfsee drive command” on page 302 for more information about the command.
3. Stop IBM Spectrum Archive EE by executing the `ltfsee stop` command.
4. Replace the physical old tape library with the new tape library.
5. Physically move all of the tape cartridges to the new library.
6. If you are moving the tape drives from the old tape library to the new library, physically remove the tape drives from the old tape library and install them in the new library. The drives need to be attached to the same node that they were attached to before.
7. The `REPLACE_LIBRARY` command associates the new library’s serial number with the old library’s ID. Example 4-16 shows the output from running the `ltfsee_config -m REPLACE_LIBRARY` command to associate the new library with the old library serial number.

*Example 4-16 ltfsee\_config -m REPLACE\_LIBRARY*

---

```
[root@server1 ~]# ltfsee_config -m REPLACE_LIBRARY
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
REPLACE_LIBRARY
```

REPLACE\_LIBRARY mode starts .

```
## 1. Check to see if the cluster is already created ##  
The cluster is already created and the configuration file  
ltfsee_config.filesystem exists.
```

\*\* Configuration change for library: lib1

```
The number of logical libraries with the assigned control node: 2  
The number of logical libraries available from this node: 1  
The number of logical libraries available from this node and with assigned  
control node: 0
```

\*\* Select the tape library from the following list  
and input the corresponding number. Then press Enter.

```
      Model      Serial Number  
1. 03584L32    0000013FA052041B  
q. Exit from this Menu
```

```
Input Number > 1  
is going to be set to library lib1.  
Do you want to continue (y/n)?  
Input > y
```

Restarting HSM daemon on server1

```
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m  
RESTART_HSM  
RESTART_HSM mode starts .
```

```
## 1. Check to see if the cluster is already created ##  
The cluster is already created and the configuration file  
ltfsee_config.filesystem exists.
```

```
Deactivating failover operations on the node.  
Restarting Space Management service...  
Stopping the HSM service.  
Terminating dsmwatchd.....  
Starting the HSM service.  
Starting dsmmigfs.....  
Activating failover operations on the node.
```

Restarting HSM daemon on server2

```
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m  
RESTART_HSM  
RESTART_HSM mode starts .
```

```
## 1. Check to see if the cluster is already created ##  
The cluster is already created and the configuration file  
ltfsee_config.filesystem exists.
```

```
Deactivating failover operations on the node.  
Restarting Space Management service...  
Stopping the HSM service.  
Terminating dsmwatchd.....  
Starting the HSM service.
```

```
Starting dsmmigfs.....  
Activating failover operations on the node.
```

```
REPLACE_LIBRARY mode completed .
```

---

**Note:** In releases before IBM Spectrum Archive EE version 1.2.6, the serial number of a library was used as the library ID. Beginning with version 1.2.6, the library ID can have a unique value other than the serial number. A UUID is assigned as the library ID in libraries that are configured with version 1.2.6 and subsequent releases.

With the separation of the library ID from the library serial number, IBM Spectrum Archive EE can replace a library by changing the library serial number that is associated with the library ID. The configuration of the cartridges, and pools, and drives in the old library is transferred to the new library.

8. Verify library serial change by running the `ltfsee_config -m LIST_LIBRARIES`.
9. Start IBM Spectrum Archive EE by using the `ltfsee start` command.
10. If you are using new drives in the new library, configure the drives with the `ltfsee drive add` command. See “The ltfsee drive command” on page 302 for more information about the command.

## 4.6.2 Pool relocation procedure

Use this procedure to logically and physically move a tape pool from one library (the source) to another library (the destination), within an IBM Spectrum Archive EE cluster. During the procedure, the tape pool that is being moved is disabled. However, during most of the procedure, the IBM Spectrum Archive EE system remains online, and other tape pools remain available for operations.

**Note:** Pool relocation will fail if any files are migrated to pools from IBM Spectrum Archive EE v1.1. To determine if any pools have files from v1.1, run the following command:

```
ltfsee_count_v11_files -p <pool_name> -l <library_name>
```

Pool relocation only works in a multi-library IBM Spectrum Archive EE cluster. Therefore, if you have a single-library environment and want to move a pool, you first need to have a second IBM Spectrum Archive EE node added to the cluster and an identical media type secondary tape library attached.

For the pool replacement procedure the following assumptions are made:

- ▶ IBM Spectrum Archive EE is configured as multi-library configuration.
  - ▶ A pool is selected to move from one library (called the *source library*) to another library (called the *destination library*).
  - ▶ All physical tape cartridges in the selected pool are moved to the destination library manually.
1. Get the pool information that is being relocated by running the `ltfsee info pools` command.
  2. Stop IBM Spectrum Archive EE.

3. Prepare the selected pool for relocation with the `ltfsee_config -m PREPARE_MOVE_POOL` command. Example 4-17 shows the output from running the command on pool1 from library lib1 to library lib2.

*Example 4-17 ltfsee\_config -m PREPARE\_MOVE\_POOL*

---

```
[root@server1 ~]# ltfsee_config -m PREPARE_MOVE_POOL -p pool1 -s lib1 -d lib2
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
PREPARE_MOVE_POOL -p pool1 -s lib1 -d lib2
PREPARE_MOVE_POOL mode starts .
```

```
## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file
ltfsee_config.filesystem exists.
```

Start the `mmapplypolicy GPFS` command to get a list of migrated files that are in IBM Spectrum Archive V1.1 format.

Starting IBM Spectrum Archive EE to refresh the index cache of the libraries.

```
Library name: lib2, library serial: 0000013FA052041F, control node (ltfsee_md)
IP address: 9.11.244.23.
```

```
Running start command - sending request : lib2.
```

```
Library name: lib1, library serial: 0000013FA052041B, control node (ltfsee_md)
IP address: 9.11.244.22.
```

```
Running start command - sending request : lib1.
```

```
Running start command - waiting for completion : lib2.
```

```
.....
```

```
Started the IBM Spectrum Archive EE services for library lib2 with good status.
```

```
Running start command - waiting for completion : lib1.
```

```
Started the IBM Spectrum Archive EE services for library lib1 with good status.
```

```
Stopping IBM Spectrum Archive EE.
```

```
Library name: lib2, library serial: 0000013FA052041F, control node (ltfsee_md)
IP address: 9.11.244.23.
```

```
Running stop command - sending request and waiting for the completion.
```

```
Library name: lib1, library serial: 0000013FA052041B, control node (ltfsee_md)
IP address: 9.11.244.22.
```

```
Running stop command - sending request and waiting for the completion...
```

```
Stopped the IBM Spectrum Archive EE services for library lib2.
```

```
Stopped the IBM Spectrum Archive EE services for library lib1.
```

```
Checking tapes with a duplicated VOLSER in library lib1 and lib2.
```

```
Copying pool definitions from library lib1 to lib2.
```

```
Saving index cache of the tapes in pool pool1.
```

```
PREPARE_MOVE_POOL mode completed.
```

---

**Note:** Pools with the same name appear in the source and destination libraries after this command. (The pool name specified for relocation must not exist in the destination library before this command). The mode attribute of the source pool is set to **relocation\_source**, and the mode attribute of the destination pool is set to **relocation\_destination**.

- List pools that have been or are in process of being moved, as shown in Example 4-18.

*Example 4-18 List relocated pools*

```
[root@server1 ~]# ltfsee_config -m LIST_MOVE_POOLS
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
LIST_MOVE_POOLS
LIST_MOVE_POOLS mode starts .
```

```
## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file
ltfsee_config.filesystem exists.
```

Source pool	Destination pool	Activated
pool1@lib1	pool1@lib2	no

- Start IBM Spectrum Archive EE by using the **ltfsee start** command.

**Note:** Jobs related to the tapes in the pool with “source” or “destination” in the Mode attribute will be rejected.

- Make a list of tapes which belong to the pool for relocation with the **ltfsee info tapes** command.
- Move tapes within the selected pool to the IE slot to remove from the previous library using the **ltfsee tape move ieslot -p <pool> -l <library> -t <tape1, [tape2, tape3,...]>** command.
- Stop IBM Spectrum Archive.
- Activate the pool for relocation by running the **ltfsee\_config -m ACTIVATE\_MOVE\_POOL** command. Example 4-19 illustrates activating the relocated pool after all tapes have been physically moved to the new tape library.

*Example 4-19 The ltfsee\_config -m ACTIVATE\_MOVE\_POOL command*

```
[root@server1 ~]# ltfsee_config -m ACTIVATE_MOVE_POOL -p pool1 -s lib1 -d lib2
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
ACTIVATE_MOVE_POOL -p pool1 -s lib1 -d lib2
ACTIVATE_MOVE_POOL mode starts .
```

```
## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file
ltfsee_config.filesystem exists.
```

Starting IBM Spectrum Archive EE to refresh the index cache of the libraries.

```
Library name: lib2, library serial: 0000013FA052041F, control node (ltfsee_md)
IP address: 9.11.244.23.
Running start command - sending request : lib2.
```

```

Library name: lib1, library serial: 0000013FA052041B, control node (ltfsee_md)
IP address: 9.11.244.22.
Running start command - sending request : lib1.
Running start command - waiting for completion : lib2.
.....
Started the IBM Spectrum Archive EE services for library lib2 with good status.
Running start command - waiting for completion : lib1.
.....
Started the IBM Spectrum Archive EE services for library lib1 with good status.

Stopping IBM Spectrum Archive EE.

Library name: lib2, library serial: 0000013FA052041F, control node (ltfsee_md)
IP address: 9.11.244.23.
Running stop command - sending request and waiting for the completion.
Library name: lib1, library serial: 0000013FA052041B, control node (ltfsee_md)
IP address: 9.11.244.22.
Running stop command - sending request and waiting for the completion.
Stopped the IBM Spectrum Archive EE services for library lib2.
.
Stopped the IBM Spectrum Archive EE services for library lib1.

Checking whether all the tapes in pool pool1 are properly moved from library
lib1 to lib2.

Updating the index cache of library lib2.

ACTIVATE_MOVE_POOL mode completed .

```

10. Verify that the pool has been activated, as shown in Example 4-20.

*Example 4-20 List relocated pools*

```

[root@server1 ~]# ltfsee_config -m LIST_MOVE_POOLS
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
LIST_MOVE_POOLS
LIST_MOVE_POOLS mode starts .

## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file
ltfsee_config.filesystem exists.

```

Source pool	Destination pool	Activated
pool1@lib1	pool1@lib2	yes

11. Start IBM Spectrum Archive EE by using the **ltfsee start** command.

12. Verify that the pool and tapes are valid, as shown in Example 4-21.

*Example 4-21 Verify pool and tapes in new tape library*

```

[root@server1 ~]# ltfsee info pools

Pool Name Total(TiB) Used(TiB) Free(TiB) Reclaimable(TiB) Tapes Type
Library Node Group

```

```

pool2          7.4      0.0      7.4          0.0      2  LTO
lib2          G0
pool1          9.6      0.0      9.6          0.0      3  LTO
lib2          G0

```

```
[root@server1 ~]# ltfsee info tapes
```

Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)
Pool	Library	Address	Drive	Appendable		
DV1999L7	Valid	L7	5338	5	5333	0
pool1	lib2	1059	-	yes		
DV0062L6	Valid	L6	2242	0	2242	0
pool1	lib2	1058	-	yes		
DV1995L7	Valid	L7	5338	0	5338	0
pool2	lib2	1057	-	yes		
DV0060L6	Valid	L6	2242	5	2237	0
pool2	lib2	1056	-	yes		
IM0732L6	Valid	L6	2242	0	2242	0
pool1	lib2	1055	-	yes		



## Upgrading from version 1.1.x

This chapter provides information about upgrading from IBM Spectrum Archive version 1.1.x to IBM Spectrum Archive version 1.2.

This chapter includes the following topics:

- ▶ Overview of the IBM Spectrum Archive EE upgrade from version 1.1.x
- ▶ The `ltfsee_config_save` command
- ▶ The `ltfsee_config_upgrade` command
- ▶ Upgrading from IBM Spectrum Archive Enterprise Edition version 1.1.x
- ▶ Upgrade processing output examples

## 5.1 Overview of the IBM Spectrum Archive EE upgrade from version 1.1.x

This section provides an overview of the process to upgrade IBM Spectrum Archive Enterprise Edition version 1.1.x to version 1.2.x or a subsequent version.

Before you begin the upgrade process, ensure that your system conforms to the IBM Spectrum Archive EE version 1.2.x system requirements. For more information, see System Requirements in the IBM Spectrum Archive Enterprise Edition online documentation:

[https://www.ibm.com/support/knowledgecenter/ST9MBR\\_1.2.0/1tfs\\_ee\\_system\\_reqs.html](https://www.ibm.com/support/knowledgecenter/ST9MBR_1.2.0/1tfs_ee_system_reqs.html)

When you upgrade from IBM Spectrum Archive EE version 1.1.x to version 1.2.x or a subsequent version, you use two command-line utilities:

- ▶ The **1tfsee\_config\_save** command, which archives the current version 1.1.x configuration.
- ▶ The **1tfsee\_config\_upgrade** command, which reads the information in the archived configuration file, and configures the new version with the same configuration settings.

Using the two commands ensures that the configuration settings are identical, even if the operating system of the server is also upgraded. However, you must ensure that the hardware configuration remains the same as it was for version 1.1.x, including these details:

- ▶ The GPFS mount point
- ▶ The number of nodes
- ▶ The names of the nodes
- ▶ The drives that are connected to each node must all remain the same

The **1tfsee\_config\_save** command saves the configuration information to a file named `1tfsee_saved_config_<YYYYMMDDHHmmSS>.tar.gz`. The command creates the `./1tfsee/upgrade` directory in your `<metadata_filesystem_name>` and stores the file there unless a different location is explicitly provided. The `<metadata_filesystem_name>` is the name of the GPFS file system that you selected to store the IBM Spectrum Archive EE metadata when you installed the IBM Spectrum Archive EE product.

The **1tfsee\_config\_upgrade** command sets up the required configuration records (files and databases) in the style that is used by version 1.2.x, while it keeps the configuration the same as used in version 1.1.x. The tape library that is used in version 1.1.x is assigned the name `lib0`. A nodegroup is created and named `G0`. All of the tape pools and nodes from version 1.1.x are assigned to nodegroup `G0` in the configuration database.

Custom parameters in the `1tfsee.config` file are not automatically inherited during the upgrade process, except for the pool and drive role attributes. If manual changes were made to the version 1.1.x configuration file, then after the upgrade process completes, manually copy those custom parameters from the version 1.1.x configuration file at `<metadata_filesystem>/1tfsee/config/1tfsee.config.V11` to the new configuration file. This is a rare occurrence.

The tape metadata that is stored on disk is converted during the upgrade operation. For planning purposes, estimate 10 - 20 minutes of processing time per million files for files that are already migrated to tapes. While the function is running, the progress is shown as the number of total and remaining tape cartridges.

Any empty tape pools in your IBM Spectrum Archive EE system are not preserved during the upgrade process.

If the `ltfsee_config_save` operation or the `ltfsee_config_upgrade` operation detects an error, the operation fails. An error message is displayed, and is also written to a log file. The log file is created in a subdirectory called `/var/opt/ibm/ltfsee/local/config_history`. For detailed instructions on how to handle errors that are detected during the upgrade process, see the 5.4.3, “Upgrading IBM Spectrum Archive Enterprise Edition” on page 84.

## 5.2 The `ltfsee_config_save` command

This section describes how to use the `ltfsee_config_save` command when you upgrade IBM Spectrum Archive Enterprise Edition from version 1.1.x to version 1.2.x, or a subsequent version.

Before you use the `ltfsee_config_save` command when you upgrade your IBM Spectrum Archive EE system from version 1.1.x to version 1.2.x, or a subsequent version, read 5.1, “Overview of the IBM Spectrum Archive EE upgrade from version 1.1.x” on page 80. Then, follow the step-by-step upgrade procedures that are documented in 5.4, “Upgrading from IBM Spectrum Archive Enterprise Edition version 1.1.x” on page 82.

This command is not used for everyday operations. This command is intended only for upgrading from IBM Spectrum Archive EE version 1.1.x.

You must have root user authority to use this command:

- ▶ **`ltfsee_config_save [--all]`**  
Collects the configuration files from all of the nodes on which IBM Spectrum Archive EE and LE+ are installed, and creates a saved configuration package file. The process places the file in the GPFS `metadata_fs/.ltfsee/upgrade` directory. This is the default execution.
- ▶ **`ltfsee_config_save [--all] <path/filename>`**  
Creates the saved configuration package file, and places it in the specified file name and path.
- ▶ **`ltfsee_config_save --single <path/filename>`**  
Collects configuration files from only the node on which you are running the command. The file name must be provided. This option can be used when the network between nodes is down. You must merge the collected configuration files into one tree structure so that the upgrade function can find the proper files.
- ▶ **`ltfsee_config_save --redo <path/filename>`**  
Sets the mid-of-configuration flag file after the command checks that the `<path/filename>` provided file is a saved configuration package file. Normally, you do not use this command.

When the `ltfsee_config_save` command is successful, it returns a zero. When an error is found, the command returns a nonzero value.

During the `ltfsee_config_save` process, a file that is named `mid_of_ltfsee_config` is created. The file is located in the `.ltfsee/upgrade` directory. The file indicates to the IBM Spectrum Archive EE system that upgrade operations are ongoing. Any attempt to use the `ltfsee start` command fails until the `mid_of_ltfsee_config` file is removed, at the end of successful `ltfsee_config_upgrade` processing.

## 5.3 The `ltfsee_config_upgrade` command

The `ltfsee_config_upgrade` command uses the saved configuration information from IBM Spectrum Archive Enterprise Edition version 1.1.x to create the same configuration in version 1.2.x, but using the 1.2.x style of configuration files.

Before you use the `ltfsee_config_upgrade` command when you upgrade your IBM Spectrum Archive EE system from version 1.1.x to version 1.2.x, or a subsequent version, read 5.1, “Overview of the IBM Spectrum Archive EE upgrade from version 1.1.x” on page 80. Then, follow the step-by-step upgrade procedures in 5.4, “Upgrading from IBM Spectrum Archive Enterprise Edition version 1.1.x” on page 82.

You must have root user authority to use this command.

You must run the `ltfsee_config_save` command before you run the `ltfsee_config_upgrade` command:

```
ltfsee_config_upgrade <path/filename>
```

This command creates version 1.2.x style of configuration files from the saved version 1.1.x configuration package file. The `<path/filename>` must be provided to access the saved\_configuration package file.

The `ltfsee_config_upgrade` command returns a zero when it completes successfully. When an error is found, the command returns nonzero values.

At the end of the successful completion of the `ltfsee_config_upgrade` command, the `mid_of_ltfsee_config` file is removed. If the file is not removed, the upgrade was not successful. The `ltfsee_start` command fails to restart your system until your upgrade is successful.

## 5.4 Upgrading from IBM Spectrum Archive Enterprise Edition version 1.1.x

This section covers how to upgrade your IBM Spectrum Archive Enterprise Edition from version 1.1.x to version 1.2.x, or a subsequent version.

The upgrade procedure has four parts, which must be completed in the following sequence:

1. Prerequisite tasks for the upgrade
2. Upgrading your operating system
3. Upgrading IBM Spectrum Archive Enterprise Edition
4. Post-upgrade tasks

### 5.4.1 Prerequisite tasks for the upgrade

Complete these tasks before you change IBM Spectrum Archive Enterprise Edition or your operating system:

1. Before you uninstall IBM Spectrum Archive EE version 1.1.x, identify any tapes that are used for recall tasks that do not belong to a tape pool by running the `ltfsee info tapes` command. If such tapes exist, create a file that lists the tape names and the tape pools to which you want them assigned. Name the file `upgrade_pools.lst`, and save it.

The file is used in a subsequent step. The list contains the <tape\_id> followed by a comma, and the <pool\_name>, as in the following example:

```
YMT001L7,poolA
YMT002L7,poolA
YMT003L7,poolB
YMT004L7,poolB
```

Migrated files on tapes that do not belong to a tape pool, and that do not get included on the `upgrade_pools.lst`, cannot be recalled until they are added to a tape pool. If such tapes are found after the upgrade to IBM Spectrum Archive EE version 1.2.x, they can be added to a tape pool by using the `ltfsee pool add` command. For an example of how to create this file, see the creating and moving the `upgrade_offline.lst` section in Example 5-5 on page 88.

2. Before you uninstall IBM Spectrum Archive EE version 1.1.x, identify any offline exported tapes that do not belong to a tape pool by running the `ltfsee info tapes` command. If such tapes exist, create a file to list the tape names and the tape pool to which you want them assigned. Name the file `upgrade_offline.lst`, and save it. The file is used in a subsequent step. The list contains the <tape\_id> followed by a comma, and the <pool\_name>, as in the following example:

```
YMT005L7,poolC
YMT006L7,poolC
YMT007L7,poolD
YMT008L7,poolD
```

Tapes that are in the state `offline exported`, do not belong to a tape pool, and do not get included on the `upgrade_offline.lst` are detected by the upgrade operation. For all IBM Spectrum Archive EE version 1.1.x releases, except for the first release, the upgrade function adds the <tape\_id> to the `upgrade_offline.lst` file. Then, upgrade processing ends. You must then edit the `upgrade_offline.lst` file to add a tape pool assignment for each tape, then restart the upgrade process.

For users who are upgrading from the first release of IBM Spectrum Archive EE, offline exported tapes that are not manually added to the `upgrade_offline.lst` do not get automatically detected. If they are not added to the `upgrade_offline.lst`, the tapes become inaccessible.

In the future, if you find that you need an offline exported tape that you neglected to put on the `upgrade_offline.lst`, you need to contact your IBM service support representative for assistance. For an example on how to create this file, see 5.5.2, “Example of creating and moving the `upgrade_offline.lst` in the home directory to <metadata\_filesystem\_name>/ltfsee/upgrade” on page 88.

3. Issue the `ltfsee stop` command to stop IBM Spectrum Archive EE from running on all nodes.

4. Unmount the tape library from each node by using the following command:

```
$ umount /ltfs
```

Use the `ps` command to determine whether the unmount process is complete:

```
$ ps aux | grep ltfs
```

5. Extract the most current version of the IBM Spectrum Archive EE installation package, but do not install it. For more information, see Extracting binary rpm files from an installation package in the online IBM Spectrum Archive Enterprise Edition documentation:

[https://www.ibm.com/support/knowledgecenter/ST9MBR\\_1.2.0/ltfs\\_ee\\_installing\\_linux\\_enterprise\\_package.html](https://www.ibm.com/support/knowledgecenter/ST9MBR_1.2.0/ltfs_ee_installing_linux_enterprise_package.html)

The installation package includes the `l1tfsee_config_save` and the `l1tfsee_config_upgrade` command-line utilities. The `l1tfsee_config_save` utility is placed in the same location as the `l1tfsee_install` tool, which is in the `rpm` directory that is created by the IBM Spectrum Archive EE installer. The `l1tfsee_config_upgrade` utility is placed in the `/opt/ibm/l1tfsee/bin` directory.

6. Run the `l1tfsee_config_save` command on one node:

```
./l1tfsee_config_save
```

The command creates a saved configuration file, and displays the name and location of the file in the window.

If an error is detected during `l1tfsee_config_save` processing, an error message is displayed, and also written to the log file at `/var/opt/ibm/l1tfsee/local/config_history`. Resolve the problem and rerun the `l1tfsee_config_save` command.

7. If you have an `upgrade_pools.lst` file, or an `upgrade_offline.lst` file, place them in the `<metadata_filesystem_name>/l1tfsee/upgrade` directory.

### What to do next

If you have any changes that you want to make to your operating system, system applications, or to IBM Spectrum Scale, go to 5.4.2, “Upgrading your operating system” on page 84. If you do not need any operating system changes, go to 5.4.3, “Upgrading IBM Spectrum Archive Enterprise Edition” on page 84.

## 5.4.2 Upgrading your operating system

This section describes upgrading your operating system before you upgrade IBM Spectrum Archive Enterprise Edition.

### Upgrade procedure

Make all of your planned changes on every server. Make backup copies of any files that might be deleted or altered. Changes might include operating system upgrades, IBM Spectrum Scale upgrades, or other application upgrades.

**Note:** When you upgrade IBM Spectrum Archive EE, ensure that the files in the GPFS file system (a component of IBM Spectrum Scale) remain in the same location. Therefore, if you are upgrading IBM Spectrum Scale, ensure that the upgrade method does not cause location changes to the files in the GPFS file system.

## 5.4.3 Upgrading IBM Spectrum Archive Enterprise Edition

This section describes how to upgrade from IBM Spectrum Archive EE version 1.1.x.

### Before you begin

Ensure that you complete all prerequisite tasks, as listed in 5.4.1, “Prerequisite tasks for the upgrade” on page 82, and any other operating system, application, or IBM Spectrum Scale upgrades that you planned to make.

## Upgrade procedure

To upgrade, complete these steps:

1. After all of your server changes are complete, install the 1.2.x version (or a subsequent version) of IBM Spectrum Archive EE by using the **l1tfsee\_install --upgrade** command. For more information, see *Installing, upgrading, or uninstalling IBM Spectrum Archive Enterprise Edition* in the online IBM Spectrum Archive Enterprise Edition documentation.
2. Run the **l1tfsee\_config\_upgrade <path/filename>** command on the node that you want to be the control node. The `<path/filename>` specifies the location and file name of the configuration information that you saved when you ran the **l1tfsee\_config\_save** command. The file name is `l1tfsee_saved_config_<YYYYMMDDHHmmSS>.tar.gz`, which is located in the `<metadata_filesystem_name>/l1tfsee/upgrade` directory (unless you specified a different location).

The **l1tfsee\_config\_upgrade** utility is located in the `/opt/ibm/l1tfsee/bin` directory.

The upgrade process creates the configuration records (files and databases) in the style that is used by IBM Spectrum Archive EE version 1.2.x (or a subsequent version), while it keeps the same configuration that was used in version 1.1.x.

For more information about control nodes, see Introduction to IBM Spectrum Archive Enterprise Edition in the online IBM Spectrum Archive Enterprise Edition documentation:

[https://www.ibm.com/support/knowledgecenter/ST9MBR\\_1.2.0/l1tfs\\_ee\\_intro.html](https://www.ibm.com/support/knowledgecenter/ST9MBR_1.2.0/l1tfs_ee_intro.html)

3. During the `l1tfsee_config_upgrade` processing, a “detected configuration” step prints the detected configuration for your review. If the detected configuration is as you intended, respond to the system with a Y to continue. If the configuration is not as you intended, resolve the error, and see step 5 for error recovery actions.
4. The `l1tfsee_config_upgrade` process produces an upgraded configuration when it successfully completes. A message is displayed in the window, and written to a log, stating that the `dcache` directory is no longer needed, and can be deleted. The log file is created in a directory at `/var/opt/ibm/l1tfsee/local/config_history`.
5. If any problems are detected during the upgrade process, an error message is displayed in the window, and written to the log file at `/var/opt/ibm/l1tfsee/local/config_history`. Choose from the following error recovery actions, according to where you were in the upgrade process when the error occurred:
  - When an error is detected after you run the **l1tfsee\_config\_save** command in step 6 on page 84, you need to resolve the problem and rerun the **l1tfsee\_config\_save** command.
  - When an error is detected after you run the **l1tfsee\_config\_upgrade** command, and after you respond to the “detected configuration” query in step 3, resolve the problem and rerun the **l1tfsee\_config\_upgrade** command.
  - When an error is detected after you run the **l1tfsee\_config\_save** command, but before you answer the query about the detected configuration during the **l1tfsee\_config\_upgrade** command in step 3, complete the following actions:
    - i. Correct the source of the error.
    - ii. Run the **l1tfsee\_config\_save** command to create another saved-configuration package file.
    - iii. Compare the contents of the previous saved-configuration package file and the new one. If differences exist between them, select the correct file.
    - iv. Update the saved-configuration package file with the correct file.
    - v. Rerun the **l1tfsee\_config\_upgrade** command.

6. If for any reason you want to return to the previous configuration, first reinstall version 1.1.x. Then, manually return the saved version 1.1.x configuration to its original location if it was previously moved to a different location. Otherwise, this file should be in the <gpfs file system> in a folder called .Spaceman.

## 5.4.4 Post-upgrade tasks

This section describes the tasks to complete after the successful upgrade to IBM Spectrum Archive EE version 1.2.x, or a subsequent version.

Complete these post-upgrade procedures:

1. If you want to make configuration changes after the upgrade completes, such as adding another tape library, or creating nodes or tape pools, use the applicable `ltfsee_config` or `ltfsee` commands. To verify that any additional configuration changes are as expected, display the current configuration by using the `ltfsee_config -m INFO` command. See the `ltfsee_config -m INFO` example in Example 5-8 on page 93.
2. Start IBM Spectrum Archive EE when your configuration changes are complete. To start IBM Spectrum Archive EE, mount the library on all nodes, and run the `ltfsee start` command on the control node.
3. Review the version 1.2.x `ltfsee` commands to become familiar with the differences in the command-line options and output that exist between IBM Spectrum Archive Enterprise Edition versions 1.1.x and 1.2.x. Most of the `ltfsee` commands have changes to their parameters to support more libraries, the new node groups, and new functions.

The compatibility of version 1.1 input parameters is maintained as much as possible. However, the version 1.2 documentation lists only version 1.2 command syntax, as the older version 1.1 syntax will soon become obsolete.

4. Become familiar with changes that affect the policy files. For example, adding `OR PATH_NAME LIKE '%/.ltfsee/%` in the exclude condition. Also, changing the `OPTS` parameter to have `-p` at the beginning of the command, followed by the tape pool name and the tape library name. For example, `OPTS '-p Tapepool12@1library1`.

## 5.5 Upgrade processing output examples

This section shows examples of output when you upgrade from IBM Spectrum Archive 1.1.x, using the `ltfsee_config_save` and `ltfsee_config_upgrade` commands.

### 5.5.1 Example output from `ltfsee_config_save`

Example 5-1 shows the output from `ltfsee_config_save`.

*Example 5-1 Output from `ltfsee_config_save`*

---

```
[root@ltfssn2 rpm.1210_10229]# ./ltfsee_config_save
The EE config save script is starting: ./ltfsee_config_save
Creating a temporary directory '/var/opt/ibm/ltfsee/local/saved_config'.
Collection V1.1 configuration files.
Copying local file '/ibm/gpfs/.SpaceMan/ltfsee.config' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/common'.
Copying local directory '/ibm/gpfs/0000013FA0020408/cart_repos' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/common/0000013FA0020408'.
Collecting information from nodes in the cluster.
```

```
node: ltfssn2.tuc.stglabs.ibm.com
Copying file '/etc/ltfs.conf.local' from node 'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
File '/etc/syslog.conf' does not exist in node 'ltfssn2.tuc.stglabs.ibm.com'.
Copying file '/etc/rsyslog.conf' from node 'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
Copying file '/etc/logrotate.d/syslog' from node 'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
File '/etc/syslog-ng/syslog-ng.conf' does not exist in node
'ltfssn2.tuc.stglabs.ibm.com'.
Copying file '/etc/rc.gpfshsm' from node 'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
Copying file '/opt/tivoli/tsm/client/hsm/bin/rc.gpfshsm' from node
'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
Copying file '/opt/tivoli/tsm/client/ba/bin/dsm.opt' from node
'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
Copying file '/opt/tivoli/tsm/client/ba/bin/dsm.sys' from node
'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
Copying file '/var/opt/ibm/ltfsee/local/ltfsee_config.filesystem' from node
'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
Copying file '/var/opt/ibm/ltfsee/local/ltfsee_msg_priority.conf' from node
'ltfssn2.tuc.stglabs.ibm.com' to
'/var/opt/ibm/ltfsee/local/saved_config/V11/ltfssn2.tuc.stglabs.ibm.com'.
File '/var/opt/ibm/ltfsee/local/ltfseetrace.conf' does not exist in node
'ltfssn2.tuc.stglabs.ibm.com'.
Archiving collected files.
Archive file
/ibm/gpfs/.ltfsee/upgrade/ltfsee_saved_config_V11_20160629115503.tar.gz created.
Removing collected files.
```

All configuration files are collected and archived into the following file.  
>>> /ibm/gpfs/.ltfsee/upgrade/ltfsee\_saved\_config\_V11\_20160629115503.tar.gz <<<  
Keep this filename to be used with the ltfsee\_config\_upgrade command later.

**NOTE:**

If there are tapes that are used for recall, but that do not belong to any pool, the tapes

must be assigned to a pool in the upgraded configuration.

If you have such tapes, prepare file

/ibm/gpfs/.ltfsee/upgrade/upgrade\_pools.lst

to provide the pool name where each tape should be assigned during the upgrade process.

If there are offline exported tapes, you must prepare file

/ibm/gpfs/.ltfsee/upgrade/upgrade\_offline.lst.

When the ltfsee\_config\_upgrade command detects that there are offline exported tapes but the file

is not there, the command fails and creates a template file.

In that case, fill out the template file and rerun the ltfsee\_config\_upgrade command.

## 5.5.2 Example of creating and moving the upgrade\_offline.lst in the home directory to <metadata\_filesystem\_name>/.ltfsee/upgrade

The following examples display the use of ltfsee info tapes and grepping for “exported offline” tapes for later insertion into the upgrade\_offline.lst file. The same principles can be applied to create the upgrade\_pools.lst file.

Example 5-2 displays the tapes that are exported offline by running the **ltfsee info tapes** command.

*Example 5-2 ltfsee info tapes | grep “(exported offline)”*

---

```
JD0334JD      TS1150      9022GB      9022GB      0GB      1035      -  
-              -              Valid LTFS (exported  
offline)
```

---

Example 5-3 creates the upgrade\_offline.lst file by using the information provided above. For example purposes, only the tape JD0334JD was originally in pool primary and thus going to be readded to pool primary as an exported offline tape. Modify the grep expression and the pool to correspond to your scenario.

*Example 5-3 Creating the upgrade\_offline.lst*

---

```
[root@ltfssn2 ~]# ltfsee info tapes | grep “(JD0334JD|[tape2]|[tape3]|etc..)” | awk  
'{print $1",primary"}' >>  
upgrade_offline.lst
```

---

Example 5-4 shows the content of the upgrade\_offline.lst file.

*Example 5-4 The content of the upgrade\_offline.lst file*

---

```
[root@ltfssn2 ~]# cat upgrade_offline.lst  
JD0334JD,primary
```

---

Example 5-5 shows moving the upgrade\_offline.lst file in the home directory to <metadata\_filesystem\_name>/.ltfsee/upgrade.

*Example 5-5 Moving the upgrade\_offline.lst to <metadata\_filesystem\_name>/.ltfsee/upgrade*

---

```
[root@ltfssn2 upgrade]# mv ~/upgrade_offline.lst .  
[root@ltfssn2 upgrade]# ls -ltr  
total 8  
-rw-r--r-- 1 root root 17 May 17 15:05 upgrade_offline.lst  
-rw-r--r-- 1 root root 4689 May 17 15:07 ltfsee_saved_config_V11_20160517150754.tar.gz  
-rw-r--r-- 1 root root 0 May 17 15:07 mid_of_ltfsee_config
```

---

## 5.5.3 Example output from ltfsee\_install --upgrade

Example 5-6 shows the output from **ltfsee\_install --upgrade** processing.

*Example 5-6 Output from ltfsee\_install --upgrade*

---

```
[root@ltfssn2 rpm.1210_10229]# ./ltfsee_install --upgrade
```

Upgrade installation option is selected.

This option uninstalls the following preinstalled packages before installation:

- LE+ component
- HSM component
- LTFS Migration Driver

These packages will be upgraded to the version that IBM Spectrum Archive Enterprise Edition provides.

Do you want to continue? (y/n) y  
Checking rpm installation and version.

The prerequisites checking completed successfully.

Stopping the space management daemon.

```

-----
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 1.3
  Client date/time: 06/29/2016 12:09:59
(c) Copyright by IBM Corporation and other(s) 1990, 2014. All Rights Reserved.

```

```

-----
Terminated the space management daemon.
The package ltfs-mig has been installed.
The package ltfs-mig is uninstalled successfully.
The package TIVsm-HSM has been installed.
Removing HSM from GPFS startup script
The package TIVsm-HSM is uninstalled successfully.
The package TIVsm-BA has been installed.
The package TIVsm-BA is uninstalled successfully.
The package TIVsm-API64 has been installed.
The package TIVsm-API64 is uninstalled successfully.
The package gskssl64 has been installed.
The package gskssl64 is uninstalled successfully.
The package gskcrypt64 has been installed.
The package gskcrypt64 is uninstalled successfully.
The package ltfs-library has been installed.
Waiting for LTFS processes to finish: [ OK ]
The package ltfs-library is uninstalled successfully.
The package ltfs has been installed.
Waiting for LTFS processes to finish: [ OK ]
warning: /etc/ltfs.conf.local saved as /etc/ltfs.conf.local.rpmsave
The package ltfs is uninstalled successfully.

```

```

Preparing... ##### [100%]
  1:ltfsle ##### [100%]
Preparing... ##### [100%]
  1:ltfsle-library ##### [100%]
Preparing... ##### [100%]
  1:ltfsle-library-plus ##### [100%]
Preparing... ##### [100%]
  1:gskcrypt64 ##### [100%]
Preparing... ##### [100%]
  1:gskssl64 ##### [100%]
Preparing... ##### [100%]
  1:TIVsm-API64 ##### [100%]

```

```

Preparing... ##### [100%]
  1:TIVsm-BA ##### [100%]
Preparing... ##### [100%]
  1:TIVsm-HSM ##### [100%]
Stopping the space management daemon.
-----
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 6.1
  Client date/time: 06/29/2016 12:12:15
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

-----
Terminated the space management daemon.
Preparing... ##### [100%]
  1:ltfs-mig ##### [100%]
Starting the space management daemon.
-----
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 6.1
  Client date/time: 06/29/2016 12:12:29
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

-----
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 6.1
  Client date/time: 06/29/2016 12:12:29
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

TSMThe node failover option is already enabled in mode ENABLED.
All rpm packages are installed successfully.
Complete the configuration using /opt/ibm/ltfsee/bin/ltfsee_config
command.

```

---

## 5.5.4 Example output from ltfsee\_config\_upgrade

Example 5-7 shows the output from ltfsee\_config\_upgrade processing.

*Example 5-7 Output from ltfsee\_config\_upgrade*

---

```

[root@ltfssn2 upgrade]# ltfsee_config_upgrade
/ibm/gpfs/.ltfsee/upgrade/ltfsee_saved_config_V11_20160629115503.tar.gz
Starting the IBM Spectrum Archive EE configuration upgrade tool
(['opt/ibm/ltfsee/bin/ltfsee_config_upgrade',
'/ibm/gpfs/.ltfsee/upgrade/ltfsee_saved_config_V11_20160629115503.tar.gz']).
Detecting the configuration from the input file.
Detected the following configuration.
=====
Metadata Filesystem:
  /ibm/gpfs
HSM Filesystems:
  /ibm/gpfs

```

```

Library: Name=lib0, S/N=0000013FA0020408
Node Group: Name=GO
Node: ltfssn2.tuc.stglabs.ibm.com
Drive: S/N=000003000576, Attribute='mrg'
Pool: primary
Tape: Barcode=JD0339JD
Tape: Barcode=JD3592JD
Tape: Barcode=JD3595JD
Tape: Barcode=JD3604JD
Pool: copy
Tape: Barcode=JD0334JD, Offline Message='offline for upgrade test'
Tape: Barcode=JD0337JD

```

```
=====
```

```

Upgrade the configuration? (y/n) > y
Validating the configuration.
Set ltfssn2.tuc.stglabs.ibm.com as the control node.
The version of EE on localhost is ltf-mig-1.2.1.0-10229.
The version of EE on ltfssn2.tuc.stglabs.ibm.com is ltf-mig-1.2.1.0-10229.
Library 0000013FA0020408 is found on ltfssn2.tuc.stglabs.ibm.com.
The drive 000003000576 is found on ltfssn2.tuc.stglabs.ibm.com.
Applying the configuration.
Configuring the cluster.
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
CLUSTER /tmp/tmpLSdRSc
CLUSTER mode starts .

```

```

## 1. Check to see if the cluster is already created ##
Cluster is already created. The configuration file ltfsee_config.filesystem will
be overridden.

```

```

## 2. Check prerequisite on cluster ##
Cluster name: ltfssn2.tuc.stglabs.ibm.com
ID: 7186459686228696309
Successfully validated the prerequisites.

```

```

## 3. Show the file systems to be configured by autopilot ##
File systems for LE+ component configuration and internal data
CONFIGANDMETA /dev/gpfs ltfssn2_nsd1 /ibm/gpfs
File systems for Space Management
SPACEMAN /dev/gpfs ltfssn2_nsd1 /ibm/gpfs

```

```

## 4. Configure space management ##
Disabling unnecessary daemons...
Editing Space Management Client settings...
Restarting Space Management service...
Terminating dsmwatchd.....
Terminating dsmwatchd.....
Starting dsmmigfs.....
Configured space management.

```

```

## 5. Add selected file systems to Space Management ##
Added the selected file systems to the space management.

```

```

## 6. Store the file systems configuration and dispatch it to all nodes ##
Storing the file systems configuration.

```

```

Copying ltfsee_config.filesystem file.
Stored the cluster configuration and dispatched the configuration file.

## 7. Create metadata directories and the configuration parameters file. ##
Created metadata directories and the configuration parameters file.
Disabling runtime AFM file state checking.

CLUSTER mode completed.
The cluster configuration completed.
Configuring the control node on ltfssn2.tuc.stglabs.ibm.com.
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
ADD_CTRL_NODE 0000013FA0020408 lib0 :000003000576@7
ADD_CTRL_NODE mode starts .

## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file ltfsee_config.filesystem
exists.

## 2. Check prerequisite on node ##
Successfully validated the prerequisites.

## 3. IBM Spectrum Scale (GPFS) Configuration for Performance Improvement ##
Setting worker1Threads=400
Setting dmapiWorkerThreads=64
Configured IBM Spectrum Scale (GPFS) preformance related settings.

## 4. Configure space management ##
Disabling unnecessary daemons...
Editing Space Management Client settings...
Restarting Space Management service...
Terminating dsmwatchd.....
Terminating dsmwatchd.....
Starting dsmmigfs.....
Configured space management.

## 5. Add this node to a tape library ##
Added this node (ltfssn2.tuc.stglabs.ibm.com, node id 1) to library lib0 as its
control node.

## 6. Add this node to a node group ##
Added this node (ltfssn2.tuc.stglabs.ibm.com, node id 1) to node group G0.

## 7. Add drives to this node ##
Selected drives: :000003000576@7.
Added the selected drives to this node (ltfssn2.tuc.stglabs.ibm.com, node id 1).

## 8. Configure the LE+ component ##
Creating mount point...
Mount point folder '/ltfs' exists.
Use this folder for the LE+ component mount point as LE+ component assumes this
folder.
Configured the LE+ component.

## 9. Enabling the system log ##
Restarting rsyslog...

```

System log (rsyslog) is enabled for IBM Spectrum Archive Enterprise Edition.

```
ADD_CTRL_NODE mode completed .
The control node configuration completed on ltfssn2.tuc.stglabs.ibm.com.
Configurating a pool copy.
The pool configuration for copy completed.
Configurating a pool primary.
The pool configuration for primary completed.
Updating the metadata.
Converting 7 schema files on /ibm/gpfs/.ltfsee/meta/0000013FA0020408/volume_cache
Starting /ibm/gpfs/.ltfsee/meta/0000013FA0020408/volume_cache/JD3592JD.schema
Starting /ibm/gpfs/.ltfsee/meta/0000013FA0020408/volume_cache/JD0334JD.schema
Starting /ibm/gpfs/.ltfsee/meta/0000013FA0020408/volume_cache/JD0337JD.schema
Finished JD3592JD.schema: Progress -> 1 / 7
Starting /ibm/gpfs/.ltfsee/meta/0000013FA0020408/volume_cache/JD0336JD.schema
Starting /ibm/gpfs/.ltfsee/meta/0000013FA0020408/volume_cache/JD3604JD.schema
Starting /ibm/gpfs/.ltfsee/meta/0000013FA0020408/volume_cache/JD0339JD.schema
Finished JD0337JD.schema: Progress -> 2 / 7
Finished JD0334JD.schema: Progress -> 3 / 7
Starting /ibm/gpfs/.ltfsee/meta/0000013FA0020408/volume_cache/JD3595JD.schema
Finished JD3604JD.schema: Progress -> 4 / 7
Finished JD3595JD.schema: Progress -> 5 / 7
Finished JD0339JD.schema: Progress -> 6 / 7
Finished JD0336JD.schema: Progress -> 7 / 7
The upgrade procedure completed. The old metadata directory can be removed by
using the command "rm -rf /ibm/gpfs/0000013FA0020408". It might take a long time.
```

---

**Note:** Before you remove the folder 0000013FA0020408 (this folder will be named the Serial Number (S/N) of your library) in /ibm/gpfs, make sure you can start EE and that everything is operating correctly. After you are comfortable with the operation and everything is running like it should, you can then delete the folder.

After the upgrade has finished, run `ltfsee_config -m INFO` to display the configuration of your environment, and verify that it matches the way you had set it up for the upgrade. Example 5-8 shows the output of the command.

*Example 5-8 ltfsee\_config -m INFO*

---

```
[root@ltfssn2 ~]# ltfsee_config -m INFO
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m INFO
INFO mode starts .

## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file ltfsee_config.filesystem
exists.
Metadata Filesystem:
  /ibm/gpfs
HSM Filesystems:
  /ibm/gpfs
Library: Name=lib0, S/N=0000013FA0020408
Node Group: Name=G0
Node: ltfssn2.tuc.stglabs.ibm.com
Drive: S/N=000003000576, Attribute='mrg'
Pool: Name=primary, ID=5c86992c-a98d-4dbd-bd70-ec06576300a8
Tape: Barcode=JD0339JD
```

Tape: Barcode=JD3592JD  
Tape: Barcode=JD3595JD  
Tape: Barcode=JD3604JD  
Pool: Name=copy, ID=73077cfd-0cdf-43cf-9866-56a664b12968  
Tape: Barcode=JD0334JD, Offline Message='offline for upgrade test'  
Tape: Barcode=JD0337JD

---



# Configuration

This chapter provides information about the postinstallation configuration of the IBM Spectrum Archive Enterprise Edition (IBM Spectrum Archive EE).

This chapter includes the following topics:

- ▶ Configuration prerequisites
- ▶ Configuring IBM Spectrum Archive EE
- ▶ First-time start of IBM Spectrum Archive EE
- ▶ Configuring IBM Spectrum Archive EE with IBM Spectrum Scale AFM

**Note:** In the lab setup for this book, we used a Red Hat-based Linux system. The screen captures within this chapter are based on Version 1 Release 2 of the product. Although the steps that you will perform are the same, you might see slightly different output responses depending on your currently used version and release of the product.

## 6.1 Configuration prerequisites

This section describes the tasks that must be completed before IBM Spectrum Archive EE is configured.

Ensure that the following prerequisites are met before IBM Spectrum Archive EE is configured. For more information, see 6.2, “Configuring IBM Spectrum Archive EE” on page 104.

- ▶ The Configuration worksheet is completed and available during the configuration process.

**Tip:** Table 6-1 on page 97, Table 6-2 on page 97, Table 6-3 on page 97, Table 6-4 on page 98, Table 6-5 on page 99, and Table 6-6 on page 99 provide a set of sample configuration worksheets. You can print and use these samples for your own purpose during your configuration of IBM Spectrum Archive EE.

- ▶ The key-based login with OpenSSH is configured.
- ▶ The IBM Spectrum Scale system is prepared and ready for use on your Linux server system.
- ▶ The control paths (CPs) to the tape library logical libraries are configured and enabled. You need at least one CP per node.

**Important:** You can configure the tape library for redundant CPs Control Path Failover feature (CPF). Other configuration steps must be performed with the IBM tape device driver for Linux (`lin_tape`) for this kind of configuration.

### 6.1.1 Configuration worksheet tables

Print Table 6-1 on page 97, Table 6-2 on page 97, Table 6-3 on page 97, Table 6-4 on page 98, Table 6-5 on page 99, and Table 6-6 on page 99 and use them as worksheets or as a template to create your own worksheets to record the information you need to configure IBM Spectrum Archive EE.

For more information, see 6.1.2, “Obtaining configuration information” on page 100 and follow the steps to obtain the information that is required to complete your worksheet.

The information in the following tables is required to configure IBM Spectrum Archive EE. Complete Table 6-4 on page 98, Table 6-5 on page 99, and Table 6-6 on page 99 with the required information and refer to this information as necessary during the configuration process, as described in 6.2, “Configuring IBM Spectrum Archive EE” on page 104.

Table 6-1, Table 6-2, and Table 6-3 show example configuration worksheets with the parameters completed for the lab setup that was used to write this book.

Table 6-1 shows the file systems.'

Table 6-1 Example IBM Spectrum Scale file systems

IBM Spectrum Scale file systems			
File system name	Mount point	Need space management? (Yes or No)	Reserved for IBM Spectrum Archive EE? (Yes or No)
gpfs	/ibm/glues	YES	YES

Table 6-2 shows the logical tape library.

Table 6-2 Example logical tape library

Logical Tape library			
<b>Tape library information</b>			
Tape library (L-Frame) Serial Number		78-A4274	
Starting SCSI Element Address of the logical tape library for IBM Spectrum Archive EE (decimal and hex)		1033dec = 409hex	
Logical tape library serial number (L-Frame S/N + "0" + SCSI starting element address in hex)		78A4274-0-409 = 78A42740409	
<b>Tape Drive information</b>			
Drive Serial number	Assigned IBM Spectrum Scale node	CP? (Yes or No)	Linux device name in the node
9A700M0029	htohru9	YES	/dev/IBMtape13
1068000073	htohru9	NO	/dev/IBMtape11

Table 6-3 shows the nodes.

Table 6-3 Example IBM Spectrum Scale nodes

IBM Spectrum Scale nodes			
IBM Spectrum Scale node name	Installing IBM Spectrum Archive EE? (Yes or No)	Tape drives assigned to this node (Serial number)	CP enabled tape drive (Serial number)
htohru9	YES	9A700M0029, 1068000073	9A700M0029

Figure 6-1 shows an example of a TS3500 GUI window that you use to display the starting SCSI element address of a TS3500 logical library. You must record the decimal value (starting address) to calculate the associated logical library serial number, as shown in Table 6-2 on page 97. You can open this window if you check for the details for a specific logical library.

Logical Library LTFS_L05	
Type	LTO
Drive element address range	257 - 258
Drives	 Frame 4 Row 5  Frame 4 Row 6
Storage element address range	1033 - 2747
Cartridges	6
Virtual I/O element address range	789 - 1023

Figure 6-1 Obtain the starting SCSI element address of a TS3500 logical library

Table 6-4 shows a blank file systems worksheet.

Table 6-4 Example IBM Spectrum Scale file systems

IBM Spectrum Scale file systems			
File system name	Mount point	Need space management? (Yes or No)	Reserved for IBM Spectrum Archive EE? (Yes or No)

Table 6-5 shows a blank logical tape library worksheet.

Table 6-5 Example logical tape library

Logical Tape library			
<b>Tape library information</b>			
Tape library (L-Frame) Serial Number			
Starting SCSI Element Address of the logical Tape Library for IBM Archive EE (decimal and hex)			
Logical tape library serial number (L-Frame S/N + "0" + SCSI starting element address in hex)			
<b>Tape Drive information</b>			
Drive Serial number	Assigned IBM Spectrum Scale node	CP? (Yes or No)	Linux device name in the node

Table 6-6 shows a blank nodes worksheet.

Table 6-6 Example IBM Spectrum Scale nodes

IBM Spectrum Scale nodes			
IBM Spectrum Scale node name	Installing IBM Spectrum Archive EE? (Yes or No)	Tape drives assigned to this node (Serial number)	CP enabled tape drive (Serial number)

## 6.1.2 Obtaining configuration information

To obtain the information about your environment that is required for configuring IBM Spectrum Archive EE, complete the following steps:

1. Log on to the operating system as a root user.
2. Start GPFS (if it is not started already) by running the following command (see Example 6-1 on page 101):  

```
# mmstartup -a
```
3. Mount GPFS (if it is not already mounted) by running the following command (see Example 6-1 on page 101):  

```
# mmmount all
```
4. Obtain a list of all GPFS file systems that exist in the IBM Spectrum Scale cluster by running the following command (see Example 6-1 on page 101):  

```
# mm1sfs all
```
5. Go to the Configuration worksheet (provided in 6.1.3, “Configuring key-based login with OpenSSH” on page 102) and enter the list of file system names in the GPFS file systems table.
6. Plan the GPFS file system that was used to store IBM Spectrum Archive EE internal data. For more information, see 6.1.4, “Preparing the IBM Spectrum Scale file system for IBM Spectrum Archive EE” on page 102.
7. Go to the Configuration worksheet and enter the GPFS file system that is used to store IBM Spectrum Archive EE internal data into Table 6-1 on page 97.
8. Obtain a list of all IBM Spectrum Scale nodes in the IBM Spectrum Scale cluster by running the following command (see Example 6-1 on page 101):  

```
# mm1snode
```
9. Go to the Configuration worksheet and enter the list of IBM Spectrum Scale nodes and whether the IBM Spectrum Archive EE is installed on the node in Logical Tape Library table.
10. Obtain the logical library serial number, as described in the footnote of Table 6-2 on page 97. For more information and support, see the IBM Knowledge Center for your specific tape library.
11. Go to the Configuration worksheet and enter the logical library serial number that was obtained in the previous step into the IBM Spectrum Scale nodes table.
12. Obtain a list of all tape drives in the logical library that you plan to use for the configuration of IBM Spectrum Archive EE. For more information, see the IBM Knowledge Center for your specific tape library.
13. Go to the Configuration worksheet and enter the tape drive serial numbers that were obtained through the previous step into the IBM Spectrum Scale nodes table.
14. Assign each drive to one of the IBM Spectrum Archive EE nodes that are listed in the Logical Library Table in the Configuration worksheet and add that information to the IBM Spectrum Scale nodes table.
15. Assign at least one CP to each of the IBM Spectrum Archive EE nodes and enter whether each drive is a CP drive in the IBM Spectrum Scale nodes section of the Configuration worksheet.
16. Go to the Configuration worksheet and update the Logical Tape Library table with the tape drive assignment and CP drive information by adding the drive serial numbers in the appropriate columns.

Keep the completed configuration worksheet available for reference during the configuration process. Example 6-1 shows how to obtain the information for the worksheet.

*Example 6-1 Obtain the IBM Spectrum Scale required information for the configuration worksheet*

---

```
[root@ltfs97 ~]# mmstartup -a
Fri Apr 5 14:02:32 JST 2013: mmstartup: Starting GPFS ...
htohru9.ltd.sdl: The GPFS subsystem is already active.
```

```
[root@ltfs97 ~]# mmmount all
Fri Apr 5 14:02:50 JST 2013: mmmount: Mounting file systems ...
```

```
[root@ltfs97 ~]# mmlsfs all
```

File system attributes for /dev/gpfs:

=====

flag	value	description
-f	8192	Minimum fragment size in bytes
-i	512	Inode size in bytes
-I	16384	Indirect block size in bytes
-m	1	Default number of metadata replicas
-M	2	Maximum number of metadata replicas
-r	1	Default number of data replicas
-R	2	Maximum number of data replicas
-j	cluster	Block allocation type
-D	nfs4	File locking semantics in effect
-k	all	ACL semantics in effect
-n	1	Estimated number of nodes that will mount file system
-B	262144	Block size
-Q	none	Quotas enforced
	none	Default quotas enabled
--filesetdf	No	Fileset df enabled?
-V	13.23 (3.5.0.7)	File system version
--create-time	Tue Mar 26 17:06:14 2013	File system creation time
-u	Yes	Support for large LUNs?
-z	Yes	Is DMAPI enabled?
-L	4194304	Logfile size
-E	Yes	Exact mtime mount option
-S	No	Suppress atime mount option
-K	whenpossible	Strict replica allocation option
--fastea	Yes	Fast external attributes enabled?
--inode-limit	154112	Maximum number of inodes
-P	system	Disk storage pools in file system
-d	ltfs97nsd	Disks in file system
--perfileset-quota	no	Per-fileset quota enforcement
-A	yes	Automatic mount option
-o	none	Additional mount options
-T	/ibm/glues	Default mount point
--mount-priority	0	Mount priority

```
[root@ltfs97 ~]# mmlsnode
GPFS nodeset Node list
```

---

htohru9	htohru9
---------	---------

---

### 6.1.3 Configuring key-based login with OpenSSH

IBM Spectrum Archive EE uses the Secure Shell (SSH) protocol for secure file transfer and requires key-based login with OpenSSH for the root user.

To use key-based login with OpenSSH, it is necessary to generate SSH key files and append the public key file from each node (including the local node) to the `authorized_keys` file in the `~root/.ssh` directory.

The following points must be considered:

- ▶ This procedure must be performed on *all* IBM Spectrum Archive EE nodes.
- ▶ After completing this task, a root user on any node in an IBM Spectrum Archive EE cluster can run any commands on any node remotely without providing the password for the root on the remote node. It is preferable that the cluster is built on a closed network. If the cluster is within a firewall, all ports can be opened. For more information, see 4.3.1, “Extracting binary rpm files from an installation package” on page 58 and 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60.

To configure key-based login with OpenSSH, complete the following steps:

1. If the `~root/.ssh` directory does not exist, create it by running the following command:

```
mkdir ~root/.ssh
```

2. If the root user does not have SSH keys, generate them by running the `ssh-keygen` command and entering CR at all prompts.

**Important:** You can verify whether the root user has a public key by locating the `id_rsa` and `id_rsa.pub` files under the `/root/.ssh/` directory. If these files do not exist, you must generate them.

3. After the key is generated, copy the key to each server that requires a key-based login for OpenSSH by running the following command:

```
ssh-copy-id root@<server>
```

4. Repeat these steps on each IBM Spectrum Archive EE node.

### 6.1.4 Preparing the IBM Spectrum Scale file system for IBM Spectrum Archive EE

Complete this task to create and mount the IBM Spectrum Scale file system before IBM Spectrum Archive EE is configured.

Before you make any system upgrades or major configuration changes to your GPFS or IBM Spectrum Scale cluster, review your GPFS or IBM Spectrum Scale documentation and consult IBM Spectrum Scale frequently asked question (FAQ) information that applies to your version of IBM Spectrum Scale. For more information about the IBM Spectrum Scale FAQ, see the Cluster products IBM Knowledge Center at this website:

[https://www.ibm.com/support/knowledgecenter/en/STXKQY\\_4.2.3/com.ibm.spectrum.scale.v4r23.doc/b1lins\\_1scfg.htm](https://www.ibm.com/support/knowledgecenter/en/STXKQY_4.2.3/com.ibm.spectrum.scale.v4r23.doc/b1lins_1scfg.htm)

Before you begin this procedure, ensure that the following prerequisites are met:

- ▶ IBM Spectrum Scale is installed on each of the IBM Spectrum Archive EE nodes.
- ▶ The IBM Spectrum Scale cluster is created and all of the IBM Spectrum Archive EE nodes belong to the cluster.

IBM Spectrum Archive EE requires space for the file metadata, which is stored in the LTFS metadata directory. The metadata directory can be stored in its own GPFS file system, or it can share the GPFS file system that is being space-managed with IBM Spectrum Archive EE.

The file system that is used for the LTFS metadata directory must be created and mounted before the IBM Spectrum Archive EE configuration is performed. The following requirements apply to the GPFS file system that is used for the LTFS metadata directory:

- ▶ The file system must be mounted and accessible from all of the IBM Spectrum Archive EE nodes in the cluster.
- ▶ The GPFS file system (or systems) that are space-managed with IBM Spectrum Archive EE must be DMAPI enabled.

To create and mount the GPFS file system, complete the following steps:

1. Create a network shared disk (NSD), if necessary, by running the following command. It is possible to share an existing NSD with another GPFS file system.

```
# mmcrnsd -F nsd.list -v no
<<nsd.list>>
%nsd: device=/dev/dm-3
      nsd=nsd00
      servers=ltfs01, ltfs02, ltfs03, ltfs04
      usage=dataAndMetadata
```

2. Start the GPFS service (if it is not started already) by running the following command:

```
# mmstartup -a
```

3. Create the GPFS file system by running the following command. For more information about the file system name and mount point, see 6.1.1, “Configuration worksheet tables” on page 96.

```
# mmcrfs /dev/gpfs nsd00 -z yes -T /ibm/glues
```

In this example, /dev/gpfs is the file system name and /ibm/glues is the mount point. For a separate file system that is used only for the LTFS metadata directory, you do not need to use the -z option. Generally, if a GPFS file system is not intended to be IBM Spectrum Archive EE managed, it should not be DMAPI-enabled, so the -z option should not be specified. The preferred configuration is to have one file system with DMAPI-enabled.

4. Mount the GPFS file system by running the following command:

```
# mmmount gpfs -a
```

For more information about the **mmmount** command, see either of the following resources:

- ▶ *General Parallel File System Version 4 Release 1.0.4 Advanced Administration Guide*, SC23-7032-01
- ▶ *IBM Spectrum Scale: Administration Guide*, which is available at this website:  
<https://www.ibm.com/support/knowledgecenter/STXKQY>

## 6.1.5 Enabling control paths in the logical library

As explained in 1.4, “IBM Spectrum Archive EE cluster configuration introduction” on page 13, each of the IBM Spectrum Archive EE nodes must have its own control path in the logical library. For this reason, it is necessary to configure the required number of control paths in the tape library. For more information, see the IBM Knowledge Center for your specific tape library.

To identify which tape drives should be enabled as control paths, see Table 6-5 on page 99 in the IBM Spectrum Archive EE Configuration worksheet.

Enable the appropriate control paths through your tape library management GUI.

**Important:** If you want to set up multipathing (CPF) for the tape drives that are attached to the IBM Spectrum Archive EE nodes, you must first configure the tape device driver `lin_tape` to enable control path failover. Follow the procedure in the *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130 from the IBM Support portal.

## 6.2 Configuring IBM Spectrum Archive EE

The topics in this section describe how to use the `ltfsee_config` command to configure IBM Spectrum Archive EE in a single node or multiple node environment. Instructions for removing a node from an IBM Spectrum Archive EE configuration are also provided.

### 6.2.1 The `ltfsee_config` utility

Use the `ltfsee_config` command-line utility to configure the IBM Spectrum Archive EE for single node or multiple node environment. You must have root user authority to use this command. This command also can be used to check an existing IBM Spectrum Archive EE configuration. The utility operates in interactive mode and guides you step-by-step through the required information that you must provide.

**Reminder:** All of the command examples use the command without the full file path name because we added the IBM Spectrum Archive EE directory (`/opt/ibm/ltfsee/bin`) to the `PATH` variable.

The `ltfsee_config` command-line tool is shown in the following example and includes the following options:

```
ltfsee_config -m <mode> [options]
```

► **-m**

<mode> and [options] can be one of the following items:

– **CLUSTER [-c]**

Creates an IBM Spectrum Archive EE cluster environment and configures a user-selected IBM Spectrum Scale (GPFS) file system to be managed by the IBM Spectrum Archive or used for its metadata. The user must run this command one time from one of the IBM Spectrum Archive nodes. Running the command a second time modifies the file systems settings of the existing cluster.

– **ADD\_CTRL\_NODE [-g | -c]**

Adds the local node as the control (MMM) node to a tape library in an existing IBM Spectrum Archive EE environment, and configures its drives and node group. There can be one or two control nodes per tape library.

**Note:** Even if you configure two control nodes per tape library, you still only run ADD\_CTRL\_NODE once per tape library.

– **ADD\_NODE [-g | -c]**

Adds the local node (as a non-control node) to a tape library, and configure its drives and node group. You choose whether or not the node is a control node as redundancy.

– **SET\_CTRL\_NODE**

Configure or reconfigure one or two control nodes and select one node to be active at the next start of IBM Spectrum Archive EE.

– **REMOVE\_NODE [-N <node\_id>] [-f]**

Removes the node and the drives configured for that node from the existing configuration.

– **DELETE\_DB**

Renames the global.db and libresources.<library\_id>.db files under the IBM Spectrum Scale <CONFIGANDMETA directory>/l1tfsee/config directory by appending .<yyyymmddHHMMSS>.save to their file names to serve as backups. ADD\_CTRL\_NODE, ADD\_NODE, or both are required on all of the nodes after this option is invoked.

– **INFO**

Shows the current configuration of this cluster.

► **Options:**

– **-c**

Check and show the existing cluster or node configuration, without configuring or modifying it.

– **-g**

Assign the node to a node group that is selected or specified by user. If **-g** is not used, the node is added to the default node group, which is named G0 if it did not exist before.

– **-N**

Remove a non-local node by specifying its node ID. If **-N** is not used, the local node is removed.

– **-f**

Force node removal. If **-f** is not used, an attempt to remove a control node fails and the configuration remains unchanged. When a control node is removed by using **-f**, other nodes from the same library and the drives that are configured for those nodes are also removed. To avoid removing multiple nodes, consider first setting another configured non-control node from the same library as the control node (**SET\_CTRL\_NODE**).

## 6.2.2 Configuring a single node cluster

Before you begin this procedure, ensure that all of the tasks that are described in 6.1, “Configuration prerequisites” on page 96 are met. Figure 6-2 shows an illustration of a single-node configuration that is described in this section.

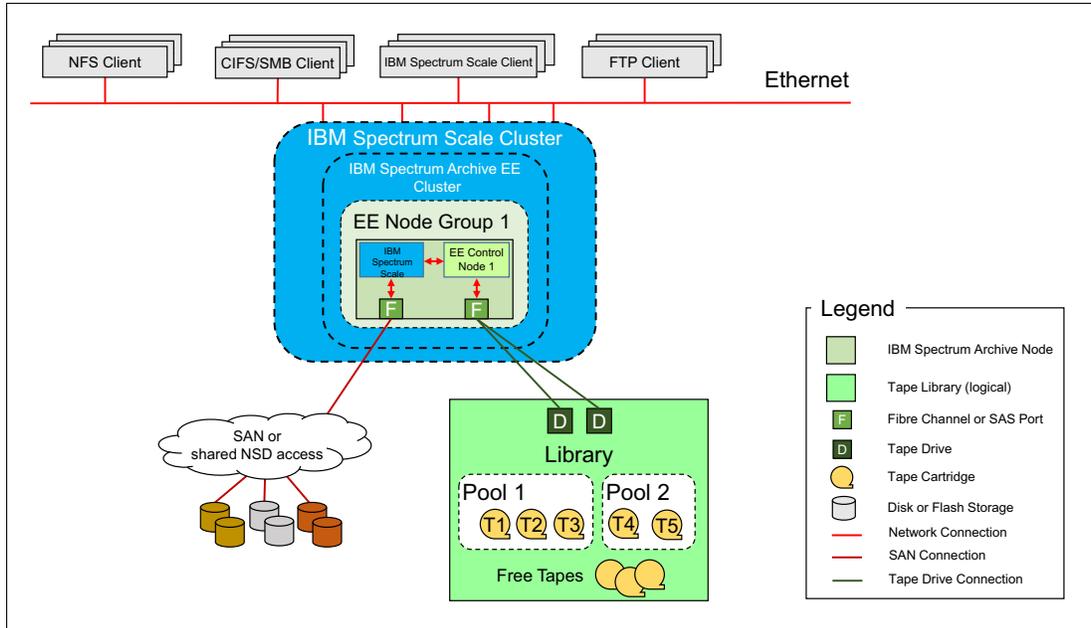


Figure 6-2 IBM Spectrum Archive single-node configuration

The steps in this section must be performed only on one node of an IBM Spectrum Archive EE cluster environment. If you plan to have only one IBM Spectrum Archive EE node, this is a so-called *single-node cluster* setup.

If you plan to set up a multi-node cluster environment for IBM Spectrum Archive EE, this configuration mode must be performed once, and only on a node of your choice of your cluster environment. All other nodes must be added. To do so, see 6.2.3, “Configuring a multiple-node cluster” on page 110.

To configure a single-node cluster for IBM Spectrum Archive EE, complete the following steps:

1. Log on to the operating system as a root user.
2. Start GPFS (if it is not already started) by running the following command:

```
# mmstartup -a
```
3. Mount the GPFS file system (if it is not already mounted) by running the following command:

```
# mmmount all
```
4. Start the IBM Spectrum Archive EE configuration utility with the `-m CLUSTER` option by running the following command and answering the prompted questions:

```
# ltfsee_config -m CLUSTER
```

Example 6-2 shows the successful run of the `ltfsee_config -m CLUSTER` command during the initial IBM Spectrum Archive EE configuration on the lab setup that was used for this book.

*Example 6-2 Run the ltfsee\_config -m CLUSTER command*

---

```
[root@ltfsm11 ~]# /opt/ibm/ltfsee/bin/ltfsee_config -m CLUSTER
CLUSTER mode starts .

## 1. Check whether the cluster is already created ##
Cluster is not configured, configuring the cluster.

## 2. Check prerequisite on cluster ##
Cluster name: ltfsm12-ltfsm11.tuc.stglabs.ibm.com
ID: 12003238441805965800
Successfully validated the prerequisites.

## 3. List file systems in the cluster ##
Retrieving IBM Spectrum Scale (GPFS) file systems...
** Select a file system for storing IBM Spectrum Archive Enterprise Edition
configuration and internal data.
    Input the corresponding number and press Enter
    or press q followed by Enter to quit.

    File system
    1. /dev/gpfs  Mount point(/ibm/gpfs)  DMAPI(Yes)
    q. Quit

Input number > 1

** Select file systems to configure for IBM Spectrum Scale (GPFS) file system
for Space Management.
    Input the corresponding numbers and press Enter
    or press q followed by Enter to quit.
    Press a followed by Enter to select all file systems.
    Multiple file systems can be specified using comma or white space
    delimiters.

    File system
    1. /dev/gpfs  Mount point(/ibm/gpfs)
    a. Select all file systems
    q. Quit

Input number > 1

## 4. Configure Space Management ##
Disabling unnecessary daemons...
Editing Space Management Client settings...
Restarting Space Management service...
Terminating dsmwatchd.....
Terminating dsmwatchd.....
Starting dsmmigfs.....
Configured space management.

## 5. Add selected file systems to the Space Management ##
Added the selected file systems to the space management.
```

```
## 6. Store the file systems configuration and dispatch it to all nodes ##
Storing the file systems configuration...
Copying ltfsee_config.filesystem file...
Stored the cluster configuration and dispatched the configuration file.

## 7. Create metadata directories and configuration parameters file ##
Created metadata directories and configuration parameters file.

CLUSTER mode completed .
```

---

**Important:** During the first run of the `ltfsee_config -m CLUSTER` command, you see the following error:

```
No file system is DMAPI enabled.
At least one file system has to be DMAPI enabled to use IBM Spectrum Archive
Enterprise Edition.
Enable DMAPI of more than one IBM Spectrum Scale (GPFS) file systems and try
again.
```

Ensure that DMAPI is turned on correctly, as described in 6.1.4, “Preparing the IBM Spectrum Scale file system for IBM Spectrum Archive EE” on page 102. You can use the following command sequence to enable DMAPI support for your GPFS file system (here the GPFS file system name that is used is `gpfs`):

```
# mmumount gpfs
mmumount: Unmounting file systems ...
# mmchfs gpfs -z yes
# mmmount gpfs
mmmount: Mounting file systems ...
```

5. Run the IBM Spectrum Archive EE configuration utility by running the following command and answering the prompted questions:

```
# ltfsee_config -m ADD_CTRL_NODE
```

Example 6-3 shows the successful run of the `ltfsee_config -m ADD_CTRL_NODE` command during initial IBM Spectrum Archive EE configuration on the lab setup that was used for this book.

*Example 6-3 Run the `ltfsee_config -m ADD_CTRL_NODE` command*

---

```
[root@l1tfsml1 ~]# /opt/ibm/ltfsee/bin/ltfsee_config -m ADD_CTRL_NODE
ADD_CTRL_NODE mode starts .
```

```
## 1. Check whether the cluster is already created ##
Cluster is already created and configuration file ltfsee_config.filesystem
exists.
```

```
## 2. Check prerequisite on node ##
Successfully validated the prerequisites.
```

```
## 3. IBM Spectrum Scale (GPFS) Configuration for Performance Improvement ##
Setting worker1Threads=400
Setting dmapWorkerThreads=64
Configured IBM Spectrum Scale (GPFS) performance related settings.
```

```
## 4. Configure Space Management ##
```

```
Disabling unnecessary daemons...
Editing Space Management Client settings...
Restarting Space Management service...
Terminating dsmwatchd.....
Terminating dsmwatchd.....
Starting dsmmigfs.....
Configured space management.
```

```
## 5. Add this node to a tape library ##
```

```
Number of logical libraries with assigned control node: 0
Number of logical libraries available from this node: 1
Number of logical libraries available from this node and with assigned control
node: 0
```

```
** Select the tape library from the following list
and input the corresponding number. Then, press Enter.
```

```
Model      Serial Number
1. 3576-MTL 000001300228_LLC
q. Return to previous menu
```

```
Input Number > 1
Input Library Name (alpha numeric or underscore, max 16 characters) >
lib_ltfsm11
Added this node (ltfsm11.tuc.stglabs.ibm.com, node id 2) to library lib_ltfsm11
as its control node.
```

```
## 6. Add this node to a node group ##
Added this node (ltfsm11.tuc.stglabs.ibm.com, node id 2) to node group G0.
```

```
## 7. Add drives to this node ##
```

```
** Select tape drives from the following list.
Input the corresponding numbers and press Enter
or press q followed by Enter to quit.
Multiple tape drives can be specified using comma or white space delimiters.
```

```
Model      Serial Number
1. ULT3580-TD6 1013000655
2. ULT3580-TD6 1013000688
3. ULT3580-TD6 1013000694
a. Select all tape drives
q. Exit from this Menu
```

```
Input Number > a
Selected drives: 1013000655:1013000688:1013000694.
Added the selected drives to this node (ltfsm11.tuc.stglabs.ibm.com, node id
2).
```

```
## 8. Configure LE+ component ##
Creating mount point...
Mount point folder '/ltfs' exists.
Use this folder for the LE+ component mount point as LE+ component assumes this
folder.
Configured LE+ component.
```

```
## 9. Enabling system log ##
Restarting rsyslog...
System log (rsyslog) is enabled for IBM Spectrum Archive Enterprise Edition.

ADD_CTRL_NODE mode completed .
```

To summarize, EE Node 1 must run `ltfsee_config -m CLUSTER` and `ltfsee_config -m ADD_CTRL_NODE` to complete this single-node configuration.

If you are configuring multiple nodes for IBM Spectrum Archive EE, continue to 6.2.3, “Configuring a multiple-node cluster” on page 110.

### 6.2.3 Configuring a multiple-node cluster

To add nodes to form a multiple-node cluster configuration after the first node is configured, complete this task. With the release of IBM Spectrum Archive EE V1.2.4.0, a redundant control node can be set for failover scenarios.

When configuring any multiple-node clusters, set a secondary node as a redundant control node for availability features. The benefits of having redundancy are explained in 7.6, “IBM Spectrum Archive EE automatic node failover” on page 145.

Figure 6-3 shows an illustration of a multiple-node cluster configuration that is described in this section.

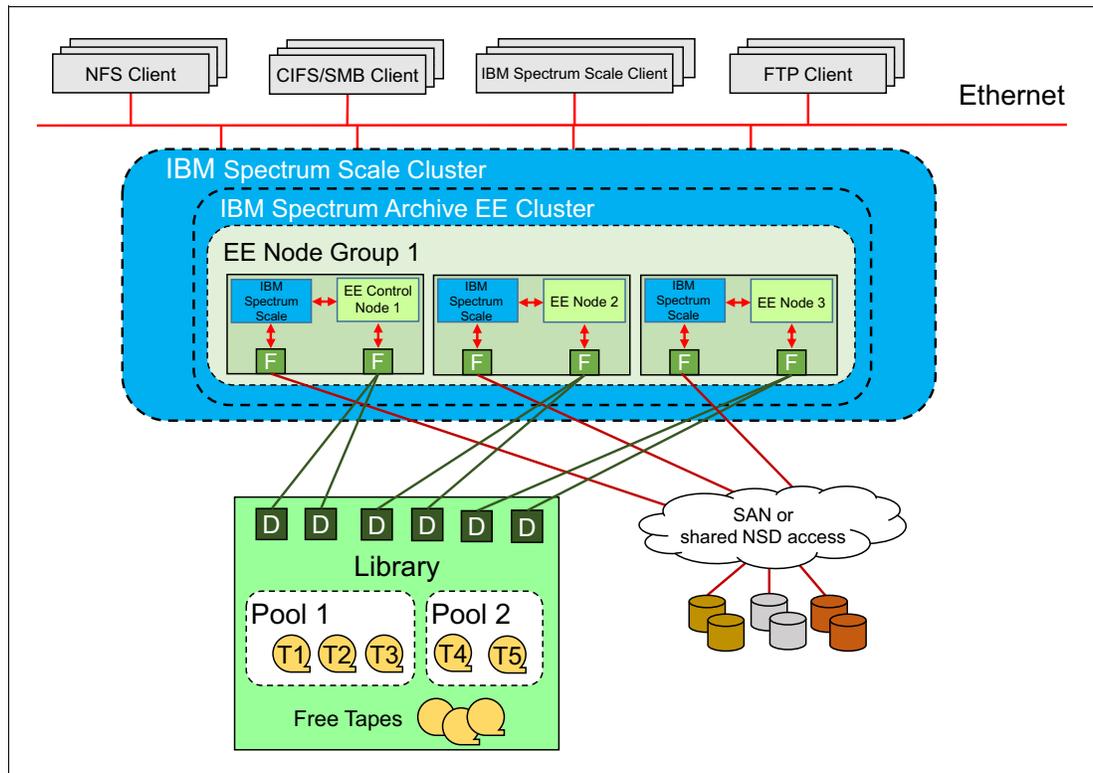


Figure 6-3 IBM Spectrum Archive multiple-node cluster configuration

Before configuring more nodes, ensure that all tasks that are described in 6.1, “Configuration prerequisites” on page 96 are completed and that the first node of the cluster environment is configured, as described in 6.2.2, “Configuring a single node cluster” on page 106.

To configure another node for a multi-node cluster setup for IBM Spectrum Archive EE, complete the following steps:

1. Log on to the operating system as a root user.
2. Start GPFS (if it is not already started) by running the following command:  

```
# mmstartup -a
```
3. Mount the GPFS file system on all nodes in the IBM Spectrum Scale cluster (if it is not already mounted) by running the following command:  

```
# mmmount all -a
```
4. Start the IBM Spectrum Archive EE configuration utility with the `-m ADD_NODE` option by running the following command and answering the prompted questions:  

```
# /opt/ibm/ltfsee/bin/ltfsee_config -m ADD_NODE
```

**Important:** This step must be performed on all nodes except for the first node that was configured in 6.2.2, “Configuring a single node cluster” on page 106.

Example 6-4 shows how to add a secondary node and set it as a redundant control node by running `ltfsee_config -m ADD_NODE`. In step 5 of the command, after selecting which library to add the node to, a prompt will appear asking to make the node a redundant control node. Enter `y` to make the second node a redundant control node. Only two nodes per library can be control nodes. If there are more than two nodes added to the cluster, enter `n` for each additional node.

*Example 6-4 Adding secondary node as a redundant control node*

---

```
[root@l1fsm12 ~]# ltfsee_config -m ADD_NODE
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
ADD_NODE
ADD_NODE mode starts .

## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file ltfsee_config.filesystem
exists.

## 2. Check prerequisite on node ##
Successfully validated the prerequisites.

## 3. IBM Spectrum Scale (GPFS) Configuration for Performance Improvement ##
Setting workerThreads=512
Setting dmapiWorkerThreads=64
Configured IBM Spectrum Scale (GPFS) performance related settings.

## 4. Configure space management ##
Disabling unnecessary daemons...
Editing Space Management Client settings...
Deactivating failover operations on the node.
Restarting Space Management service...
Stopping the HSM service.
Terminating dsmwatchd.....
Starting the HSM service.
Starting dsmmigfs.....
Activating failover operations on the node.
Configured space management.
```

## 5. Add this node to a tape library ##

The number of logical libraries with the assigned control node: 2

The number of logical libraries available from this node: 1

The number of logical libraries available from this node and with assigned control node: 1

\*\* Select the tape library from the following list  
and input the corresponding number. Then press Enter.

Library id	Library name	Control node
1. 0000013FA0520411	ltfsee_lib	9.11.120.198

q. Exit from this Menu

Input Number > 1

Add this node as a control node for control node redundancy(y/n)?

Input >y

The node ltfsm12(9.11.120.201) has been added as a control node for control node redundancy

Added this node (ltfsm12, node id 2) to library ltfsee\_lib.

## 6. Add this node to a node group ##

Added this node (ltfsm12, node id 2) to node group G0.

## 7. Add drives to this node ##

\*\* Select tape drives from the following list.

Input the corresponding numbers and press Enter  
or press 'q' followed by Enter to quit.

Multiple tape drives can be specified using comma or white space delimiters.

	Model	Serial Number
1.	ULT3580-TD5	1068093078
2.	ULT3580-TD5	1068093084

a. Select all tape drives  
q. Exit from this Menu

Input Number > a

Selected drives: 1068093078:1068093084.

Added the selected drives to this node (ltfsm12, node id 2).

## 8. Configure the LE+ component ##

Creating mount point...

Mount point folder '/ltfs' exists.

Use this folder for the LE+ component mount point as LE+ component assumes this folder.

Former saved configuration file exists which holds the following information:

=== difference /etc/ltfs.conf.local.rpmsave from /etc/ltfs.conf.local ===

=== end of difference ===

Do you want to use the saved configuration (y/n)?

Input > y

The LE+ component configuration is restored from a saved configuration.

Configured the LE+ component.

ADD\_NODE mode completed .

---

To summarize, you ran the following configuration options on EE Node 1 in 6.2.2, “Configuring a single node cluster” on page 106:

- ▶ `ltfsee_config -m CLUSTER`
- ▶ `ltfsee_config -m ADD_CTRL_NODE`

For each additional IBM Spectrum Archive node in EE Node Group 1, run the `ltfsee_config -m ADD_NODE` command. For example, in Figure 6-3 on page 110, you must run `ltfsee_config -m ADD_NODE` on both EE Node 2 and EE Node 3.

If you require multiple tape library attachments, go to 6.2.4, “Configuring a multiple-node cluster with two tape libraries” on page 113.

## 6.2.4 Configuring a multiple-node cluster with two tape libraries

Starting with IBM Spectrum Archive V1R2, IBM Spectrum Archive supports the Multiple Tape Library Attachment feature in a single IBM Spectrum Scale cluster. This feature allows for data replication to pools in separate libraries for more data resiliency, and allows for total capacity expansion beyond a single library limit.

The second tape library can be the same tape library model as the first tape library or can be a different tape library model. These two tape libraries can be connected to a IBM Spectrum Scale cluster in a single site or can be placed in metro distance (less than 300 km) locations through IBM Spectrum Scale synchronous mirroring (stretched cluster).

For more information about synchronous mirroring by using IBM Spectrum Scale replication, see the following website:

[https://www.ibm.com/support/knowledgecenter/en/STXKQY\\_4.2.3/com.ibm.spectrum.scale.v4r23.doc/b11adv\\_continuous\\_replication\\_SSdata.htm](https://www.ibm.com/support/knowledgecenter/en/STXKQY_4.2.3/com.ibm.spectrum.scale.v4r23.doc/b11adv_continuous_replication_SSdata.htm)

**Important:** Stretched cluster is available for distances shorter than 300 km. For longer distances, the Active File Management (AFM) feature of IBM Spectrum Scale should be used with IBM Spectrum Archive. The use of AFM is with two different IBM Spectrum Scale clusters with one instance of IBM Spectrum Archive at each site. For more details about IBM Spectrum Scale AFM support, see 2.3.5, “Active File Management” on page 29.

To add nodes to form a multiple-node cluster configuration with two tape libraries after the first node is configured, complete this task. Figure 6-4 shows an illustration of the configuration with two tape libraries.

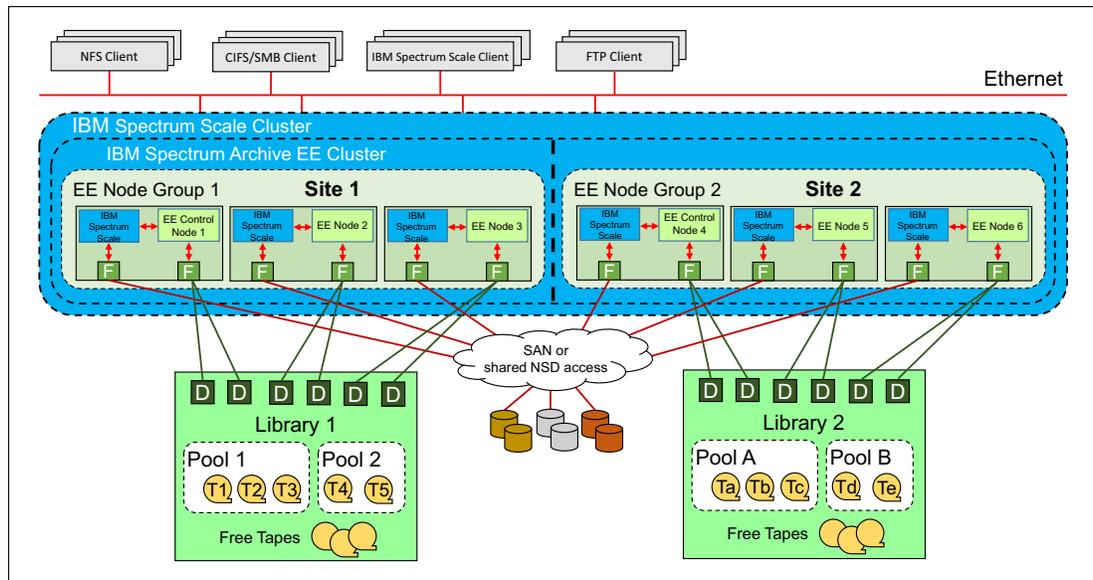


Figure 6-4 IBM Spectrum Archive multiple-node cluster configuration across two tape libraries

Before configuring more EE nodes, ensure that all tasks that are described in 6.1, “Configuration prerequisites” on page 96 are completed and that the first node of the cluster environment is configured, as described in 6.2.2, “Configuring a single node cluster” on page 106.

To configure the nodes at the other location for a multiple-node two-tape library cluster setup for IBM Spectrum Archive EE, complete the following steps:

1. Run the IBM Spectrum Archive EE configuration utility by running the following command and answering the prompted questions:

```
# /opt/ibm/ltfsee/bin/ltfsee_config -m ADD_CTRL_NODE
```

Using Figure 6-4 as an example, the `ltfsee_config -m ADD_CTRL_NODE` command is run on EE Node 4.

2. Run the IBM Spectrum Archive EE configuration utility on all the remaining EE nodes at the other location by running the following command and answering the prompted questions:

```
# /opt/ibm/ltfsee/bin/ltfsee_config -m ADD_NODE
```

Using Figure 6-4 as an example, the `ltfsee_config -m ADD_NODE` command is run on EE Node 5 and EE Node 6.

## 6.2.5 Modifying a multiple-node configuration for control node redundancy

If users are upgrading to the IBM Spectrum Archive EE V1.2.4.0 from a previous version and have a multiple-node configuration, users must manually set a redundant control node. See 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60 on how to perform upgrades. To modify the configuration to set a secondary node to be a redundant control node, IBM Spectrum Archive EE must not be running.

Run **ltfsee stop**, to stop IBM Spectrum Archive EE and LE+ and check that the processes have stopped by running the **pidof mmm** command on the control node and the **pidof ltf** command on each node. After IBM Spectrum Archive EE has stopped, run **ltfsee\_config -m SET\_CTRL\_NODE** to modify the configuration to add a redundant control node.

Example 6-5 shows the output of **ltfsee\_config -m SET\_CTRL\_NODE** to create a redundant control node. In this example, the cluster has two libraries connected and will only perform setting a redundant control node on one of the two libraries. Repeat the same steps and select the second library to make a redundant control node for the second library.

*Example 6-5 Setting existing node to be redundant control node*

---

```
[root@ltfsm11 ~]# ltfsee_config -m SET_CTRL_NODE
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
SET_CTRL_NODE
SET_CTRL_NODE mode starts .

## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file ltfsee_config.filesystem
exists.
## 2. Control node configuration.

** Select a library to set control nodes.

    Libraries
    1:0000013FA0520411
    2:0000013FA0520412
    q.Quit

Input number >1
Set the control nodes for the library 0000013FA0520411.

** Select 1 or 2 nodes for redundant control nodes from the following list.
    They can be specified using comma or white space delimiters.
    Nodes marked [x] are the current redundant configured nodes.

    Nodes
    1:[x]ltfsm11
    2:[_]ltfsm12
    q.Quit

Input number >1 2

## 3. Select control node to be active ##
    The following nodes are selected as redundant nodes.
    Select a node that will be active in the next LTFs-EE run.

    Nodes
    1:ltfsm11
    2:ltfsm12
    q.Quit

Input number >1

The node ebisu(9.11.120.198) has been set to be active for library ltfsee_lib
```

---

After successfully setting up redundant control nodes for each library cluster, start Spectrum Archive EE by running the `ltfsee start` command. Then, run `ltfsee info nodes` to verify that each node was started up properly and is available. You should also see that there are two control nodes per library. Example 6-6 shows output from a multiple-node configuration with two tape libraries

Example 6-6 `ltfsee info nodes`

---

```
[root@ltfsm11 ~]# ltfsee info nodes
```

Node ID	Status	Node IP	Drives	Ctrl Node	Library	Node Group
4	Available	9.11.120.224	2	Yes	ltfsee_lib2	G0
ltfsm14						
3	Available	9.11.120.207	2	Yes(Active)	ltfsee_lib2	G0
ltfsm13						
2	Available	9.11.120.201	2	Yes	ltfsee_lib1	G0
ltfsm12						
1	Available	9.11.120.198	2	Yes(Active)	ltfsee_lib1	G0
ltfsm11						

---

## 6.3 First-time start of IBM Spectrum Archive EE

To start IBM Spectrum Archive EE the first time, complete the following steps:

1. Identify the Linux device name for the logical tape library (IBM SCSI changer name) you plan to use and that you configured in 6.2.2, “Configuring a single node cluster” on page 106 or 6.2.3, “Configuring a multiple-node cluster” on page 110 by running the following command:

```
# cat /proc/scsi/IBMchanger*
```

**Note:** Starting with IBM Spectrum Archive V1.2.4.0, `ltfsee start` handles starting up both LE+ components and EE components. In addition to starting the two main components, it also starts a monitoring daemon on each node.

2. Check that the following embedded, customized Tivoli Storage Manager for Space Management (HSM) client components are running:

```
# ps -ef|grep dsm
```

You should see three recall daemons running.

3. Start the IBM Spectrum Archive EE program by running the following command:

```
/opt/ibm/ltfsee/bin/ltfsee start
```

**Important:** If the `ltfsee start` command does not return after several minutes, it might be either because tapes are being unloaded or because the firewall is running. The firewall service must be disabled on the IBM Spectrum Archive EE nodes. For more information, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60.

You might be able to skip steps 1 on page 116 and 2 on page 116 the next time you start IBM Spectrum Archive EE because they are not functional-related, just informational.

Example 6-7 shows all of the steps and the output when IBM Spectrum Archive EE was started the first time. During the first start, you might discover a warning message, as shown in the following example:

```
The authenticity of host 'localhost (:::1)' can't be established.
RSA key fingerprint is 8f:56:95:fe:9c:eb:37:7f:95:b1:21:b9:45:d6:91:6b.
Are you sure you want to continue connecting (yes/no)?
```

This message is normal during the first start and you can easily continue by entering yes and pressing **Enter**.

*Example 6-7 Start IBM Spectrum Archive EE the first time*

---

```
[root@ltfsm11 ~]# cat /proc/scsi/IBM*
lin_tape version: 3.0.13
lin_tape major number: 246
Attached Changer Devices:
Number  model      SN                HBA                SCSI                FO Path
0       03584L32   0000013FA0520411 lpfc                11:0:1:1            NA
1       03584L32   0000013FA0520411 lpfc                11:0:2:1            NA
lin_tape version: 3.0.13
lin_tape major number: 246
Attached Tape Devices:
Number  model      SN                HBA                SCSI                FO Path
0       ULT3580-TD5 1068093078        lpfc                11:0:0:0            NA
1       ULT3580-TD5 1068093069        lpfc                11:0:1:0            NA
2       ULT3580-TD5 1068092972        lpfc                11:0:2:0            NA
3       ULT3580-TD5 1068093084        lpfc                11:0:3:0            NA
[root@ltfsm11 ~]# ps -afe | grep dsm
root    14351    1  0 15:33 ?        00:00:01
/opt/tivoli/tsm/client/hsm/bin/dsmwatchd nodetach
root    15131 30301  0 16:33 pts/0    00:00:00 grep --color=auto dsm
root    17135    1  0 15:33 ?        00:00:00 dsmrecalld
root    17160 17135  0 15:33 ?        00:00:00 dsmrecalld
root    17161 17135  0 15:33 ?        00:00:00 dsmrecalld

[root@ltfsm11 ~]# ltfsee start
Library name: lib_ebisu, library id: 0000013FA0520411, control node (ltfsee_md) IP
address: 9.11.120.198.
GLES648I(00277): Running start command - sending request : lib_ebisu
GLES648I(00297): Running start command - waiting for completion : lib_ebisu
.....
GLES639I(00136): Started the Spectrum Archive EE services for library lib_ebisu
with good status.
[root@ltfsm11 ~]# pidof ltf
32482
[root@ltfsm11 ~]# pidof mmm
823
[root@ltfsm12 ~]# pidof ltf
13872
```

---

Now, IBM Spectrum Archive EE is started and ready for basic usage. For further handling, managing, and operations of IBM Spectrum Archive EE (such as creating pools, adding and formatting tapes, and setting up migration policies), see Chapter 7, “Operations” on page 129.

## 6.4 Configuring IBM Spectrum Archive EE with IBM Spectrum Scale AFM

This section walks through how to set up IBM Spectrum Archive EE and IBM Spectrum Scale AFM to create either a Centralized Archive Repository, or an Asynchronous Archive Replication solution. The steps shown in this section assume that the user has already installed and configured IBM Spectrum Archive EE.

If IBM Spectrum Archive EE has not been previously installed and configured, set up AFM first and then follow the instructions in Chapter 4, “Installation” on page 49 to install Spectrum Archive EE and then in Chapter 6, “Configuration” on page 95. If performed in that order, you can skip this section. See 8.10.3, “IBM Spectrum Archive EE migration policy with AFM” on page 238 for information about creating migration policy on cache nodes.

**Important:** Starting with IBM Spectrum Archive EE V1.2.3.0, IBM Spectrum Scale AFM is supported. This support is limited to only one cache mode, independent writer (IW).

For a more detailed explanation on configuring IBM Spectrum Scale AFM, see the following documentation:

[https://www.ibm.com/support/knowledgecenter/STXKQY\\_4.2.0/com.ibm.spectrum.scale.v4r2.adv.doc/b11adv\\_afmsystemsetup.htm](https://www.ibm.com/support/knowledgecenter/STXKQY_4.2.0/com.ibm.spectrum.scale.v4r2.adv.doc/b11adv_afmsystemsetup.htm)

### 6.4.1 Configuring a Centralized Archive Repository solution

A Centralized Archive Repository solution consists of having IBM Spectrum Archive EE at just the home cluster of IBM Spectrum Scale AFM. The steps in this section show how to set up a home site with IBM Spectrum Archive EE, and how to set up the cache site and link them. For more information on use cases, see Figure 9-5 on page 270.

Steps 1 - 5 demonstrate how to set up a IBM Spectrum Scale AFM home cluster and start IBM Spectrum Archive EE. Steps 6 - 9 show how to set up the IW caches for IBM Spectrum Scale AFM cache clusters:

1. If IBM Spectrum Scale is not already active and GPFS is not already mounted, start IBM Spectrum Scale and wait until the cluster becomes active. Then, mount the file system if it is not set to mount automatically using the commands in Example 6-8.

*Example 6-8 Starting and mounting IBM Spectrum Scale and GPFS file system*

```
[root@l1fseehomesrv ~]# mmstartup -a
Tue Mar 21 14:37:57 MST 2017: mmstartup: Starting GPFS ...
[root@l1fseehomesrv ~]# mmgetstate -a

Node number  Node name      GPFS state
-----
           1      l1fseehomesrv  arbitrating
[root@l1fseehomesrv ~]# mmgetstate -a

Node number  Node name      GPFS state
-----
           1      l1fseehomesrv  active
[root@l1fseehomesrv ~]# mmmount all -a
Tue Mar 21 14:40:36 MST 2017: mmmount: Mounting file systems ...
[root@l1fseeccachesrv ~]# systemctl start hsm
```

```
[root@ltfseecachesrv ~]# dsmmigfs start
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 0.0
  Client date/time: 03/22/2017 13:41:36
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.
```

```
[root@ltfseecachesrv ~]# dsmmigfs enablefailover
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 0.0
  Client date/time: 03/22/2017 13:41:41
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.
```

Automatic failover is enabled on this node in mode ENABLED.

---

**Note:** Step 2 assumes that the user has already created their file set and linked it to the GPFS file system. The following examples use IWhome as the home file set.

2. After IBM Spectrum Scale is active and the GPFS file system is mounted, edit the NFS exports file (/etc/exports) to include the new file set. It is important that the no\_root\_squash, sync, and rw arguments are used. Example 6-9 shows example content of the exports file for file set IWhome.

*Example 6-9 Contents of an exports file*

---

```
[root@ltfseehomesrv ~]# cat /etc/exports
/ibm/glues/IWhome
*(rw,sync,no_root_squash,nohide,insecure,no_subtree_check,fsid=125)
```

---

**Note:** The fsid in the exports file needs to be a unique number different than any other export clause within the exports file.

3. After the exports file has been modified to include the file set, start the NFS service. Example 6-10 shows an example of starting and checking the NFS service.

*Example 6-10 Starting and checking the status of NFS service*

---

```
[root@ltfseehomesrv ~]# systemctl start nfs
[root@ltfseehomesrv ~]# systemctl status nfs
? nfs-server.service - NFS server and services
   Loaded: loaded (/usr/lib/systemd/system/nfs-server.service; enabled; vendor
   preset: disabled)
   Active: active (exited) since Tue 2017-03-21 15:58:43 MST; 2s ago
     Process: 1895 ExecStopPost=/usr/sbin/exportfs -f (code=exited,
   status=0/SUCCESS)
     Process: 1891 ExecStopPost=/usr/sbin/exportfs -au (code=exited,
   status=0/SUCCESS)
     Process: 1889 ExecStop=/usr/sbin/rpc.nfsd 0 (code=exited, status=0/SUCCESS)
     Process: 10062 ExecStart=/usr/sbin/rpc.nfsd $RPCNFSDARGS (code=exited,
   status=0/SUCCESS)
     Process: 10059 ExecStartPre=/usr/sbin/exportfs -r (code=exited,
   status=0/SUCCESS)
   Main PID: 10062 (code=exited, status=0/SUCCESS)
     CGroup: /system.slice/nfs-server.service
```

```
Mar 21 15:58:43 ltfseehomesrv.tuc.stglabs.ibm.com systemd[1]: Starting NFS
server and services...
Mar 21 15:58:43 ltfseehomesrv.tuc.stglabs.ibm.com systemd[1]: Started NFS
server and services.
```

---

4. After NFS has properly started the final step to configure IBM Spectrum Scale AFM at the home cluster is to enable the exported path. Run **mmafmconfig enable <path-to-fileset>** to enable the exported file set. Example 6-11 shows the execution of the **mmafmconfig** command with the IWhome file set.

*Example 6-11 Execution of mmafconfig enable <path-to-fileset>*

---

```
[root@ltfseehomesrv ~]# mmafconfig enable /ibm/glues/IWhome/
[root@ltfseehomesrv ~]#
```

---

5. After the file set has been enabled for AFM proceed by starting up IBM Spectrum Archive if it has not been started previously by running **ltfsee start**.

Run the next four steps on the designated cache nodes:

6. Before starting IBM Spectrum Scale on the cache clusters, determine which nodes will become the gateway nodes and then run the **mmchnode --gateway -N <node1,node2,etc...>** command to create gateway nodes. Example 6-12 shows the output of running **mmchnode** on one cache node.

*Example 6-12 Setting gateway nodes for cache clusters*

---

```
[root@ltfseecachesrv ~]# mmchnode --gateway -N ltfseecachesrv
Tue Mar 21 16:49:16 MST 2017: mmchnode: Processing node
ltfseecachesrv.tuc.stglabs.ibm.com
[root@ltfseecachesrv ~]#
```

---

7. After all the gateway nodes have been set, start IBM Spectrum Scale and mount the file system if it is not done automatically, as shown in Example 6-13:

- a. **mmmstartup -a**
- b. **mmgetstate -a**
- c. **mmm mount all -a** (optional, only if the GPFS file system is not mounted automatically)

*Example 6-13 Starting and mounting IBM Spectrum Scale and GPFS file system*

---

```
[root@ltfseecachesrv ~]# mmstartup -a
Tue Mar 21 14:37:57 MST 2017: mmstartup: Starting GPFS ...
[root@ltfseecachesrv ~]# mmgetstate -a
```

```
Node number  Node name          GPFS state
-----
          1      ltfseecachesrv  arbitrating
[root@ltfseecachesrv ~]# mmgetstate -a
```

```
Node number  Node name          GPFS state
-----
          1      ltfseecachsv1  active
[root@ltfseecachesrv ~]# mmmount all -a
Tue Mar 21 14:40:36 MST 2017: mmmount: Mounting file systems ...
[root@ltfseecachesrv ~]#
```

---

- After IBM Spectrum Scale has been started and the GPFS file system is mounted, then create the cache fileset by using `mmcrfileset` with the `afmTarget`, `afmMode`, and `inode-space` parameters. Example 6-14 shows the execution of `mmcrfileset` to create a cache fileset.

*Example 6-14 Creating a cache fileset that targets the home fileset*

---

```
[root@l1tfseeecachesrv ~]# mmcrfileset gpfs iwcache -p afmmode=iw -p
afmtarget=l1tfseeehomesrv:/ibm/glues/IWhome --inode-space=new
Fileset iwcache created with id 1 root inode 4194307.
```

---

- After the fileset is created, it can be linked to a directory in the GPFS file system by running the `mmlinkfileset <device> <fileset> -J <gpfs file system/fileset name>` command. Example 6-15 shows output of running `mmlinkfileset`.

*Example 6-15 Linking the GPFS fileset to a directory on the GPFS file system*

---

```
[root@l1tfseeecachesrv glues]# mmlinkfileset gpfs iwcache -J /ibm/glues/iwcache
Fileset iwcache linked at /ibm/glues/iwcache
```

---

Steps 6 - 9 need to be run on each cache cluster that will be linked to the home cluster. After completing these steps, IBM Spectrum Scale AFM and IBM Spectrum Archive EE are set up on the home cluster and IBM Spectrum Scale AFM is set up on each cache cluster. The system is ready to perform centralized archiving and caching.

## 6.4.2 Configuring an Asynchronous Archive Replication solution

An Asynchronous Archive Replication solution consists of having IBM Spectrum Archive EE at both the home and cache cluster for IBM Spectrum Scale AFM. This section demonstrates how to set up IBM Spectrum Scale AFM with IBM Spectrum Archive EE to create an Asynchronous Archive Replication solution. For more information on use cases, see 9.2.2, “Asynchronous Archive Replication” on page 270.

Steps 1 - 5 demonstrate how to set up a IBM Spectrum Scale AFM home cluster and start IBM Spectrum Archive EE. Steps 6 - 11 demonstrate how to set up the cache clusters, and steps 12 - 15 demonstrate how to reconfigure Spectrum Archive EE's configuration to work with IBM Spectrum Scale AFM.

- If IBM Spectrum Archive is not already active and GPFS is not already mounted, start the file system and wait until the file system becomes active. Then, mount the file system if it is not set to mount automatically using the commands in Example 6-16.

*Example 6-16 Starting and mounting IBM Spectrum Scale and GPFS file system*

---

```
[root@l1tfseeehomesrv ~]# mmstartup -a
Tue Mar 21 14:37:57 MST 2017: mmstartup: Starting GPFS ...
[root@l1tfseeehomesrv ~]# mmgetstate -a
```

```
Node number  Node name      GPFS state
-----
          1      l1tfseeehomesrv  arbitrating
[root@l1tfseeehomesrv ~]# mmgetstate -a
```

```
Node number  Node name      GPFS state
-----
          1      l1tfseeehomesrv  active
[root@l1tfseeehomesrv ~]# mmmount all -a
```

```

Tue Mar 21 14:40:36 MST 2017: mmmount: Mounting file systems ...
[root@ltfseehomesrv ~]# systemctl start hsm
[root@ltfseehomesrv ~]# dsmmigfs start
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 0.0
  Client date/time: 03/22/2017 13:41:36
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

[root@ltfseehomesrv ~]# dsmmigfs enablefailover
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 0.0
  Client date/time: 03/22/2017 13:41:41
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

Automatic failover is enabled on this node in mode  ENABLED.

```

---

**Note:** Step 2 assumes that the user has already created their fileset and linked it to the GPFS file system. The following examples use IWhome as the home fileset.

2. After IBM Spectrum Scale is active and the GPFS file system is mounted, edit the NFS exports file (/etc/exports) to include the new fileset. It is important that the no\_root\_squash, sync, and rw arguments are used. Example 6-17 shows example content of the exports file for fileset IWhome.

*Example 6-17 Contents of an exports file*

---

```

[root@ltfseehomesrv ~]# cat /etc/exports
/ibm/glues/IWhome
*(rw,sync,no_root_squash,nohide,insecure,no_subtree_check,fsid=125)

```

---

**Note:** The fsid in the exports file needs to be a unique number different than any other export clause within the exports file.

3. After the exports file has been modified to include the fileset, start the NFS service. Example 6-18 shows an example of starting and checking the NFS service.

*Example 6-18 Starting and checking the status of NFS service*

---

```

[root@ltfseehomesrv ~]# systemctl start nfs
[root@ltfseehomesrv ~]# systemctl status nfs
? nfs-server.service - NFS server and services
   Loaded: loaded (/usr/lib/systemd/system/nfs-server.service; enabled; vendor
   preset: disabled)
   Active: active (exited) since Tue 2017-03-21 15:58:43 MST; 2s ago
   Process: 1895 ExecStopPost=/usr/sbin/exportfs -f (code=exited,
   status=0/SUCCESS)
   Process: 1891 ExecStopPost=/usr/sbin/exportfs -au (code=exited,
   status=0/SUCCESS)
   Process: 1889 ExecStop=/usr/sbin/rpc.nfsd 0 (code=exited, status=0/SUCCESS)
   Process: 10062 ExecStart=/usr/sbin/rpc.nfsd $RPCNFSDARGS (code=exited,
   status=0/SUCCESS)
   Process: 10059 ExecStartPre=/usr/sbin/exportfs -r (code=exited,
   status=0/SUCCESS)

```

```
Main PID: 10062 (code=exited, status=0/SUCCESS)
CGroup: /system.slice/nfs-server.service
```

```
Mar 21 15:58:43 ltfseehomesrv.tuc.stglabs.ibm.com systemd[1]: Starting NFS
server and services...
```

```
Mar 21 15:58:43 ltfseehomesrv.tuc.stglabs.ibm.com systemd[1]: Started NFS
server and services.
```

---

4. After NFS has properly started, the final step to configure IBM Spectrum Scale AFM at the home cluster is to enable the exported path. Run **mmafmconfig enable <path-to-fileset>** to enable the exported fileset. Example 6-19 shows the execution of the **mmafmconfig** command with the IWhome fileset.

*Example 6-19 Execution of mmafmconfig enable <path-to-fileset>*

---

```
[root@ltfseehomesrv ~]# mmafmconfig enable /ibm/glues/IWhome/
[root@ltfseehomesrv ~]#
```

---

5. After the fileset has been enabled for AFM, start Spectrum Archive if it has not been started previously by running **ltfsee start**.

After the home cluster is set up and an NFS export directory is enabled for IBM Spectrum Scale AFM, steps 6 - 11 demonstrate how to set up a Spectrum Scale AFM IW cache fileset at a cache cluster and connect the cache's fileset with the home's fileset. Steps 12 - 15 show how to modify IBM Spectrum Archive EE's configuration to allow cache filesets.

6. If IBM Spectrum Archive EE is active, properly shut it down by using the commands in Example 6-20.

*Example 6-20 Shutting down IBM Spectrum Archive EE*

---

```
[root@ltfseecachesrv ~]# ltfsee stop
Library name: AnacondaQ, library id: 0000013400190402, control node (MMM) IP
address: 9.11.121.209.
Stopped LTFS EE service (MMM) for library AnacondaQ.
[root@ltfseecachesrv ~]# pidof mmm
[root@ltfseecachesrv ~]# umount /ltfs
[root@ltfseecachesrv ~]# pidof ltfs
[root@ltfseecachesrv ~]#
```

---

7. If IBM Spectrum Scale is active, properly shut it down by using the commands in Example 6-21.

*Example 6-21 Shutting down IBM Spectrum Scale*

---

```
[root@ltfseecachesrv ~]# dsmmigfs disablefailover
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 0.0
  Client date/time: 03/22/2017 13:31:14
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.
```

Automatic failover is disabled on this node.

```
[root@ltfseecachesrv ~]# dsmmigfs stop
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 0.0
  Client date/time: 03/22/2017 13:31:19
```

(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

```
[root@ltfseecachesrv ~]# systemctl stop hsm
[root@ltfseecachesrv ~]# mmumount all -a
Wed Mar 22 13:31:44 MST 2017: mmumount: Unmounting file systems ...
[root@ltfseecachesrv ~]# mmshutdown -a
Wed Mar 22 13:31:56 MST 2017: mmshutdown: Starting force unmount of GPFS file
systems
Wed Mar 22 13:32:01 MST 2017: mmshutdown: Shutting down GPFS daemons
ltfseecachesrv.tuc.stglabs.ibm.com: Shutting down!
ltfseecachesrv.tuc.stglabs.ibm.com: 'shutdown' command about to kill process
24101
ltfseecachesrv.tuc.stglabs.ibm.com: Unloading modules from
/lib/modules/3.10.0-229.el7.x86_64/extra
ltfseecachesrv.tuc.stglabs.ibm.com: Unloading module mmfs26
ltfseecachesrv.tuc.stglabs.ibm.com: Unloading module mmfslinux
Wed Mar 22 13:32:10 MST 2017: mmshutdown: Finished
```

---

8. With IBM Spectrum Archive EE and IBM Spectrum Scale both shut down, set the gateway nodes if they have not been set when IBM Spectrum Scale was configured by using the command in Example 6-22.

*Example 6-22 Setting a gateway node*

---

```
[root@ltfseecachesrv ~]# mmchnode --gateway -N ltfseecachesrv
Tue Mar 21 16:49:16 MST 2017: mmchnode: Processing node
ltfseecachesrv.tuc.stglabs.ibm.com
[root@ltfseecachesrv ~]#
```

---

9. Properly start IBM Spectrum Scale by using the commands in Example 6-23.

*Example 6-23 Starting IBM Spectrum Scale*

---

```
[root@ltfseecachesrv ~]# mmstartup -a
Wed Mar 22 13:41:02 MST 2017: mmstartup: Starting GPFS ...
[root@ltfseecachesrv ~]# mmmount all -a
Wed Mar 22 13:41:22 MST 2017: mmmount: Mounting file systems ...
[root@ltfseecachesrv ~]# systemctl start hsm
[root@ltfseecachesrv ~]# dsmmigfs start
IBM Spectrum Protect
Command Line Space Management Client Interface
Client Version 8, Release 1, Level 0.0
Client date/time: 03/22/2017 13:41:36
```

(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

```
[root@ltfseecachesrv ~]# dsmmigfs enablefailover
IBM Spectrum Protect
Command Line Space Management Client Interface
Client Version 8, Release 1, Level 0.0
Client date/time: 03/22/2017 13:41:41
```

(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

Automatic failover is enabled on this node in mode ENABLED.

---

10. Create the independent-writer fileset by using the command in Example 6-24.

*Example 6-24 Creating an IW fileset*

---

```
[root@ltfseecachesrv ~]# mmcrfileset gpfs iwcache -p afmnode=independent-writer
-p afmtarget=ltfseehomesrv:/ibm/glues/IWhome --inode-space=new
Fileset iwcache created with id 1 root inode 4194307.
[root@ltfseecachesrv ~]#
```

---

11. Link the fileset to a directory on the node's GPFS file system by using the command in Example 6-25.

*Example 6-25 Linking an IW fileset*

---

```
[root@ltfseecachesrv ~]# mmlinkfileset gpfs iwcache -J /ibm/glues/iwcache
Fileset iwcache linked at /ibm/glues/iwcache
[root@ltfseecachesrv ~]#
```

---

IBM Spectrum Scale AFM is now configured and has a working home and IW cache clusters.

12. With IBM Spectrum Archive EE still shut down, obtain the metadata and HSM file systems IBM Spectrum Archive EE by using the command in Example 6-26.

*Example 6-26 Obtaining metadata and HSM file system(s)*

---

```
[root@ltfseecachesrv ~]# ltfsee_config -m INFO
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
INFO
INFO mode starts .
```

```
## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file
ltfsee_config.filesystem exists.
Metadata Filesystem:
  /ibm/glues
HSM Filesystems:
  /ibm/glues
Library: Name=ltfsee_lib1, S/N=0000013400190402
Node Group: Name=ltfseecachesrv
Node: ltfseecachesrv.tuc.stglabs.ibm.com
  Drive: S/N=00078D00BC, Attribute='mrg'
  Drive: S/N=00078D00BD, Attribute='mrg'
Pool: Name=copy_cache, ID=902b097a-7a34-4847-a346-0e6d97444a21
  Tape: Barcode=DV1982L7
  Tape: Barcode=DV1985L7
Pool: Name=primary_cache, ID=14adb6cf-d1f5-46ef-a0bb-7b3881bdb4ec
  Tape: Barcode=DV1983L7
  Tape: Barcode=DV1984L7
```

---

13. Modify IBM Spectrum Archive EE's configuration by using the command in Example 6-27 with the same file systems recorded from step 7.

*Example 6-27 Modify IBM Spectrum Archive EE configuration for IBM Spectrum Scale AFM*

---

```
[root@ltfseecachesrv ~]# ltfsee_config -m CLUSTER
The EE configuration script is starting: /opt/ibm/ltfsee/bin/ltfsee_config -m
CLUSTER
CLUSTER mode starts .
```

---

```

## 1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file
lrfsee_config.filesystem exists.

## 2. Check prerequisite on cluster ##
Cluster name: lrfseecachesrv.tuc.stglabs.ibm.com
ID: 14631719600452284991
Successfully validated the prerequisites.

## 3. List file systems in the cluster ##
Retrieving IBM Spectrum Scale (GPFS) file systems...
** Select a file system for storing IBM Spectrum Archive Enterprise Edition
configuration and internal data.
    Input the corresponding number and press Enter
    or press 'q' followed by Enter to quit.

    File system
    1. /dev/gpfs Mount point(/ibm/glues) DMAPI(Yes)
    q. Quit

Input number > 1

** Select file systems to configure for IBM Spectrum Scale (GPFS) file system
for Space Management.
    Input the corresponding numbers and press Enter
    or press 'q' followed by Enter to quit.
    Press a followed by Enter to select all file systems.
    Multiple file systems can be specified using comma or white space
    delimiters.

    File system
    1. /dev/gpfs Mount point(/ibm/glues)
    a. Select all file systems
    q. Quit

Input number > 1

## 4. Configure space management ##
Disabling unnecessary daemons...
Editing Space Management Client settings...
Deactivating failover operations on the node.
Restarting Space Management service...
Stopping the HSM service.
Terminating dsmwatchd.....
Starting the HSM service.
Starting dsmmigfs.....
Activating failover operations on the node.
Configured space management.

## 5. Add selected file systems to Space Management ##
Added the selected file systems to the space management.

## 6. Store the file systems configuration and dispatch it to all nodes ##
Storing the file systems configuration.

```

Copying ltfsee\_config.filesystem file.  
Stored the cluster configuration and dispatched the configuration file.

## 7. Create metadata directories and the configuration parameters file. ##  
Created metadata directories and the configuration parameters file.  
Enabling runtime AFM file state checking.

CLUSTER mode completed .

---

14. Start IBM Spectrum Archive EE by using the commands in Example 6-28.

*Example 6-28 Start IBM Spectrum Archive EE*

---

```
[root@ltfseecachesrv ~]# ltfsee start
Library name: AnacondaQ, library ID: 0000013400190402, control node (ltfsee_md)
IP address: 9.11.121.209.
Running start command - sending request : AnacondaQ.
Running start command - waiting for completion : AnacondaQ.
.....
Started the Spectrum Archive Enterprise Edition services for library AnacondaQ
with good
status.
```

---

15. At the start of IBM Spectrum Archive EE, the AFMSKIPUNCACHEDFILES flag inside the /opt/tivoli/tsm/client/ba/bin/dsm.sys file should be set to yes. It can be checked by using the command in Example 6-29. If it has not been properly set, modify the file so that the AFMSKIPUNCACHEDFILES is set to yes.

*Example 6-29 Validating AFMSKIPUNCACHEDFILES is set to yes*

---

```
[root@ltfseecachsrv ~]# grep AFMSKIPUNCACHEDFILES
/opt/tivoli/tsm/client/ba/bin/dsm.sys
AFMSKIPUNCACHEDFILES YES
```

---

After successfully completing these steps, IBM Spectrum Archive EE and IBM Spectrum Scale AFM are set up at both the home and cache cluster. They can now be used as an Asynchronous Archive Replication solution.





# Operations

In this chapter, the day-to-day management of the IBM Spectrum Archive Enterprise Edition (IBM Spectrum Archive EE) environment is described.

This chapter includes the following topics:

- ▶ Overview
- ▶ Status information
- ▶ Upgrading components
- ▶ Starting and stopping IBM Spectrum Archive EE
- ▶ IBM Spectrum Archive EE database backup
- ▶ IBM Spectrum Archive EE automatic node failover
- ▶ Tape library management
- ▶ Tape storage pool management
- ▶ Pool capacity monitoring
- ▶ Migration
- ▶ Premigration
- ▶ Preserving file system objects on tape
- ▶ Restoring non-empty regular files and file system objects from tape
- ▶ Recall
- ▶ Repairing files to their resident state
- ▶ Handling of write-failure tapes
- ▶ Handling of read-failure tapes
- ▶ Reconciliation
- ▶ Reclamation
- ▶ Checking and repairing
- ▶ Importing and exporting
- ▶ Drive Role settings for job assignment control
- ▶ Tape drive intermix support
- ▶ WORM support for the IBM TS1155, TS1150, and TS1140 tape drives
- ▶ Obtaining the location of files and data
- ▶ Obtaining inventory, job, and scan status
- ▶ Cleaning up a scan or session
- ▶ Monitoring the system with SNMP
- ▶ Configuring Net-SNMP

- ▶ IBM Spectrum Archive REST API
- ▶ Data Migration

**Note:** In our lab setup for writing this book, we used a Red Hat-based Linux system. The screen captures within this chapter are based on the Version 1 Release 2 of the product. Although the steps that you will perform are the same, you might see slightly different output responses in your window depending on your version and release of the product.

## 7.1 Overview

The following terms specific to IBM Spectrum Archive EE operations are used in this chapter:

<b>Migration</b>	The movement of files from the IBM Spectrum Scale file system on disk to IBM Linear Tape File System tape cartridges, which leaves behind a stub file.
<b>Premigration</b>	The movement of files from GPFS file systems on disk to LTFS tape cartridges without replacing them with stub files on the GPFS file system. Identical copies of the files are on the GPFS file system and in LTFS storage.
<b>Recall</b>	The movement of migrated files from tape cartridges back to the originating GPFS file system on disk, which is the reverse of migration.
<b>Reconciliation</b>	The process of synchronizing a GPFS file system with the contents of an LTFS tape cartridge and removing old and obsolete objects from the tape cartridge. You must run reconciliation when a GPFS file is deleted, moved, or renamed.
<b>Reclamation</b>	The process of defragmenting a tape cartridge. The space on a tape cartridge that is occupied by deleted files is not reused during normal LTFS operations. New data is always written after the last index on tape. The process of reclamation is similar to the same named process in IBM Tivoli Storage Manager from the IBM Spectrum Archive family. All active files are consolidated onto a second tape cartridge, which improves overall tape usage.
<b>Retrieval</b>	The process of triggering IBM Spectrum Archive EE to retrieve information about physical resources from the tape library. This process is scheduled to occur automatically at regular intervals, but can be run manually.
<b>Tape check</b>	The process of checking a tape for errors when it is added to a tape cartridge pool.
<b>Tape repair</b>	The process of repairing any errors that are found on the tape during the tape check process.
<b>Import</b>	The addition of an LTFS tape cartridge to IBM Spectrum Archive EE.
<b>Export</b>	The removal of an LTFS tape cartridge from IBM Spectrum Archive EE.
<b>Data migration</b>	The new method of technology migration of tape drives and cartridges. There are two ways to perform data migration: <ul style="list-style-type: none"><li>- within a pool</li><li>- pool to pool</li></ul>

### 7.1.1 IBM Spectrum Archive EE command summaries

Use IBM Spectrum Archive EE commands to configure IBM Spectrum Archive EE tape cartridge pools and perform IBM Spectrum Archive EE administrative tasks. The commands use the `ltfsee <options>` format.

The following `ltfsee` command options are available. All options, except `info`, can be run only with root user permissions:

- `ltfsee cleanup`

Use this command to clean up scan results that are not removed after a scan or session finishes. For example, you can use this command to clean up any scan numbers that remain after a user runs a **ltfsee migrate** command, then closes the command window before the command completes.

▶ **ltfsee datamigrate**

Use this command to move data by running reclaim from one pool to a different pool, but within the same node group. This function can be used to move data on older technology tapes to newer technology tapes in a different pool.

▶ **ltfsee drive**

Use this command to add a tape drive to or remove a tape drive from the IBM Spectrum Archive EE system.

▶ **ltfsee export**

Use this command to remove one or more tape cartridges from the IBM Spectrum Archive EE system by removing files on them from the IBM Spectrum Scale namespace. After you use this command, run the **ltfsee tape move ieslot** command to move a tape cartridge to the I/O station in the library. The **ltfsee export** command has two export modes: Normal Export and Offline Export. Normal Export removes the metadata on the GPFS file system and removes the tape from the storage pool. Offline Export keeps the metadata on the file system and keeps the tape in the storage pool as Offline.

▶ **ltfsee failover**

Use this command to manually cause an active control node to perform a failover to the redundant control node.

▶ **ltfsee fsopt**

Use this command to query or update the file system level settings stub size, read start recalls, and preview size.

▶ **ltfsee import**

Use this command to add one or more tape cartridges to the IBM Spectrum Archive EE system and reinstantiate the files in the IBM Spectrum Scale namespace.

▶ **ltfsee info**

Use this command to list current information about LTFs EE jobs and scans, and resource inventory for tape cartridges, drives, nodes, tape cartridge pools, and files.

▶ **ltfsee library**

Use this command to list the settings of each library.

▶ **ltfsee migrate**

Use this command to migrate files to tape cartridge pools.

▶ **ltfsee node**

Use this command to enable or disable nodes and list the settings of each node.

▶ **ltfsee pool**

Use this command to create, delete, and modify IBM Spectrum Archive EE tape cartridge pools, and to format tape cartridges before they are added to the tape cartridge pools.

▶ **ltfsee premigrate**

Use this command to premigrate files to tape cartridge pools.

▶ **ltfsee rebuild**

This command rebuilds a GPFS file system into the specified directory with the files and file system objects that are found on the specified tapes.

▶ **ltfsee recall**

Use this command to perform selective bulk recalls from the tape cartridges.

▶ **ltfsee reclaim**

Use this command to free tape space by removing unreferenced files and unreferenced content from tape cartridges in a tape cartridge pool. The **ltfsee reconcile** command is a prerequisite for efficient tape reclamation.

▶ **ltfsee reconcile**

Use this command to perform reconciliation tasks between the IBM Spectrum Scale namespace and the LTFSS namespace. This command is a prerequisite for efficient tape reclamation.

▶ **ltfsee recover**

Use this command to recover files from tape or to remove a tape from IBM Spectrum Archive EE when the tape is in the Critical state or the Write Fenced state.

▶ **ltfsee repair**

Use this command to repair a file or file system object by changing its state to Resident.

▶ **ltfsee retrieve**

Use this command to synchronize data in the IBM Spectrum Archive EE inventory with the tape library. The IBM Spectrum Archive EE system synchronizes data automatically, but it might be useful to explicitly trigger this operation if the configuration changes, for example, if a drive is added or removed.

▶ **ltfsee save**

Use this command to save file system objects (symbolic links, empty files, and empty directories) to tape cartridge pools.

▶ **ltfsee start**

Use this command to start the IBM Spectrum Archive EE and LE+ system. HSM components must be running before you can use this command. You can run the **ltfsee start** command on any IBM Spectrum Archive EE node in the cluster.

**Important:** If the **ltfsee start** command does not return after several minutes, it might be because the firewall is running or tapes are being unmounted from the drives. The firewall service must be disabled on the IBM Spectrum Archive EE nodes. For more information, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60.

▶ **ltfsee status**

Use this command to identify the node where the multi-tape management module (MMM) service is started.

▶ **ltfsee stop**

Use this command to stop the IBM Spectrum Archive EE system. You can run this command on any IBM Spectrum Archive EE node in the cluster.

▶ **ltfsee tape**

Use this command to move one or more tape cartridges to an IE slot in the I/O station or to a home slot.

► **ltfsee threshold**

Use this command to specify the limit at which migrations are preferred over recalls.

## 7.1.2 Using the command-line interface

The IBM Spectrum Archive EE system provides a command-line interface (CLI) that supports the automation of administrative tasks, such as starting and stopping the system, monitoring its status, and configuring tape cartridge pools. The CLI is the primary method for administrators to manage IBM Spectrum Archive EE. There is no GUI available as of this writing.

In addition, the CLI is used by the IBM Spectrum Scale `mmapplypolicy` command to trigger migrations or premigrations. When this action occurs, the `mmapplypolicy` command calls IBM Spectrum Archive EE when an IBM Spectrum Scale scan occurs, and passes the file name of the file that contains the scan results and the name of the target tape cartridge pool.

The `ltfsee` command uses the following syntax:

```
ltfsee <options>
```

**Reminder:** All of the command examples use the command without the full file path name because we added the IBM Spectrum Archive EE directory (`/opt/ibm/ltfsee/bin`) to the `PATH` variable.

For more information, see 11.1, “Command-line reference” on page 300.

## 7.2 Status information

This section describes the process that is used to determine whether each of the major components of IBM Spectrum Archive EE is running correctly. For more information about troubleshooting IBM Spectrum Archive EE, see Chapter 10, “Troubleshooting” on page 273.

The components should be checked in the order that is shown here because a stable, active GPFS file system is a prerequisite for starting IBM Spectrum Archive EE.

### 7.2.1 IBM Spectrum Scale

The following IBM Spectrum Scale commands are used to obtain cluster state information:

- The `mmdiag` command obtains basic information about the state of the GPFS daemon.
- The `mmgetstate` command obtains the state of the GPFS daemon on one or more nodes.
- The `mmiscluster` and `mmisconfig` commands show detailed information about the GPFS cluster configuration.

This section describes how to obtain GPFS daemon state information by running the GPFS command `mmgetstate`. For more information about the other GPFS commands, see the *General Parallel File System Version 4 Release 1.0.4 Advanced Administration Guide*, SC23-7032-01, or see the *IBM Spectrum Scale: Administration Guide*, which is available at this website:

<https://www.ibm.com/support/knowledgecenter/STXKQY>

The node on which the `mmgetstate` command is run must have the GPFS mounted. The node must also run remote shell commands on any other node in the GPFS/IBM Spectrum Scale cluster without the use of a password and without producing any extraneous messages.

Example 7-1 shows how to get status about the GPFS/IBM Spectrum Scale daemon on one or more nodes.

*Example 7-1 Check the GPFS/IBM Spectrum Scale status*

---

```
[root@lufs97 ~]# mmgetstate -a
Node number Node name      GPFS state
-----
          1    htohr9      down
```

---

The `-a` argument shows the state of the GPFS/IBM Spectrum Scale daemon on all nodes in the cluster.

**Permissions:** Retrieving the status for GPFS/IBM Spectrum Scale requires root user permissions.

The following GPFS/IBM Spectrum Scale states are recognized and shown by this command:

- ▶ Active: GPFS/IBM Spectrum Scale is ready for operations.
- ▶ Arbitrating: A node is trying to form a quorum with the other available nodes.
- ▶ Down: GPFS/IBM Spectrum Scale daemon is not running on the node or is recovering from an internal error.
- ▶ Unknown: Unknown value. The node cannot be reached or some other error occurred.

If the GPFS/IBM Spectrum Scale state is not active, attempt to start GPFS/IBM Spectrum Scale and check its status, as shown in Example 7-2.

*Example 7-2 Start GPFS/IBM Spectrum Scale*

---

```
[root@lufs97 ~]# mmstartup -a
Tue Apr  2 14:41:13 JST 2013: mmstartup: Starting GPFS ...
[root@lufs97 ~]# mmgetstate -a
Node number Node name      GPFS state
-----
          1    htohr9      active
```

---

If the status is active, also check the GPFS/IBM Spectrum Scale mount status by running the command that is shown in Example 7-3.

*Example 7-3 Check the GPFS/IBM Spectrum Scale mount status*

---

```
[root@lufs97 ~]# mmlsmount all
File system gpfs is mounted on 1 nodes.
```

---

The message confirms that the GPFS file system is mounted.

## 7.2.2 IBM Spectrum Archive Library Edition Plus component

IBM Spectrum Archive EE constantly checks to see whether the IBM Spectrum Archive Library Edition Plus (LE+) component is running. If the IBM Spectrum Archive LE+ component is running correctly, you can see whether the LTFS file system is mounted by running the **mount** command or the **df** command, as shown in Example 7-4. The IBM Spectrum Archive LE+ component must be running on all EE nodes.

*Example 7-4 Check the IBM Spectrum Archive LE+ component status (running)*

---

```
[root@ltfs97 ~]# df -m
Filesystem            1M-blocks      Used Available Use% Mounted on
/dev/mapper/Vo1Group-lv_root
                        33805          5081      27007    16% /
tmpfs                  1963            0         1963     0% /dev/shm
/dev/vda1              485             36         424     8% /boot
/dev/gpfs              153600          8116     145484    6% /ibm/glues
ltfs:/dev/IBMchanger8
                        2147483648      0 2147483648    0% /ltfs
```

---

With IBM Spectrum Archive EE V1.2.4.0 and later, IBM Spectrum Archive LE+ is started automatically when running the **ltfsee start** command. If errors occur during the start of the IBM Spectrum Archive EE system, run the **ltfsee info nodes** command to display which component failed to start. For more information about the updated **ltfsee info nodes** command, see 7.6, “IBM Spectrum Archive EE automatic node failover” on page 145.

## 7.2.3 Hierarchical Space Management

Hierarchical Space Management (HSM) must be running before you start IBM Spectrum Archive EE. You can verify that HSM is running by checking whether the watch daemon (**dsmwatchd**) and at least three recall daemons (**dsmrecalld**) are active. Query the operating system to verify that the daemons are active by running the command that is shown in Example 7-5.

*Example 7-5 Check the HSM status by running ps*

---

```
[root@ltfs97 /]# ps -ef|grep dsm
root      1355      1  0 14:12 ?        00:00:01
/opt/tivoli/tsm/client/hsm/bin/dsmwatchd nodetach
root      5657      1  0 14:41 ?        00:00:00
/opt/tivoli/tsm/client/hsm/bin/dsmrecalld
root      5722  5657  0 14:41 ?        00:00:00
/opt/tivoli/tsm/client/hsm/bin/dsmrecalld
root      5723  5657  0 14:41 ?        00:00:00
/opt/tivoli/tsm/client/hsm/bin/dsmrecalld
```

---

The **dsmmigfs** command also provides the status of HSM, as shown by the output in Example 7-6.

*Example 7-6 Check the HSM status by using dsmmigfs*

---

```
[root@ltfs97 ~]# dsmmigfs query -detail
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 6.3
  Client date/time: 10/18/2016 23:17:15
```

(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

The local node has Node ID: 1  
The failover environment is active on the local node.  
The recall distribution is enabled.

```
File System Name:    /ibm/gpfs
High Threshold:     90
Low Threshold:      80
Premig Percentage:  10
Quota:              227272
Stub Size:          0
Read Starts Recall: no
Preview Size:       0
Server Name:        SERVER_A
Max Candidates:     100
Max Files:          0
Read Event Timeout: 600
Stream Seq:         0
Min Partial Rec Size: 0
Min Stream File Size: 0
MinMigFileSize:    0
Preferred Node:     1tfs97.tuc.stglabs.ibm.com    Node ID: 1
Owner Node:        1tfs97.tuc.stglabs.ibm.com    Node ID: 1
```

---

You can also ensure that the GPFS file system (named `gpfs` in this example) is managed by HSM by running the command that is shown in Example 7-7.

*Example 7-7 Check GPFS file system*

```
[root@1tfs97 /] mm1sfs gpfs|grep DMAPI
-z                Yes                Is DMAPI enabled?
```

---

To manage a file system with IBM Spectrum Archive EE, it must be data management application programming interface (DMAPI) enabled. A file system is managed by IBM Spectrum Archive EE by running the `1tfsee_config` command, which is described in 6.2, “Configuring IBM Spectrum Archive EE” on page 104.

**Permissions:** Starting HSM requires root user permissions.

If the HSM watch daemon (`dsmwatchd`) is not running, Example 7-8 shows you how to start it.

*Example 7-8 Start the HSM watch daemon*

```
[root@1tfsr11 ~]# systemctl start hsm.service
[root@1tfsr11 ~]# ps -afe | grep dsm
root    7687    1  0 08:46 ?        00:00:00
/opt/tivoli/tsm/client/hsm/bin/dsmwatchd nodetach
root    8405  6621  0 08:46 pts/1    00:00:00 grep --color=auto
dsm
```

---

If the HSM recall daemons (`dsmrecalld`) are not running, Example 7-9 shows you how to start them.

*Example 7-9 Start HSM*

---

```
[root@l1tfsm11 ~]# dsmmigfs start
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 4.80
  Client date/time: 02/11/2016 10:51:11
(c) Copyright by IBM Corporation and other(s) 1990, 2015. All Rights Reserved.
```

---

If failover operations within the IBM Spectrum Scale cluster are wanted on the node, run the `dsmmigfs enablefailover` command after you run the `dsmmigfs start` command.

## 7.2.4 IBM Spectrum Archive EE

After IBM Spectrum Archive EE is started, you can retrieve details about the node that the multi-tape management module (MMM) service was started on by running the `l1tfsee status` command. You can also use this command to determine whether the MMM service is running. The MMM is the module that manages configuration data and physical resources of IBM Spectrum Archive EE.

**Permissions:** Retrieving the status for the MMM service requires root user permissions.

If the MMM service is running correctly, you see a message similar to the message shown in Example 7-10.

*Example 7-10 Check the IBM Spectrum Archive EE status*

---

```
[root@l1tfsm11 ~]# l1tfsee status
Ctrl Node    MD      MMM      Library
9.11.120.207 Active  Active  lib_l1tfsm12
9.11.120.198 Active  Active  lib_l1tfsm11
```

---

If the MMM service is not running correctly, you see a message that is similar to the message shown in Example 7-11.

*Example 7-11 Check the IBM Spectrum Archive EE status*

---

```
[root@l1tfsm11 ~]# l1tfsee status
Ctrl Node    MD      MMM      Library
9.11.120.207 Inactive Inactive lib_l1tfsm12
9.11.120.198 Inactive Inactive lib_l1tfsm11
```

---

The Inactive status text says that IBM Spectrum Archive EE is not running on that particular EE Control Node. In our example, IBM Spectrum Archive EE is not running and must be started. Example 7-12 shows how to start IBM Spectrum Archive EE.

*Example 7-12 Start IBM Spectrum Archive EE*

---

```
[root@l1tfsm11 ~]# l1tfsee start
```

```
Library name: lib_l1tfsm12, library id: 0000013FA0520412, control node (l1tfsee_md)
IP address: 9.11.120.207.
```

```

GLES648I(00277): Running start command - sending request : lib_ltfsm12
Library name: lib_ltfsm11, library id: 0000013FA0520411, control node (ltfsee_md)
IP address: 9.11.120.198.
GLES648I(00277): Running start command - sending request : lib_ltfsm11
GLES648I(00297): Running start command - waiting for completion : lib_ltfsm12
.....
GLES639I(00136): Started the Spectrum Archive EE services for library lib_ltfsm12
with good status.
GLES648I(00297): Running start command - waiting for completion : lib_ltfsm11

GLES639I(00136): Started the Spectrum Archive EE services for library lib_ltfsm11
with good status.

```

**Note:** Starting with IBM Spectrum Archive EE V1.2.4.0, automatic failover is available for control nodes. This feature is limited to having two control nodes set. For more information about failover features, see 7.6, “IBM Spectrum Archive EE automatic node failover” on page 145.

**Important:** If the `ltfsee start` command does not return after several minutes, it might be because the firewall is running or tapes are being unmounted from the drives. The firewall service must be disabled on the IBM Spectrum Archive EE nodes. For more information, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60.

## 7.3 Upgrading components

The following sections describe the process that is used to upgrade IBM Spectrum Scale and other components of IBM Spectrum Archive EE.

### 7.3.1 IBM Spectrum Scale

Complete this task if you must update your version of IBM Spectrum Scale that is used with IBM Spectrum Archive EE.

Before any system upgrades or major configuration changes are made to your IBM Spectrum Scale cluster, review your IBM Spectrum Scale documentation and consult the IBM Spectrum Scale frequently asked question (FAQ) information that applies to your version of IBM Spectrum Scale. For more information about the IBM Spectrum Scale FAQ, see the Cluster products IBM Knowledge Center and select the IBM Spectrum Scale release under the Cluster product libraries topic in the navigation pane that applies to your installation:

[http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale\\_welcome.html](http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html)

To update IBM Spectrum Scale, complete the following steps:

1. Stop IBM Spectrum Archive EE by running the command that is shown in Example 7-13.

*Example 7-13 Stop IBM Spectrum Archive EE*

```
[root@l1tfsm11 ~]# ltfsee stop
```

```
Library name: lib_ltfsm12, library id: 0000013FA0520412, control node
(ltfsee_md) IP address: 9.11.120.207.
```

```

GLES648I(00144): Running stop command - sending request and waiting for the
completion
Library name: lib_ltfsm11, library id: 0000013FA0520411, control node
(ltfsee_md) IP address: 9.11.120.198.
GLES648I(00144): Running stop command - sending request and waiting for the
completion
GLES647I(00195): Stopped the Spectrum Archive EE services for library
lib_ltfsm12.
..
GLES647I(00195): Stopped the Spectrum Archive EE services for library
lib_ltfsm11.

```

---

2. Run the **pidof mmm** command on all EE Control Nodes until all MMM processes have been terminated.
3. Run the **pidof ltf**s command on all EE nodes until all ltf processes have been terminated.
4. Disable DSM failover by running the command that is shown in Example 7-14.

*Example 7-14 Disable failover*

---

```

[root@ltfsm11 ~]# dsmmigfs disablefailover
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 0.0
  Client date/time: 04/20/2017 11:31:18
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.

Automatic failover is disabled on this node.

```

---

5. Stop the Tivoli Storage Manager for Space Management HSM by running the command that is shown in Example 7-15.

*Example 7-15 Stop HSM*

---

```

[root@ltfsm11 ~]# dsmmigfs stop
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 4.80
  Client date/time: 02/11/2016 10:57:52
(c) Copyright by IBM Corporation and other(s) 1990, 2015. All Rights Reserved.

```

---

This command must be run on every IBM Spectrum Archive EE node.

6. Stop the watch daemon by running the command that is shown in Example 7-16.

*Example 7-16 Stop the watch daemon*

---

```

[root@ltfsm11 ~]# systemctl stop hsm.service

```

---

This command must be run on every IBM Spectrum Archive EE node.

7. Unmount GPFS by running the command that is shown in Example 7-17.

*Example 7-17 Stop GPFS*

---

```

[root@ltfs97 ~]# mmumount all
Tue Apr 16 23:43:29 JST 2013: mmumount: Unmounting file systems ...

```

---

If the `mmumount all` command results show that processes are still being used (as shown in Example 7-18), you must wait for them to finish and then run the `mmumount all` command again.

*Example 7-18 Processes running that prevent the unmounting of the GPFS file system*

---

```
[root@l1tfs97 ~]# mmumount all
Tue Apr 16 23:46:12 JST 2013: mmumount: Unmounting file systems ...
umount: /ibm/glues: device is busy.
(In some cases useful info about processes that use
the device is found by lsof(8) or fuser(1))
umount: /ibm/glues: device is busy.
(In some cases useful info about processes that use
the device is found by lsof(8) or fuser(1))
```

---

8. Shut down GPFS by running the command that is shown in Example 7-19.

*Example 7-19 Shut down GPFS*

---

```
[root@l1tfs97 ~]# mmshutdown -a
Tue Apr 16 23:46:51 JST 2013: mmshutdown: Starting force unmount of GPFS file
systems
Tue Apr 16 23:46:56 JST 2013: mmshutdown: Shutting down GPFS daemons
htohru9.ltd.sdl: Shutting down!
htohru9.ltd.sdl: 'shutdown' command about to kill process 3645
htohru9.ltd.sdl: Unloading modules from
/lib/modules/2.6.32-220.el6.x86_64/extra
htohru9.ltd.sdl: Unloading module mmfs26
htohru9.ltd.sdl: Unloading module mmfslinux
htohru9.ltd.sdl: Unloading module tracedev
Tue Apr 16 23:47:03 JST 2013: mmshutdown: Finished
```

---

9. Download the IBM Spectrum Scale update from IBM Fix Central. Extract the IBM Spectrum Scale .rpm files and install the updated .rpm files by running the command that is shown in Example 7-20.

*Example 7-20 Update IBM Spectrum Scale*

---

```
rpm -Uvh *.rpm
```

---

10. Rebuild and install the IBM Spectrum Scale portability layer by running the command that is shown in Example 7-21.

*Example 7-21 Rebuild GPFS*

---

```
mmbuildgpl
```

---

11. Start GPFS by running the command that is shown in Example 7-22.

*Example 7-22 Start GPFS*

---

```
[root@l1tfs97 ~]# mmstartup -a
Tue Apr 16 23:47:42 JST 2013: mmstartup: Starting GPFS ...
```

---

12. Mount the GPFS file system by running the command that is shown in Example 7-23.

*Example 7-23 Mount GPFS file systems*

---

```
[root@l1tfs97 ~]# mmmount all
```

13. Start the watch daemon by running the command that is shown in Example 7-24.

*Example 7-24 Start the watch daemon*

---

```
[root@ltfsm11 ~]# systemctl start hsm.service
```

---

This command must be run on every IBM Spectrum Archive EE node.

14. Start the HSM by running the command that is shown in Example 7-25.

*Example 7-25 Start HSM*

---

```
[root@ltfsm11 ~]# dsmmigfs start
IBM Tivoli Storage Manager
Command Line Space Management Client Interface
  Client Version 7, Release 1, Level 4.80
  Client date/time: 02/11/2016 11:06:39
(c) Copyright by IBM Corporation and other(s) 1990, 2015. All Rights Reserved.
```

---

This command must be run on every IBM Spectrum Archive EE node.

15. Enable failover by running the command that is shown in Example 7-26.

*Example 7-26 Enable failover*

---

```
[root@ltfsm11 ~]# dsmmigfs enablefailover
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 0.0
  Client date/time: 04/20/2017 14:51:05
(c) Copyright by IBM Corporation and other(s) 1990, 2016. All Rights Reserved.
```

Automatic failover is enabled on this node in mode ENABLED.

---

16. Start IBM Spectrum Archive EE by running the command that is shown in Example 7-27.

*Example 7-27 Start IBM Spectrum Archive EE*

---

```
[root@ltfsm11 ~]# ltfsee start

Library name: lib_ltfsm12, library id: 0000013FA0520412, control node
(ltfsee_md) IP address: 9.11.120.207.
GLES648I(00277): Running start command - sending request : lib_ltfsm12
Library name: lib_ltfsm11, library id: 0000013FA0520411, control node
(ltfsee_md) IP address: 9.11.120.198.
GLES648I(00277): Running start command - sending request : lib_ltfsm11
GLES648I(00297): Running start command - waiting for completion : lib_ltfsm12
.....
GLES639I(00136): Started the Spectrum Archive EE services for library
lib_ltfsm12 with good status.
GLES648I(00297): Running start command - waiting for completion : lib_ltfsm11

GLES639I(00136): Started the Spectrum Archive EE services for library
lib_ltfsm11 with good status.
```

---

Check the status of each component. Optionally, you can check the status of each component when it is started, as described in 7.2, "Status information" on page 134.

## 7.3.2 IBM Spectrum Archive LE+ component

For more information about how to upgrade the IBM Spectrum Archive LE+ component, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60. Because the IBM Spectrum Archive LE+ component is a component of IBM Spectrum Archive EE, it is upgraded as part of the IBM Spectrum Archive EE upgrade.

## 7.3.3 Hierarchical Storage Management

For more information about how to upgrade HSM, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60. Because HSM is a component of IBM Spectrum Archive EE, it is upgraded as part of the IBM Spectrum Archive EE upgrade.

## 7.3.4 IBM Spectrum Archive EE

For more information about how to upgrade IBM Spectrum Archive EE, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60.

# 7.4 Starting and stopping IBM Spectrum Archive EE

This section describes how to start and stop IBM Spectrum Archive EE.

## 7.4.1 Starting IBM Spectrum Archive EE

Run the `ltfsee start` command to start the IBM Spectrum Archive LE+ component, the LTFSEE monitoring daemon, and the IBM Spectrum Archive EE system. The HSM components must be running before you can use this command. You can run the `ltfsee start` command on any IBM Spectrum Archive EE node in the cluster.

For example, to start IBM Spectrum Archive EE, run the command that is shown in Example 7-28.

*Example 7-28 Start IBM Spectrum Archive EE*

---

```
[root@ltfsm11 ~]# ltfsee start

Library name: lib_ltfsm12, library id: 0000013FA0520412, control node (ltfsee_md)
IP address: 9.11.120.207.
GLES648I(00277): Running start command - sending request : lib_ltfsm12
Library name: lib_ltfsm11, library id: 0000013FA0520411, control node (ltfsee_md)
IP address: 9.11.120.198.
GLES648I(00277): Running start command - sending request : lib_ltfsm11
GLES648I(00297): Running start command - waiting for completion : lib_ltfsm12
.....
GLES639I(00136): Started the Spectrum Archive EE services for library lib_ltfsm12
with good status.
GLES648I(00297): Running start command - waiting for completion : lib_ltfsm11

GLES639I(00136): Started the Spectrum Archive EE services for library lib_ltfsm11
with good status.
```

---

**Important:** If the `ltfsee start` command does not return after several minutes, it might be because the firewall is running or unmounting tapes from drives. The firewall service must be disabled on the IBM Spectrum Archive EE nodes. For more information, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60.

You can confirm that IBM Spectrum Archive EE is running by referring to the steps in Example 7-10 on page 138 or by running the command in Example 7-29.

*Example 7-29 Check the status of all available IBM Spectrum Archive EE nodes*

---

```
[root@ltfsm11 ~]# ltfsee info nodes
Node ID  Status      Node IP      Drives  Ctrl Node  Library      Node Group
Host Name
2        Available  9.11.120.224    2  Yes        lib_ltfsm12  GO ltfsm12
1        Available  9.11.120.207    2  Yes(Active) lib_ltfsm11  GO
ltfsm11
```

---

## 7.4.2 Stopping IBM Spectrum Archive EE

The `ltfsee stop` command stops the IBM Spectrum Archive EE system on all EE Control Nodes. To stop IBM Spectrum Archive EE, IBM Spectrum Archive LE+, and LTFSEE monitoring daemon on all EE Nodes, use the stop option, as shown in Example 7-30.

*Example 7-30 Stop IBM Spectrum Archive EE*

---

```
[root@ltfsm11 ~]# ltfsee stop

Library name: lib_ltfsm12, library id: 0000013FA0520412, control node (ltfsee_md)
IP address: 9.11.120.207.
GLES648I(00144): Running stop command - sending request and waiting for the
completion
Library name: lib_ltfsm11, library id: 0000013FA0520411, control node (ltfsee_md)
IP address: 9.11.120.198.
GLES648I(00144): Running stop command - sending request and waiting for the
completion
GLES647I(00195): Stopped the Spectrum Archive EE services for library
lib_ltfsm12.
..
GLES647I(00195): Stopped the Spectrum Archive EE services for library
lib_ltfsm11.
```

---

In some cases, you might see the GLESM658I informational message if there are active jobs on the job queue in IBM Spectrum Archive EE:

```
There are still scans in progress.
To terminate Spectrum Archive EE for this library,
run the ltfsee stop command with -f (force) option.
```

If you are sure that you want to stop IBM Spectrum Archive EE, run the `ltfsee stop` command with the `-f` option, which forcefully stops any running IBM Spectrum Archive EE jobs abruptly.

## 7.5 IBM Spectrum Archive EE database backup

IBM Spectrum Archive EE uses databases to store important system states and configurations for healthy operations. Starting with IBM Spectrum Archive EE V1.2.4.0, database backup is performed automatically and stores as many as three backups. The backup files can be found under the `/var/opt/ibm/ltfsee/local/dbbackup` directory.

Database backup is performed automatically whenever a significant modification is performed to the cluster that requires updating the original database. These changes include the following commands:

- ▶ `ltfsee pool create`
- ▶ `ltfsee pool delete`
- ▶ `ltfsee pool add`
- ▶ `ltfsee pool remove`
- ▶ `ltfsee drive add`
- ▶ `ltfsee drive remove`
- ▶ `ltfsee_config -m ADD_CTRL_NODE`
- ▶ `ltfsee_config -m ADD_NODE`
- ▶ `ltfsee_config -m REMOVE_NODE`
- ▶ `ltfsee_config -m SET_CTRL_NODE`

The backups are performed when IBM Spectrum Archive EE is running, part of the monitor daemon checks periodically if the original database files have been modified by the above operations. If the thread detects that there has been a modification, then it will go into a grace period. If no further modifications are performed during a set amount of time during this grace period, a backup will then occur.

This technique is intended to prevent repeating backups in a short amount of time, for example when multiple tapes are being added to a pool when users perform `ltfsee pool add` on multiple tapes. Instead, it performs a backup after the command finished executing and no further changes are made in a set time limit.

These backups are crucial in rebuilding an IBM Spectrum Archive EE cluster if a server gets corrupted or the GPFS file system needs to be rebuilt. After reinstalling the IBM Spectrum Archive EE onto the server, replace the `.global` and `.lib` database files under the `path/to/gpfs/filesystem/.ltfsee/config` directory with the database files backed up from the `/var/opt/ibm/ltfsee/local/dbbackup/` directory.

## 7.6 IBM Spectrum Archive EE automatic node failover

This section describes IBM Spectrum Archive EE automatic failover features, and the new LTFSEE monitoring daemon and the updated commands to display any nodes that are having issues.

**Note:** The features mentioned in this subsection are introduced in IBM Spectrum Archive EE V1.2.4.0.

## 7.6.1 IBM Spectrum Archive EE monitoring daemon

When IBM Spectrum Archive EE is started, a monitoring daemon is started on each node to monitor various critical components that make up the software:

- ▶ MMM
- ▶ IBM Spectrum Archive LE+
- ▶ Remote IBM Spectrum Archive EE monitoring daemon
- ▶ GPFS
- ▶ Rpcbind
- ▶ Rsyslog
- ▶ SSH

Figure 7-1 shows the IBM Spectrum Archive EE component diagram.

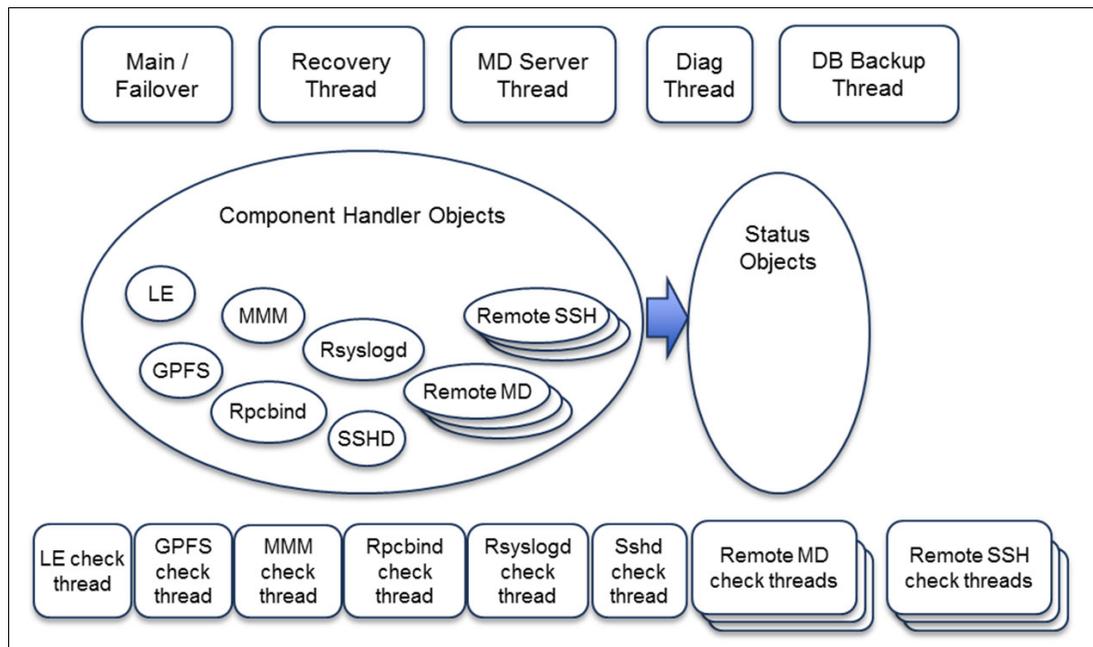


Figure 7-1 IBM Spectrum Archive EE component diagram

Components, such as MMM, IBM Spectrum Archive LE+, and remote monitoring daemons, have automatic recovery features. If one of those three components is hung, or has crashed, the monitor daemon performs a recovery to restart it. In an environment where a redundant control node is available and MMM is no longer responding or alive, an automatic failover takes place and the redundant control node becomes the new active control node.

If there is only one control node available in the cluster, then an in-place failover occurs to bring back the MMM process on that control node.

**Note:** Only the active control node and the redundant control node's monitor daemon monitors each other, while the active control node also monitors non control node's monitoring daemon.

If the monitoring daemon has hung or been killed and there is no redundant control node, restart IBM Spectrum Archive EE to start a new monitoring daemon.

As for the rest of the components, currently there are no automatic recovery actions that can be performed. In the current release, the **ltfsee info nodes** command has been updated to also report any node errors that might be occurring. If GPFS, rpcbind, or rsyslog are having problems, the issues can be viewed by using the updated command.

Example 7-31 shows output from running **ltfsee info nodes** when one node has not started **rpcbind**. To correct this error, start **rpcbind** on the designated node and EE will refresh itself and the node will become available.

*Example 7-31 The ltfsee info nodes output*

---

```
[root@ltfsm1 ~]# ltfsee info nodes
Node ID  Status      Node IP      Drives  Ctrl Node  Library  Node Group
Host Name
2        Error        9.11.120.224  2  Yes        lib_ltfsm12 G0
ltfsm1
1        Available   9.11.120.207  2  Yes(Active) lib_ltfsm1 G0
ltfsm1

Problem Detected:
Node ID  Error Modules
      2  rpcbind;
```

---

In addition to the automatic failover, there is also an option to perform a manual failover by running the **ltfsee failover** command. This command is used to fail over the MMM process to a redundant control node. This command is only available for use when a redundant control node exists. Example 7-32 shows output from running the **ltfsee failover** command.

*Example 7-32 The ltfsee failover command*

---

```
[root@ltfsm1 ~]# ltfsee failover -l lib_ltfsm1
GLESL659I(00168): failover on library lib_ltfsm1 is requested.
```

---

## 7.7 Tape library management

This section describes how to use **ltfsee** commands to add and remove tape drives and tape cartridges from your LTFSS library.

### 7.7.1 Adding tape cartridges

This section describes how to add tape cartridges in IBM Spectrum Archive EE. An unformatted tape cartridge cannot be added to the IBM Spectrum Archive EE library. However, you can format a tape when you add it to a tape cartridge pool. The process of formatting a tape in LTFSS creates the required LTFSS partitions on the tape.

After tape cartridges are added through the I/O station, or after they are inserted directly into the tape library, you might have to run an **ltfsee retrieve** command. First, run the **ltfsee info tapes** command. If the tape cartridges are missing, run the **ltfsee retrieve** command, which synchronizes the data for these changes between the IBM Spectrum Archive EE system and the tape library.

This process occurs automatically. However, if the tape does not appear within the **ltfsee info tapes** command output, you can force a rebuild of the inventory (synchronization of IBM Spectrum Archive EE inventory with the tape library's inventory).

## Data tape cartridge

To add a tape cartridge (that was previously used by LTFS) to the IBM Spectrum Archive EE system, complete the following steps:

1. Insert the tape cartridge into the I/O station.
2. Run the **ltfsee info tapes** command to see whether your tape appears in the list, as shown in Example 7-33. In this example, the **-l** option is used to limit the tapes to one tape library.

*Example 7-33 Run the ltfsee info tapes command to check whether a tape cartridge must be synchronized*

```
[root@ltfsm11 ~]# ltfsee info tapes -l TS3310
```

Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)	Pool	Library	Address	Drive	Appendable
TST200L6	Valid	L6	2242	0	2242	0	test1	TS3310	4117	-	yes
LTFS01L7	Valid	L7	5338	0	5338	0	test	TS3310	4119	-	yes
PFE001L6	Valid	L6	2242	95	2146	0	test1	TS3310	4118	-	yes

Tape cartridge 2MB346L5 is not in the list.

3. Because tape cartridge 2MB346L5 is not in the list, synchronize the data in the IBM Spectrum Archive EE inventory with the tape library by running the **ltfsee retrieve** command, as shown in Example 7-34.

*Example 7-34 Synchronize the tape*

```
[root@ltfsm11 ~]# ltfsee retrieve
```

Library name: TS4500, library serial: 0000078AA0040405, control node (MMM) IP address: 9.155.118.177.  
GLESLO36I(00184): Retrieve completed.

Library name: TS3310, library serial: 000001300013\_LLC, control node (MMM) IP address: 9.155.118.179.  
GLESLO36I(00184): Retrieve completed.

4. Repeating the **ltfsee info tapes** command shows that the inventory was corrected, as shown in Example 7-35.

*Example 7-35 Tape cartridge 2FC143L5 is synchronized*

```
[root@ltfsm11 ~]# ltfsee info tapes -l TS3310
```

Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)	Pool	Library	Address	Drive	Appendable
TST200L6	Valid	L6	2242	0	2242	0	test1	TS3310	4117	-	yes
LTFS01L7	Valid	L7	5338	0	5338	0	test	TS3310	4119	-	yes
PFE001L6	Valid	L6	2242	95	2146	0	test1	TS3310	4118	-	yes
2MB346L5	Valid	L5	1327	0	1327	0	-	TS3310	1232	-	yes

5. If necessary, move the tape cartridge from the I/O station to a storage slot by running the **ltfsee tape move** command (see Example 7-36) with the **homeslot** option.

*Example 7-36 Move tape to homeslot*

```
[root@ltfsm11 ~]# ltfsee tape move homeslot -t 2MB346L5 -l TS3310
```

Tape 2FC143L5 is moved successfully.

6. Add the tape cartridge to a tape cartridge pool. If the tape cartridge contains actual data to be added to LTFS, you must import it first before you add it. Run the **ltfsee import** command to add the tape cartridge into the IBM Spectrum Archive EE library and import the files on that tape cartridge into the IBM Spectrum Scale namespace as stub files.

If you have no data on the tape cartridge (but it is already formatted for LTFS), add it to a tape cartridge pool by running the **ltfsee pool add** command.

## Scratch cartridge

To add a scratch cartridge to the IBM Spectrum Archive EE system, complete the following steps:

1. Insert the tape cartridge into the I/O station.
2. Synchronize the data in the IBM Spectrum Archive EE inventory with the tape library by running the **ltfsee retrieve** command, as shown in Example 7-34 on page 148.
3. If necessary, move the tape cartridge from the I/O station to a storage slot by running the **ltfsee tape move** command with the **homeslot** option, as shown in Example 7-36 on page 148.
4. To format the tape cartridge while adding it to a tape cartridge pool, use the **-f** or **--format** option. For example, Example 7-37 shows the output of the **ltfsee pool add** command.

### *Example 7-37 Format a scratch tape*

---

```
[root@ltfsm11 ~]# ltfsee pool add -p test -t 2FC143L5 -l TS3310 -f
GLESLO42I(00640): Adding tape 2FC143L5 to storage pool test.
Tape 2FC143L5 successfully formatted.
Added tape 2FC143L5 to pool test successfully.
```

---

For more information about other formatting options, see 7.7.3, “Formatting tape cartridges” on page 150.

## 7.7.2 Moving tape cartridges

This section summarizes the IBM Spectrum Archive EE commands that can be used for moving tape cartridges.

### **Moving to different tape cartridge pools**

If a tape cartridge contains any files, IBM Spectrum Archive EE will not allow you to move a tape cartridge from one tape cartridge pool to another. If this move is attempted, you receive an error message, as shown in Example 7-38.

### *Example 7-38 Error message when removing a tape cartridge from a pool with migrated/saved files*

---

```
[root@ltfsm11 ~]# ltfsee pool remove -p test1 -t PFE001L6 -l TS3310
GLESLO43I(00832): Removing tape PFE001L6 from storage pool test1.
GLESLO357E(00868): Tape PFE001L6 has migrated files or saved files. It has not been
removed from the pool.
```

---

However, you can remove an empty tape cartridge from one tape cartridge pool and add it to another tape cartridge pool, as shown in Example 7-39.

### *Example 7-39 Remove an empty tape cartridge from one tape cartridge pool and add it to another*

---

```
[root@ltfsm11 ~]# ltfsee pool remove -p test1 -t TST200L6 -l TS3310
GLESLO43I(00832): Removing tape TST200L6 from storage pool test1.
Removed tape TST200L6 from pool test1 successfully.
```

```
[root@ltfsm11 ~]# ltfsee pool add -p test -t TST200L6 -l TS3310
GLESLO42I(00640): Adding tape TST200L6 to storage pool test.
Tape TST200L6 successfully checked.
Added tape TST200L6 to pool test successfully.
```

---

Before you remove a tape cartridge from one tape cartridge pool and add it to another tape cartridge pool, reclaim the tape cartridge to ensure that no files remain on the tape when it is removed. For more information, see 7.19, “Reclamation” on page 192.

### Moving to the homeslot

To move a tape cartridge from a tape drive to its homeslot in the tape library, use the command that is shown in Example 7-40. You might want to use this command in cases where a tape cartridge is loaded in a tape drive and you want to unload it.

*Example 7-40 Move a tape cartridge from a tape drive to its homeslot*

---

```
[root@ltfsm11 ~]# ltfsee tape move homeslot -t TST200L6 -p test -l TS3310
GLES373I(00636): Moving tape TST200L6
Tape TST200L6 is unmounted because it is inserted into the drive.
Tape TST200L6 is moved successfully.
```

---

### Moving to the I/O station

The command that is shown in Example 7-41 moves a tape cartridge to the ieslot (I/O station). This might be required when tape cartridges are exported or exported offline.

*Example 7-41 Move a tape cartridge to the ieslot after an export offline operation*

---

```
[root@ltfsm11 ~]# ltfsee export -p test -l TS3310 -t TST200L6 -o "LTFS EE
Redbooks"
Export of tape TST200L6 has been requested...
Export of tape TST200L6 complete.
Updated offline state of tape TST200L6.
Removing tape TST200L6 from storage pool test.

[root@ltfsm11 ~]# ltfsee tape move ieslot -t TST200L6 -l TS3310
GLES373I(00636): Moving tape 2TST200L6
Tape TST200L6 is unmounted because it is inserted into the drive.
Tape 2TST200L6 is moved successfully.
```

---

The move can be between homeslot and ieslot or tape drive and homeslot. If the tape cartridge belongs to a tape cartridge pool and online (not in the Offline state), the request to move it to the ieslot fails. After a tape cartridge is moved to ieslot, the tape cartridge cannot be accessed from IBM Spectrum Archive EE. If the tape cartridge contains migrated files, the tape cartridge should not be moved to ieslot without first exporting the tape cartridge.

A tape cartridge in ieslot cannot be added to a tape cartridge pool. Such a tape cartridge must be moved to home slot before adding it.

## 7.7.3 Formatting tape cartridges

This section describes how to format a medium in the library for the IBM Spectrum Archive EE. There are two options that can be used: `-f` or `-F`. The `-f` option formats only a scratch tape. The format fails if the tape cartridge is already LTFS formatted. If the tape cartridge is already formatted for IBM Spectrum Archive EE, the format fails, as shown in Example 7-42.

*Example 7-42 Format failure*

---

```
[root@ltfsm11 ~]# ltfsee pool add -p test1 -l 1TS3310 -t TST200L6 -f
GLES4042I(00640): Adding tape TST200L6 to storage pool test1.
```

---

GLESL091E(00683): This operation is not allowed to this state of tape. Need to check the status of Tape TST200L6 by using the `ltfsee info tapes` command.

---

When the formatting is requested, IBM Spectrum Archive EE attempts to mount the target medium to obtain the medium condition. The medium is formatted if the `mount` command finds any of the following conditions:

- ▶ The medium was not yet formatted for LTFS.
- ▶ The medium has an invalid label.
- ▶ Labels in both partitions do not have the same value.

If none of these conditions are found, LTFS assumes that the medium already is formatted, so it is not formatted by default. However, the `-F` option formats any tape, even if it is formatted but does not contain any referenced data. Tape cartridges that were used by IBM Spectrum Archive EE must be reclaimed to allow them to be reformatted because the reclamation process removes the referenced data.

Example 7-43 shows a tape cartridge being formatted by using the `-F` option.

*Example 7-43 Forced format*

---

```
[root@ltfsm11 ~]# ltfsee pool add -p test1 -l LTS3310 -t TST200L6 -F
GLESL042I(00640): Adding tape TST200L6 to storage pool test1.
Tape TST200L6 successfully formatted.
Added tape TST200L6 to pool test1 successfully.
```

---

Multiple tape cartridges can be formatted by specifying multiple tape VOLSERs. Example 7-44 shows three tape cartridges that are formatted sequentially or simultaneously.

*Example 7-44 Format multiple tape cartridges*

---

```
[root@ltfsm11 ~]# ltfsee pool add -p test1 -l LTS3310 -t TST200L6 -t LTFS01L7 -e
GLESL042I(00640): Adding tape TST200L6 to storage pool test1.
GLESL042I(00640): Adding tape TFS01L7 to storage pool test1.
.
Tape ST200L6 successfully formatted.
Added tape ST200L6 to pool test1 successfully.
Tape TFS01L7 successfully formatted.
Added tape TFS01L7 to pool test1 successfully.
```

---

When multiple format jobs are submitted, IBM Spectrum Archive EE uses all available drives with the 'g' drive attribute for the format jobs, which are done in parallel.

## Active file check before formatting tape cartridges

Some customers (such as those in the video surveillance industry) might want to retain data only for a certain retention period and then reuse the tape cartridges. Running the **reconciliation** and **reclamation** commands are the most straightforward method. However, this process might take a long time if there are billions of small files in GPFS, because the command checks every file in GPFS and deletes files on the tape cartridge one by one.

The fastest method is to manage the pool and identify data in tape cartridges that has passed the retention period. Customers can then remove the tape cartridge and add to a new pool by reformatting the entire tape cartridge. This approach saves time, but customers need to be certain that the tape cartridge does not have any active data.

To be sure that a tape cartridge is format-ready, this section introduces a new option, **-E**, to the **ltfsee pool remove** command. When run, this command checks whether the tape cartridge contains any active data. If all the files in the tape cartridge have already been deleted in GPFS, the command determines that the tape cartridge is effectively empty and removes the tape cartridge from the pool. If the tape cartridge still has active data, the command will not remove it. No **reconciliation** command is necessary before this command.

When **-E** is run, the command performs the following steps:

1. Determine whether the specified tape cartridge is in the specified pool and is not mounted.
2. Reserve the tape cartridge so that no migration will occur to the tape.
3. Read the volume cache (GPFS file) for the tape cartridge. If any file entries exist in the volume cache, check whether the corresponding GPFS stub file exists, as-is or renamed. This is the same check routine as the current reclaim **-q** (quick reconcile) command.
4. If the tape cartridge is empty or has files but all of them have already been deleted in GPFS (not renamed), remove the tape cartridge from the pool.

Example 7-45 shows the output of the **active file check -E** command with three tape cartridges. VTAPE3L5 is an empty tape cartridge. VTAPE4L5 and VTAPE5L5 have files migrated to them and all corresponding GPFS stub files deleted, but reconciliation has not yet been performed.

*Example 7-45 Removing tape cartridges from pool with active file check*

---

```
[root@io0 src]# /opt/ibm/ltfsee/bin/ltfsee pool remove -p pool1 -l lib1 -t
VTAPE3L5 -t VTAPE4L5 -t VTAPE5L5 -E
GLESL584I(00820): Reserving tapes.
GLESL135I(00851): Reserved tapes: VTAPE3L5 VTAPE4L5 VTAPE5L5 .
GLESL043I(00881): Removing tape VTAPE3L5 from storage pool pool1.
Removed tape VTAPE3L5 from pool pool1 successfully.
GLESL043I(00881): Removing tape VTAPE4L5 from storage pool pool1.
GLESL575I(00907): Tape VTAPE4L5 contains files but all of them have been
deleted in GPFS.
GLESL572I(00938): Removed tape VTAPE4L5 from pool pool1 successfully. Format
the tape when adding back to a pool.
GLESL043I(00881): Removing tape VTAPE5L5 from storage pool pool1.
GLESL575I(00907): Tape VTAPE5L5 contains files but all of them have been
deleted in GPFS.
GLESL572I(00938): Removed tape VTAPE5L5 from pool pool1 successfully. Format
the tape when adding back to a pool.
GLESL137I(00980): Removing tape reservations.
```

---

5. If the tape cartridge has a valid, active file, the check routine aborts on the first hit and goes on to the next specified tape cartridge. The command will not remove the tape cartridge from the pool.

In Example 7-46, VTAPE7L5 has files migrated to the tape cartridge, but all corresponding GPFS stub files have been deleted. VTAPE8L5 has active renamed data whereas VTAPE9L5 has active data. Upon running the **-E** command, only VTAPE7L5 is removed. VTAPE8L5 and VTAPE9L5 will not be removed.

*Example 7-46 Tape cartridges containing active data are removed from the pool*

---

```
[root@io0 archive]# ltfsee pool remove -p pool1 -t VTAPE7L5 -t VTAPE8L5 -t
VTAPE9L5 -E
GLESL584I(00820): Reserving tapes.
```

```

GLESL135I(00851): Reserved tapes: VTAPE7L5 VTAPE8L5 VTAPE9L5 .
GLESL043I(00881): Removing tape VTAPE7L5 from storage pool pool1.
GLESL575I(00907): Tape VTAPE7L5 contains files but all of them have been
deleted in GPFS.
GLESL572I(00938): Removed tape VTAPE7L5 from pool pool1 successfully. Format
the tape when adding back to a pool.
GLESL043I(00881): Removing tape VTAPE8L5 from storage pool pool1.
GLESL357E(00890): Tape VTAPE8L5 has migrated files or saved files. It has not
been removed from the pool.
GLESL043I(00881): Removing tape VTAPE9L5 from storage pool pool1.
GLESL357E(00890): Tape VTAPE9L5 has migrated files or saved files. It has not
been removed from the pool.
GLESL137I(00980): Removing tape reservations.

```

---

The active file check applies to all data types that the current IBM Spectrum Archive EE might store to a tape cartridge:

- ▶ Normal migrated files
- ▶ Saved objects such as empty directory and link files

Another approach is to run `mmapplypolicy` to list all files that have been migrated to the designated tape cartridge ID. However, if the IBM Spectrum Scale file system has over 1 billion files, the `mmapplypolicy` scan might take a long time.

## 7.7.4 Removing tape drives

When the LTFS mounts the library, all tape drives are inventoried by default. The following procedure can be started when a tape drive requires replacing or repairing and must be physically removed from the library. The same process also must be carried out when firmware for the tape drive is upgraded. If a tape is in the drive and a job is in-progress, the tape is unloaded automatically when the job completes.

After mounting the library, the user can run `ltfsee` commands to manage the library and to correct a problem if one occurs.

To remove a tape drive from the library, complete the following steps:

1. Remove the tape drive from the IBM Spectrum Archive EE inventory by running the `ltfsee drive remove` command, as shown in Example 7-47. A medium in the tape drive is automatically moved to the home slot (if one exists).

### Example 7-47 Remove a tape drive

```

[root@ltfsm11 ~]# ltfsee drive remove -d 1013000655 -l lib_ltfsm11
GLESL121I(00282): Drive serial 1013000655 is removed from LTFS EE drive list.

```

```

[root@ltfsm11 ~]# ltfsee info drives
Drive S/N  Status      Type    Role  Library  Address  Node ID  Tape  Node
Group
1013000667 Not mounted  LT06    mrg   lib_ltfsm12  256     1      -    G0
1013000110 Not mounted  LT06    mrg   lib_ltfsm12  257     1      -    G0
1013000692 Not mounted  LT06    mrg   lib_ltfsm12  258     1      -    G0
1013000694 Not mounted  LT06    mrg   lib_ltfsm11  256     2      -    G0
1013000688 Not mounted  LT06    mrg   lib_ltfsm11  257     2      -    G0
1013000655 Stock        UNKNOWN ---   lib_ltfsm11  258     -      -    -

```

2. Physically remove the tape drive from the tape library.

For more information about how to remove tape drives, see the IBM Knowledge Center for your IBM tape library.

## 7.7.5 Adding tape drives

Add the tape drive to the LTFS inventory by running the `ltfsee drive add` command, as shown in Example 7-48.

Optionally, drive attributes can be set when adding a tape drive. Drive attributes are the logical OR of the attributes: `migrate(4)`, `recall(2)`, and `generic(1)`. If the individual attribute is set, any corresponding jobs on the job queue can be run on that drive. The drive attributes can be specified after the colon (`:`) following the drive serial, and must be a decimal number.

In Example 7-48, 6 is the logical OR of `migrate(4)` and `recall(2)`, so migration jobs and recall jobs can be performed on this drive. For more information, see 7.22, “Drive Role settings for job assignment control” on page 203.

The node ID is required for the `add` command.

*Example 7-48 Add a tape drive*

```
[root@ltfsm11 ~]# ltfsee drive add -d 1013000655:6 -n 2 -l lib_ltfsm11
GLESL119I(00174): Drive 1013000655:6 added successfully.
```

```
[root@ltfsm11 ~]# ltfsee info drives
```

Drive S/N	Status	Type	Role	Library	Address	Node ID	Tape	Node Group
1013000667	Not mounted	LT06	mrg	lib_ltfsm12	256	1	-	G0
1013000110	Not mounted	LT06	mrg	lib_ltfsm12	257	1	-	G0
1013000692	Not mounted	LT06	mrg	lib_ltfsm12	258	1	-	G0
1013000694	Not mounted	LT06	mrg	lib_ltfsm11	256	2	-	G0
1013000688	Not mounted	LT06	mrg	lib_ltfsm11	257	2	-	G0
1013000655	Not mounted	LT06	mrg	lib_ltfsm11	258	2	-	G0

## 7.8 Tape storage pool management

This section describes how to use the `ltfsee pool` command to manage tape cartridge pools with IBM Spectrum Archive EE.

**Permissions:** Managing tape cartridge pools by running the `ltfsee pool` command requires root user permissions.

To perform file migrations, it is first necessary to create and define tape cartridge pools, which are the targets for migration. It is then possible to add or remove tape cartridges to or from the tape cartridge pools.

Consider the following rules and recommendations for tape cartridge pools:

- ▶ Before adding tape cartridges to a tape cartridge pool, the tape cartridge must first be in the homeslot of the tape library. For more information about moving to the homeslot, see 7.7.2, “Moving tape cartridges” on page 149.

- ▶ Multiple jobs can be performed in parallel when more than one tape cartridge is defined in a tape cartridge pool. Have multiple tape cartridges in each tape cartridge pool to increase performance.
- ▶ The maximum number of drives in a node group that is used for migration for a particular tape cartridge pool can be limited by setting the **mountlimit** attribute for the tape cartridge pool. The default is 0, which is unlimited. For more information about the **mountlimit** attribute, see 8.2, “Maximizing migration performance with redundant copies” on page 227.
- ▶ After a file is migrated to a tape cartridge pool, it cannot be migrated again to another tape cartridge pool before it is recalled.
- ▶ When tape cartridges are removed from a tape cartridge pool but not exported from IBM Spectrum Archive EE, they are no longer targets for migration or recalls.
- ▶ When tape cartridges are exported from IBM Spectrum Archive EE system by running the **ltfsee export** command, they are removed from their tape cartridge pool and the files are not accessible for recall.

### 7.8.1 Creating tape cartridge pools

This section describes how to create tape cartridge pools for use with IBM Spectrum Archive EE. Tape cartridge pools are logical groupings of tape cartridges within IBM Spectrum Archive EE. The groupings might be based on their intended function (for example, `OnsitePool` and `OffsitePool`) or based on their content (for example, `MPEGpool` and `JPEGpool`). However, you must create at least one tape cartridge pool.

You create tape cartridge pools by using the **create** option of the **ltfsee pool** command. For example, the command that is shown in Example 7-49 creates the tape cartridge pool named `MPEGpool`.

*Example 7-49 Create a tape cartridge pool*

---

```
[root@lufs97 /]# ltfsee pool create -p MPEGpool
```

---

For single tape library systems, the **-l** option (library name) can be omitted. For two tape library systems, the **-l** option is used to specify the library name.

For single node group systems, the **-g** option (node group) can be omitted. For multiple node group systems, the **-g** option is used to specify the node group.

The default tape cartridge pool type is a regular pool. If a WORM pool is wanted, supply the **--worm physical** option.

The pool names are case-sensitive and can be duplicated in different tape libraries. No informational messages are shown at the successful completion of the command. However, you can confirm that the pool was created by running the **ltfsee info pools** command.

### 7.8.2 Deleting tape cartridge pools

This section describes how to delete tape cartridge pools for use with IBM Spectrum Archive EE. Delete tape cartridge pools by using the **delete** option of the **ltfsee pool** command. For example, the command in Example 7-50 on page 156 deletes the tape cartridge pool that is named `MPEGpool`.

*Example 7-50 Delete a tape cartridge pool*

---

```
[root@l1fs97 ~]# ltfsee pool delete -p MPEGpool
```

---

For single tape library systems, the `-l` option (library name) can be omitted. For two tape library systems, the `-l` option is used to specify the library name.

When deleting a tape cartridge pool, the `-g` option (node group) can be omitted.

No informational messages are shown after the successful completion of the command.

**Important:** If the tape cartridge pool contains tape cartridges, the tape cartridge pool cannot be deleted until the tape cartridges are removed.

You cannot use IBM Spectrum Archive EE to delete a tape cartridge pool that still contains data. Example 7-51 shows an attempt to delete a tape cartridge pool with tape cartridges in it and its resulting error message.

*Example 7-51 Error message when deleting a tape cartridge pool containing data*

---

```
[root@l1fsm11 ~]# ltfsee pool delete -p copy_l1fsm11 -l lib_l1fsm11
GLESLO44E(00467): Failed to delete storage pool. The pool copy_l1fsm11 is not
empty. Tapes must be removed from the storage pool before it can be deleted.
```

---

To allow the deletion of the tape cartridge pool, you must remove all tape cartridges from it by running the `ltfsee pool remove` command, as described in 7.7.2, “Moving tape cartridges” on page 149.

## 7.9 Pool capacity monitoring

Monitoring the pool capacity is a feature that is available starting with IBM Spectrum Archive EE V1.2.4.0. This feature monitors designated pools set by the users. By using SNMP traps, the feature also informs users if the set pools are running low on space, or have completely run out of space if a migration were to fail due to capacity reasons. The pool capacity monitoring feature benefits customers by giving them enough time to plan ahead before pool space is depleted.

To enable the pool capacity monitoring feature, temporarily stop your IBM Spectrum Archive EE environment to modify the `ltfsee.config` file under your GPFS file system. Perform the following steps to stop IBM Spectrum Archive EE and modify the `ltfsee.config` file:

1. Run `ltfsee stop` to stop IBM Spectrum Archive EE.
2. In your favorite text editor, open the `ltfsee.config` file under `path/to/gpfs/filesystem/.ltfsee/config/ltfsee.config` and add the following line:  
`POOL_MON_ENABLE yes`
3. Run `ltfsee start` to start IBM Spectrum Archive EE.

With the added line to the IBM Spectrum Archive EE config file, each time IBM Spectrum Archive EE is started up, a monitoring thread is enabled so that it can monitor set pools for low space. The pool capacity monitor thread checks each set pool for a low capacity every 30 minutes, and sends a trap every 24 hours.

Example 7-52 shows the contents of the `ltfsee.config` file after modifying it to enable pool capacity monitoring. Yours might or might not have additional modifications.

*Example 7-52 Enabling pool capacity monitoring in ltfsee.config*

---

```
[root@ltfseesrv1 ~]# cat /ibm/glues/.ltfsee/config/ltfsee.config
DISABLE_AFM_CHECK yes
POOL_MON_ENABLE yes
```

---

After IBM Spectrum Archive EE is running again, set a pool capacity threshold for designated pools. Run the `ltfsee pool show -p <pool_name>` command to view the new attributes for your pools. Example 7-53 outputs the attributes of a pool from the `ltfsee pool show` command.

*Example 7-53 The ltfsee pool show command*

---

```
[root@ltfseesrv1 ~]# ltfsee pool show -p pool1
Attribute          Value
poolname           primary_home
poolid             4c6d25a1-6d1f-43fc-bbd6-d1c19afcee11
devtype            LT0
mediarestriction   none
format             DEFAULT (0xFFFFFFFF)
worm               no (0)
nodegroup          G0
fillpolicy         Default
owner              System
mountlimit         0
lowspacewarningenable   true
lowspacewarningthreshold  0
nospacewarningenable    true
mode               normal
```

**Note:** An additional pool attribute has been introduced with V1.2.6 (`media_restriction`). The `pool set / pool show` commands shall be updated to set the cartridge type and to show the current setting of the cartridge type as the `media_restriction` attribute.

The `media_restriction` value defines the cartridge type, and reflects the last two letters of the cartridge bar code. The value can be L5, L6, L7, L8, M8, JB, JC, JD, JK, JL, JY, JZ, or none. If the specified type is none, any cartridge type is accepted.

---

By default, the `lowspacewarningenable` and `nospacewarningenable` attributes are set to `true`, and `lowspacewarningthreshold` is set to 0, which indicates no traps will be sent. An SNMP trap is sent for no space remaining in the pool when migrations failed due to pool space being depleted.

The `lowspacewarningthreshold` attribute value is set in TiB. To modify the attributes in each pool, use the `ltfsee pool set -p <pool_name> -a <attribute> -v <value>` command.

Example 7-54 shows the output of modifying pool1's **lowspacewarningthreshold** attribute to 2 TiB.

*Example 7-54 Output of modifying pool1's lowspacewarningthreshold attribute*

---

```
[root@l1tfseesrv1 ~]# l1tfsee pool set -p pool1 -a lowspacewarningthreshold -v 2
[root@l1tfseesrv1 ~]# l1tfsee pool show -p pool1
Attribute                Value
poolname                 primary_home
poolid                   4c6d25a1-6d1f-43fc-bbd6-d1c19afcee11
devtype                  LTO
mediarestriction        none
format                   DEFAULT (0xFFFFFFFF)
worm                     no (0)
nodegroup                G0
fillpolicy              Default
owner                    System
mountlimit              0
lowspacewarningenable   true
lowspacewarningthreshold 2
nospacewarningenable    true
mode                     normal
```

---

With **lowspacewarningthreshold** set to 2 TiB, when pool1's capacity drops below 2 TiB, a trap will be sent to the user. Example 7-55 shows the traps generated when pool1's capacity drops below 2 TiB and when migrations have failed due to pool1's capacity being depleted.

*Example 7-55 Traps sent when pool capacity is below the set threshold or migrations fail*

---

```
[root@l1tfseesrv1 ~]# l1tfsee info pools
Pool Name      Total(TiB) Free(TiB) Reclaimable(TiB) Tapes Type Library
Node Group
pool1          1.3      0.9      0.0          1 LTO l1tfsee1ib G0

2017-06-14 14:12:01 l1tfseesrv1.tuc.stglabs.ibm.com [UDP:
[9.11.121.229]:60080->[9.11.121.229]:162]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (44294822) 5 days, 3:02:28.22
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.2.6.246.1.2.22.0.609
SNMPv2-SMI::enterprises.2.6.246.1.1.3.0 = INTEGER: 40
SNMPv2-SMI::enterprises.2.6.246.1.1.1.0 = INTEGER: 7
SNMPv2-SMI::enterprises.2.6.246.1.1.4.0 = STRING: "l1tfseesrv1.tuc.stglabs.ibm.com"
SNMPv2-SMI::enterprises.2.6.246.1.1.5.0 = STRING: "GLESM609W: Pool space is going
to be small, library: anacondaq, pool: pool1, remaining: 0.9(TiB), threshold:
2(TiB)"

[root@l1tfseesrv1 ~]# l1tfsee info pools
Pool Name      Total(TiB) Free(TiB) Reclaimable(TiB) Tapes Type Library
Node Group
pool1          1.3      0.0      0.0          1 LTO l1tfsee1ib G0

[root@l1tfseesrv1 ~]# l1tfsee migrate mig -p pool1
GLESL167I(00396): A list of files to be migrated has been sent to LTFs EE using
scan id 2724109569.
GLESL038I(00442): Migration result: 0 succeeded, 10 failed, 0 duplicate, 0
duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too
early for migration.
```

```
2017-06-15 07:14:47 ltfseesrv1.tuc.stglabs.ibm.com [UDP:
[9.11.121.229]:59312->[9.11.121.229]:162]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (5959955) 16:33:19.55
SNMPv2-MIB::snmpTrapOID.0 = OID: IBMSA-MIB::ibmsaErrM613NoSpaceForMig
IBMSA-MIB::ibmsaMessageSeverity.0 = INTEGER: error(20) IBMSA-MIB::ibmsaJob.0 =
INTEGER: other(7) IBMSA-MIB::ibmsaEventNode.0 = STRING:
"ltfseesrv1.tuc.stglabs.ibm.com" IBMSA-MIB::ibmsaMessageText.0 = STRING:
"GLESM613E: There is not enough space available on the tapes belonging to storage
pool pool1 to perform the migration."
```

---

## 7.10 Migration

The migration process is the most significant reason for using IBM Spectrum Archive EE. Migration is the movement of files from IBM Spectrum Scale (on disk) to LTF5 tape cartridges in tape cartridge pools, which leaves behind a small stub file on the disk. This process has the obvious effect of reducing the usage of IBM Spectrum Scale. You can move less frequently accessed data to lower-cost, lower-tier tape storage from where it can be easily recalled.

IBM Spectrum Scale policies are used to specify files in the IBM Spectrum Scale namespace (through a GPFS scan) to be migrated to the LTF5 tape tier. For each specified GPFS file, the file content, GPFS path, and user-defined extended attributes are stored in LTF5 so that they can be re-created at import. Empty GPFS directories are not migrated.

In addition, it is possible to migrate an arbitrary list of files directly by running the **ltfsee migrate** command. This task is done by specifying the file name of a scan list file that lists the files to be migrated and specifying the designated pools as command options.

**Important:** Running the Tivoli Storage Manager for Space Management **dsmmigrate** command directly is not supported.

To migrate files, the following configuration and activation prerequisites must be met:

- ▶ Ensure that the MMM service is running on an LTF5 node. For more information, see 7.2.4, “IBM Spectrum Archive EE” on page 138.
- ▶ Ensure that storage pools that are not empty are created and defined. For more information, see 7.8.1, “Creating tape cartridge pools” on page 155.
- ▶ Ensure that space management is turned on. For more information, see 7.2.3, “Hierarchical Space Management” on page 136.
- ▶ Activate one of the following mechanisms to trigger migration:
  - Automated IBM Spectrum Scale policy-driven migration that uses thresholds.
  - Manual policy-based migration by running the **mmapplypolicy** command.
  - Manual migration by running the **ltfsee migrate** command and a prepared list of files and tape cartridge pools.

IBM Spectrum Archive EE uses a semi-sequential fill policy for tapes that enables multiple files to be written in parallel by using multiple tape drives within the tape library. Jobs are put on the queue and the scheduler looks at the queue to decide which jobs should be run. If one tape drive is available, all of the migration goes on one tape cartridge. If there are three tape drives available, the migrations are spread among the three tape drives. This configuration improves throughput and is a more efficient usage of tape drives.

With the recent changes to IBM Spectrum Archive EE V1.2.2.0, individual files no longer get scheduled on the job queue. IBM Spectrum Archive EE now internally groups files into file lists and schedules these lists on the job queue. The lists are then distributed to available drives to perform the migrations.

The grouping is done by using two parameters: A total file size and a total number of files. The default settings for the file lists are 20 GB or 20,000 files. This requirement means that a file list can contain either 20 GB of files or 20,000 number of files, whichever fills up first, before creating a new file list. For example, if you have 10 files to migrate and each file is 10 GB in size, then when migration is kicked off, IBM Spectrum Archive EE internally generates five file lists containing two files each because the two files reach the 20 GB limit that a file list can have. It then schedules those file lists to the job queue for available drives.

For performance references, see 3.4.3, “Performance” on page 44.

**Note:** Starting with IBM Spectrum Archive EE V1.2.2.0, files that are recently created need to wait two minutes before being migrated. Otherwise, the migrations will fail.

Example 7-56 shows the output of running the `mmapplypolicy` command that uses a policy file called `sample_policy.txt`.

*Example 7-56 Output of the mmapplypolicy command*

```
[root@ltfsm11 ~]# mmapplypolicy /ibm/gpfs/production/ -P sample_policy.txt
[I] GPFS Current Data Pool Utilization in KB and %
Pool_Name          KB_Occupied      KB_Total Percent_Occupied
system             861676032      12884901888    6.687486172%
[I] 6314402 of 104857600 inodes used: 6.021883%.
[I] Loaded policy rules from sample_policy.txt.
Evaluating policy rules with CURRENT_TIMESTAMP = 2016-02-12@21:56:04 UTC
Parsed 3 policy rules.

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'COPY_POOL'
EXEC '/opt/ibm/ltfsee/bin/ltfsee'
OPTS '-p copy_ltfsm11@lib_ltfsm11'

RULE 'LTFS_EE_FILES' MIGRATE FROM POOL 'system' TO POOL 'COPY_POOL'
WHERE FILE_SIZE > 0
AND ((NOT MISC_ATTRIBUTES LIKE '%M%') OR (MISC_ATTRIBUTES LIKE '%M%' AND
MISC_ATTRIBUTES NOT LIKE '%V%'))
AND NOT ((PATH_NAME LIKE '/ibm/gpfs/.ltfsee/%' OR PATH_NAME LIKE '%/.SpaceMan/%'))
[I] 2016-02-12@21:56:05.015 Directory entries scanned: 4.
[I] Directories scan: 3 files, 1 directories, 0 other objects, 0 'skipped' files
and/or errors.
[I] 2016-02-12@21:56:05.019 Sorting 4 file list records.
[I] Inodes scan: 3 files, 1 directories, 0 other objects, 0 'skipped' files and/or
errors.
[I] 2016-02-12@21:56:05.040 Policy evaluation. 4 files scanned.
[I] 2016-02-12@21:56:05.044 Sorting 3 candidate file list records.
[I] 2016-02-12@21:56:05.045 Choosing candidate files. 3 records scanned.
[I] Summary of Rule Applicability and File Choices:
  Rule#   Hit_Cnt   KB_Hit   Chosen KB_Chosen   KB_IllRule
      0         3     46392     3     46392     ORULE 'LTFS_EE_FILES'
MIGRATE FROM POOL 'system' TO POOL 'COPY_POOL' WHERE(.)
```

[I] Filesystem objects with no applicable rules: 1.

[I] GPFS Policy Decisions and File Choice Totals:

Chose to migrate 46392KB: 3 of 3 candidates;

Predicted Data Pool Utilization in KB and %:

Pool_Name	KB_Occupied	KB_Total	Percent_Occupied
system	861629640	12884901888	6.687126122%

GLESL167I(00509): A list of files to be migrated has been sent to LTFS EE using scan id 1425344513.

GLESL038I(00555): Migration result: 3 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for migration.

[I] 2016-02-12@21:57:55.981 Policy execution. 3 files dispatched.

[I] A total of 3 files have been migrated, deleted or processed by an EXTERNAL EXEC/script;

0 'skipped' files and/or errors.

---

## 7.10.1 Managing file migration pools

A file can be migrated to one pool or to multiple pools if replicas are configured. However, after the file is in the migrated state, it cannot be migrated again to other tape cartridge pools before it is recalled and made resident again by using the **ltfsee repair** command. For more information about creating replicas, see 7.10.4, “Replicas and redundant copies” on page 170. Recalling the file into resident state invalidates the LTFS copy from the reconcile and export perspective.

## 7.10.2 Threshold-based migration

This section describes how to use IBM Spectrum Scale policies for threshold-based migrations with IBM Spectrum Archive EE.

Automated IBM Spectrum Scale policy-driven migration is a standard IBM Spectrum Scale migration procedure that allows file migration from IBM Spectrum Scale disk pools to external pools. IBM Spectrum Archive EE is configured as an external pool to IBM Spectrum Scale by using policy statements.

After you define an external tape cartridge pool, migrations or deletion rules can refer to that pool as a source or target tape cartridge pool. When the **mmapplypolicy** command is run and a rule dictates that data should be moved to an external pool, the user-provided program that is identified with the EXEC clause in the policy rule starts. That program receives the following arguments:

- ▶ The command to be run. IBM Spectrum Scale supports the following subcommands:
  - LIST: Provides arbitrary lists of files with no semantics on the operation.
  - MIGRATE: Migrates files to external storage and reclaims the online space that is allocated to the file.
  - PREMIGRATE: Migrates files to external storage, but does not reclaim the online space.
  - PURGE: Deletes files from both the online file system and the external storage.
  - RECALL: Recall files from external storage to the online storage.
  - TEST: Tests for presence and operation readiness. Returns zero for success and returns nonzero if the script should not be used on a specific node.

- ▶ The name of a file that contains a list of files to be migrated.

**Important:** IBM Spectrum Archive EE supports only the MIGRATE, PREMIGRATE, and RECALL subcommands.

- ▶ Any optional parameters that are specified with the **OPTS** clause in the rule. These optional parameters are not interpreted by the IBM Spectrum Scale policy engine, but by the method that IBM Spectrum Archive EE uses to pass the tape cartridge pools to which the files are migrated.

To set up automated IBM Spectrum Scale policy-driven migration to IBM Spectrum Archive EE, you must configure IBM Spectrum Scale to be managed by IBM Spectrum Archive EE. In addition, a migration callback must be configured.

Callbacks are provided primarily as a method for system administrators to take notice when important IBM Spectrum Scale events occur. It registers a user-defined command that IBM Spectrum Scale runs when certain events occur. For example, an administrator can use the low disk event callback to inform system administrators when a file system is getting full.

The migration callback is used to register the policy engine to be run if a high threshold in a file system pool is met. For example, after your pool usage reaches 80%, you can start the migration process. You must enable the migration callback by running the **mmaddcallback** command.

In the **mmaddcallback** command in Example 7-57, the **--command** option points to a sample script file called `/usr/lpp/mmfs/bin/mmapplypolicy`. Before you run this command, you must ensure that the specified sample script file exists. The **--event** option registers the events for which the callback is configured, such as the “low disk space” events that are in the command example.

For more information about how to create and set a fail-safe policy, see 8.10, “Real world use cases for mmapplypolicy” on page 235.

*Example 7-57 A mmaddcallback example*

---

```
mmaddcallback MIGRATION --command /usr/lpp/mmfs/bin/mmapplypolicy --event  
lowDiskSpace --parms “%fsName -B 20000 -m <2x the number of drives>  
--single-instance”
```

---

For more information, see the following publications:

- ▶ *IBM Spectrum Scale: Administration and Programming Reference Guide:*  
<https://www.ibm.com/support/knowledgecenter/STXKQY>
- ▶ *Tivoli Field Guide - TSM for Space Management for UNIX-GPFS Integration* white paper:  
<https://www.ibm.com/support/docview.wss?uid=swg27028178>

After the file system is configured to be managed by IBM Spectrum Archive EE and the migration callback is configured, a policy can be set up for the file system. The placement policy that defines the initial placement of newly created files and the rules for placement of restored data must be installed into IBM Spectrum Scale by using the **mmchpolicy** command. If an IBM Spectrum Scale file system does not have a placement policy installed, all the data is stored in the system storage pool.

You can define the file management rules and install them in the file system together with the placement rules by running the **mmchpolicy** command. You also can define these rules in a separate file and explicitly provide them to the **mmapplypolicy** command by using the **-P** option. The latter option is described in 7.10.3, “Manual migration” on page 166.

In either case, policy rules for placement or migration can be intermixed. Over the life of the file, data can be migrated to a different tape cartridge pool any number of times, and files can be deleted or restored.

The policy must define IBM Spectrum Archive EE (/opt/ibm/ltfsee/bin/ltfsee) as an external tape cartridge pool.

**Tip:** Only one IBM Spectrum Scale policy, which can include one or more rules, can be set up for a particular GPFS file system.

After a policy is entered into a text file (such as `policy.txt`), you can apply the policy to the file system by running the **mmchpolicy** command. You can check the syntax of the policy before you apply it by running the command with the **-I test** option, as shown in Example 7-58.

*Example 7-58 Test an IBM Spectrum Scale policy*

---

```
mmchpolicy /dev/gpfs policy.txt -t "System policy for LTFS EE" -I test
```

---

After you test your policy, run the **mmchpolicy** command without the **-I test** to set the policy. After a policy is set for the file system, you can check the policy by displaying it with the **mmispolicy** command, as shown in Example 7-59. This policy migrates all files in groups of 20 GiB after the IBM Spectrum Archive disk space reaches a threshold of or above 80% in the `/ibm/glues/archive` directory to tape.

*Example 7-59 List an IBM Spectrum Scale policy*

---

```
[root@l1tfs97]# mmispolicy /dev/gpfs -L
/* LTFS EE - GPFS policy file */

define(
    user_exclude_list,
    PATH_NAME LIKE '/ibm/glues/0%'
    OR NAME LIKE '%&%' )

define(
    user_include_list,
    FALSE )

define(
    exclude_list,
    NAME LIKE 'dsmerror.log')

/* define is_premigrated uses GPFS inode attributes that mark a file
   as a premigrated file. Use the define to include or exclude premigrated
   files from the policy scan result explicitly */
define(
    is_premigrated,
    MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE '%V%')
```

```

/* define is_migrated uses GPFS inode attributes that mark a file
   as a migrated file. Use the define to include or exclude migrated
   files from the policy scan result explicitly */
define(
    is_migrated,
    MISC_ATTRIBUTES LIKE '%V%')

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'Archive_files'
EXEC '/opt/ibm/ltfsee/bin/ltfsee'
OPTS -p 'pool1'
SIZE(20971520)

RULE 'ARCHIVE_FILES' MIGRATE FROM POOL 'system'
THRESHOLD(80,50)
TO POOL 'Archive_files'
WHERE PATH_NAME LIKE '/ibm/glues/archive/%'
AND NOT (exclude_list)
AND (NOT (user_exclude_list) OR (user_include_list))
AND (is_migrated OR is_premigrated)

```

---

To ensure that a specified IBM Spectrum Scale is migrated only once, run the `mmapplypolicy` command with the `--single-instance` option. If this is not done, IBM Spectrum Archive EE attempts to start another migration process every two minutes. It is called every two minutes because the threshold continues to be exceeded while migrations are still in the process of migrating the data to free up the space.

As a preferred practice, the user should not use overlapping IBM Spectrum Scale policy rules within different IBM Spectrum Scale policy files that select the same files for migration to different tape cartridge pools. If a file is already migrated, later migration attempts fail, which is the standard HSM behavior. However, this is not normally done, due to incorporating thresholds.

**Important:** If a single IBM Spectrum Scale file system is used and the metadata directory is stored in the same file system that is space-managed with IBM Spectrum Archive EE, migration of the metadata directory must be prevented. The name of metadata directory is `/ibm/gpfs/.ltfsee/meta/`.

By combining the attributes of `THRESHOLD` and `WEIGHT` in IBM Spectrum Scale policies, you can have a great deal of control over the migration process. When an IBM Spectrum Scale policy is applied, each candidate file is assigned a weight (based on the `WEIGHT` attribute). All candidate files are sorted by weight and the highest weight files are chosen to `MIGRATE` until the low occupancy percentage (based on the `THRESHOLD` attribute) is achieved, or there are no more candidate files.

Example 7-60 on page 165 shows a policy that starts migration of all files when the file system pool named “system” reaches 80% full (see the `THRESHOLD` attribute), and continues migration until the pool is reduced to 60% full or less by using a weight that is based on the date and time that the file was last accessed (refer to the `ACCESS_TIME` attribute). The file system usage is checked every two minutes.

All files to be migrated must have more than 5 MB of disk space that is allocated for the file (see the `KB_ALLOCATED` attribute). The migration is performed to an external pool, presented by IBM Spectrum Archive EE (`/opt/ibm/ltfsee/bin/ltfsee`), and the data that is migrated is sent to the IBM Spectrum Archive EE tape cartridge pool named `Tapepool1`. In addition, this example policy excludes some system files and directories.

*Example 7-60 Threshold-based migration in an IBM Spectrum Scale policy file*

---

```
define
(
  exclude_list,
  (
    PATH_NAME LIKE '%/.SpaceMan/%'
    OR PATH_NAME LIKE '%/.ctdb/%'
    OR PATH_NAME LIKE '/ibm/glues/.ltfsee/%'
    OR NAME LIKE 'fileset.quota%'
    OR NAME LIKE 'group.quota%'
  )
)
RULE EXTERNAL POOL 'ltfsee'
EXEC '/opt/ibm/ltfsee/bin/ltfsee'
OPTS -p'Tapepool1' /* This is our pool in LTFS Enterprise Edition */
SIZE(20971520)

/* The following statement is the migration rule */
RULE 'ee_sysmig' MIGRATE FROM POOL 'system'

THRESHOLD(80,60)
WEIGHT(CURRENT_TIMESTAMP - ACCESS_TIME)
TO POOL 'ltfsee'
WHERE (KB_ALLOCATED > 5120)
AND NOT (exclude_list)

/* The following statement is the default placement rule that is required for a
system migration */
RULE 'default' set pool 'system'
```

---

In addition to monitoring the file system's overall usage in Example 7-60, you can monitor how frequently a file is accessed with IBM Spectrum Scale policies. A file's *access temperature* is an attribute for a policy that provides a means of optimizing tiered storage. File temperatures are a relative attribute, which indicates whether a file is “hotter” or “colder” than the others in its pool.

The policy can be used to migrate hotter files to higher tiers and colder files to lower. The access temperature is an exponential moving average of the accesses to the file. As files are accessed, the temperature increases. Likewise, when the access stops, the file cools. File temperature is intended to optimize nonvolatile storage, not memory usage. Therefore, cache hits are not counted. In a similar manner, only user accesses are counted.

The access counts to a file are tracked as an exponential moving average. A file that is not accessed loses a percentage of its accesses each period. The loss percentage and period are set through the configuration variables `fileHeatLossPercent` and `fileHeatPeriodMinutes`. By default, the file access temperature is not tracked.

To use access temperature in policy, the tracking must first be enabled. To do this, set the following configuration variables:

► **fileHeatLossPercent**

The percentage (0 - 100) of file access temperature that is dissipated over the fileHeatPeriodMinutes time. The default value is 10.

► **fileHeatPeriodMinutes**

The number of minutes that is defined for the recalculation of file access temperature. To turn on tracking, **fileHeatPeriodMinutes** must be set to a nonzero value from the default value of 0. You use `WEIGHT(FILE_HEAT)` with a policy **MIGRATE** rule to prioritize migration by file temperature.

The following example sets fileHeatPeriodMinutes to 1440 (24 hours) and fileHeatLossPercent to 10, meaning that unaccessed files lose 10% of their heat value every 24 hours, or approximately 0.4% every hour (because the loss is continuous and “compounded” geometrically):

```
mmchconfig fileheatperiodminutes=1440,fileheatlosspercent=10
```

**Note:** If the updating of the file access time (atime) is suppressed or if relative atime semantics are in effect, proper calculation of the file access temperature might be adversely affected.

These examples provide only an introduction to the wide range of file attributes that migration can use in IBM Spectrum Scale policies. IBM Spectrum Scale provides a range of other policy rule statements and attributes to customize your IBM Spectrum Scale environment, but a full description of all these is outside the scope for this publication.

For syntax definitions for IBM Spectrum Scale policy rules, which correspond to constructs in this script (such as EXEC, EXTERNAL POOL, FROM POOL, MIGRATE, RULE, OPTS, THRESHOLD, TO POOL, WEIGHT, and WHERE), see the information about policy rule syntax definitions in *IBM Spectrum Scale: Administration Guide* at this website:

<https://www.ibm.com/support/knowledgecenter/STXKQY>

Also, see 8.4, “Setting mmapplypolicy options for increased performance” on page 229.

For more information about IBM Spectrum Scale SQL expressions for policy rules, which correspond to constructs in this script (such as CURRENT\_TIMESTAMP, FILE\_SIZE, MISC\_ATTRIBUTES, NAME, and PATH\_NAME), see the information about SQL expressions for policy rules in *IBM Spectrum Scale: Administration Guide* at this website:

<https://www.ibm.com/support/knowledgecenter/STXKQY>

### 7.10.3 Manual migration

In contrast to the threshold-based migration process that can be controlled only from within IBM Spectrum Scale, the manual migration of files from IBM Spectrum Scale to LTFs tape cartridges can be accomplished by running the **mmapplypolicy** command or the **lftsee** command. The use of these commands is documented in this section. Manual migration is more likely to be used for *ad hoc* migration of a file or group of files that do not fall within the standard IBM Spectrum Scale policy that is defined for the file system.

#### Using mmapplypolicy

This section describes how to manually start file migration while using an IBM Spectrum Scale policy file for file selection.

You can apply a manually created policy by manually running the `mmapplypolicy` command, or by scheduling the policy with the system scheduler. You can have multiple different policies, which can each include one or more rules. However, only one policy can be run at a time.

**Important:** Prevent migration of the `.SPACEMAN` directory of a GPFS file system by excluding the directory with an IBM Spectrum Scale policy rule.

You can accomplish manual file migration for an IBM Spectrum Scale file system that is managed by IBM Spectrum Archive EE by running the `mmapplypolicy` command. This command runs a policy that selects files according to certain criteria, and then passes these files to IBM Spectrum Archive EE for migration. As with automated IBM Spectrum Scale policy-driven migrations, the name of the target IBM Spectrum Archive EE tape cartridge pool is provided as the first option of the pool definition rule in the IBM Spectrum Scale policy file.

The following phases occur when the `mmapplypolicy` command is started:

1. Phase 1: Selecting candidate files

In this phase of the `mmapplypolicy` job, all files within the specified GPFS file system device (or below the input path name) are scanned. The attributes of each file are read from the file's GPFS inode structure.

2. Phase two: Choosing and scheduling files

In this phase of the `mmapplypolicy` job, some or all of the candidate files are chosen. Chosen files are scheduled for migration, accounting for the weights and thresholds that are determined in phase one.

3. Phase three: Migrating and premigrating files

In the third phase of the `mmapplypolicy` job, the candidate files that were chosen and scheduled by the second phase are migrated or premigrated, each according to its applicable rule.

For more information about the `mmapplypolicy` command and other information about IBM Spectrum Scale policy rules, see the *IBM Spectrum Scale: Administration Guide*:

<https://www.ibm.com/support/knowledgecenter/STXKQY>

**Important:** In a multicluster environment, the scope of the `mmapplypolicy` command is limited to the nodes in the cluster that owns the file system.

### **Hints and tips**

Before you write and apply policies, consider the following points:

- ▶ Always test your rules by running the `mmapplypolicy` command with the `-I test` option and the `-L 3` (or higher) option before they are applied in a production environment. This step helps you understand which files are selected as candidates and which candidates are chosen.
- ▶ To view all selected files that have been chosen for migration, run the `mmapplypolicy` command with the `-I defer` and the `-f /tmp` options. The `-I defer` option runs the actual policy without actually making any data movements, and the `-f /tmp` option specifies a directory or file to output each migration rule to. This option is helpful when dealing with lots of files.
- ▶ Do not apply a policy to an entire file system of vital files until you are confident that the rules correctly express your intentions. To test your rules, find or create a subdirectory with a modest number of files, some that you expect to be selected by your SQL policy rules and some that you expect are skipped.

Run the following command:

```
mmapplypolicy /ibm/gpfs/TestSubdirectory -L 6 -I test
```

The output shows you exactly which files are scanned and which ones match rules.

### **Testing an IBM Spectrum Scale policy**

Example 7-61 shows a `mmapplypolicy` command that tests, but does not apply, an IBM Spectrum Scale policy by using the `testpolicy` policy file.

#### *Example 7-61 Test an IBM Spectrum Scale policy*

---

```
[root@l1tfsml1 ~]# mmapplypolicy /ibm/gpfs/production/ -P sample_policy.txt -I test
[I] GPFS Current Data Pool Utilization in KB and %
Pool_Name          KB_Occupied      KB_Total Percent_Occupied
system             861650688       12884901888    6.687289476%
[I] 6314406 of 104857600 inodes used: 6.021887%.
[I] Loaded policy rules from sample_policy.txt.
Evaluating policy rules with CURRENT_TIMESTAMP = 2016-02-13@16:17:57 UTC
Parsed 3 policy rules.

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system' RULE EXTERNAL POOL 'COPY_POOL'
EXEC '/opt/ibm/l1tfsml1/bin/l1tfsml1'
OPTS '-p copy_l1tfsml1@lib_l1tfsml1'
SIZE(20971520)

RULE 'LTFS_EE_FILES' MIGRATE FROM POOL 'system'
THRESHOLD(50,0)
TO POOL 'COPY_POOL'
WHERE FILE_SIZE > 5242880
AND NAME LIKE '*.IMG'
AND ((NOT MISC_ATTRIBUTES LIKE '%M%') OR (MISC_ATTRIBUTES LIKE '%M%' AND
MISC_ATTRIBUTES NOT LIKE '%V%'))
AND NOT ((PATH_NAME LIKE '/ibm/gpfs/.l1tfsml1/%' OR PATH_NAME LIKE '%/.SpaceMan/%'))
[I] 2016-02-13@16:17:58.356 Directory entries scanned: 5.
[I] Directories scan: 4 files, 1 directories, 0 other objects, 0 'skipped' files
and/or errors.
[I] 2016-02-13@16:17:58.360 Sorting 5 file list records.
[I] Inodes scan: 4 files, 1 directories, 0 other objects, 0 'skipped' files and/or
errors.
[I] 2016-02-13@16:17:58.376 Policy evaluation. 5 files scanned.
[I] 2016-02-13@16:17:58.380 Sorting 1 candidate file list records.
[I] 2016-02-13@16:17:58.380 Choosing candidate files. 1 records scanned.
[I] Summary of Rule Applicability and File Choices:
  Rule#   Hit_Cnt   KB_Hit   Chosen KB_Chosen   KB_I11Rule
        0         1    20736     1    20736      ORULE 'LTFS_EE_FILES'
MIGRATE FROM POOL 'system' TO POOL 'COPY_POOL' WHERE(.)

[I] Filesystem objects with no applicable rules: 4.[I] GPFS Policy Decisions and
File Choice Totals:
  Chose to migrate 20736KB: 1 of 1 candidates;
Predicted Data Pool Utilization in KB and %:
Pool_Name          KB_Occupied      KB_Total Percent_Occupied
system             861629952       12884901888    6.687128544%
```

---

The policy in Example 7-61 on page 168 is configured to select files that have the file extension `.IMG` for migration to the IBM Spectrum Archive EE tape cartridge pool named `copy_ltfsm11` in library name `lib_ltfsm11` if the usage of the `/ibm/glues` file system exceeds 50% for any `.IMG` file that exceeds 5 MB.

## Using ltfsee

The `ltfsee migrate` command requires a migration list file that contains a list of files to be migrated with the name of the target tape cartridge pool. Unlike migrating files by using IBM Spectrum Scale policy, it is not possible to use wildcards in place of file names. The name and path of each file to be migrated must be specified in full. The file must be in the following format:

- ▶ `-- /ibm/glues/file1.mpeg`
- ▶ `-- /ibm/glues/file2.mpeg`

**Note:** Make sure that there is a space before and after the “--”.

Example 7-62 shows the output of running such a `migrate` command.

*Example 7-62 Manual migration by using a scan result file*

---

```
[root@lufs97 /]# ltfsee migrate gpfs-scan.txt -p MPEGpool
GLES167I(00400): A list of files to be migrated has been sent to LTFS EE using
scan id 1842482689.
GLES1038I(00448): Migration result: 2 succeeded, 0 failed, 0 duplicate, 0
duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too
early for migration.
```

---

## Using a cron job

Migrations that use the `ltfsee` and `mmapplypolicy` commands can be automated by scheduling cron jobs, which is possible by setting a cron job that periodically triggers migrations by calling `mmapplypolicy` with `ltfsee` as an external program. In this case, the full path to `ltfsee` must be specified. The following are steps to start the crond process and create a cron job:

1. Start the crond process from by running `/etc/rc.d/init.d/crond start` or `/etc/init.d/crond start`.
2. Create a crontab job by opening the crontab editor with the `crontab -e` command. If using VIM to edit the jobs, press `i` to enter insert mode to start typing.
3. Enter the frequency and command that you would like to run.
4. After entering the jobs you would like to run, exit the editor. If using VIM, press the Escape key and enter `:wq`. If using nano, press **Ctrl** + **x**. This combination opens the save options. Then, press **y** to save the file and then **Enter** to override the file name.
5. View that the cron job has been created by running `crontab -l`.

The syntax for a cron job is `m h dom mon dow command`. In this syntax, `m` stands for minutes, `h` stands for hours, `dom` stands for day of month, `mon` stands for month, and `dow` stands for day of week. The hour parameter is in a 24-hour period, so `0` represents midnight and `12` represents noon.

Example 7-63 shows how to start the **crond** process and create a single cron job that performs migrations every six hours.

*Example 7-63 Creating a cron job for migrations to run every 6 hours*

---

```
[root@l1tfseesrv ~]# /etc/rc.d/init.d/crond start
Starting crond: [ OK ]

[root@l1tfseesrv ~]# crontab -e
00 0,6,12,18 * * * /usr/lpp/mmfs/bin/mmapplypolicy gpfs -P
/root/premigration_policy.txt -B 20000 -m 16
crontab: installing new crontab

[root@l1tfseesrv ~]# crontab -l
00 0,6,12,18 * * * /usr/lpp/mmfs/bin/mmapplypolicy gpfs -P
/root/premigration_policy.txt -B 20000 -m 16
```

---

## Thresholds

By default, recall jobs have a higher priority than migration jobs. A threshold parameter is available as an option to the **l1tfsee** command. It determines the percentage usage of the IBM Spectrum Scale at which migrations are preferred over recalls. The default value is 95%.

If this value is passed for one of the managed file systems, migrations are given a higher priority. Recalls are again preferred after the file system usage drops by 5%. For example, if a threshold of 93% is selected, recalls again are preferred when the file system usage is at or below 88%. In most environments, you do not need to change this setting. Example 7-64 shows an example of reducing the threshold from 95% to 80%.

*Example 7-64 Threshold*

---

```
[root@l1tfs97 ~]# l1tfsee threshold
Current threshold: 95%

[root@l1tfs97 ~]# l1tfsee threshold 80

[root@l1tfs97 ~]# l1tfsee threshold
Current threshold: 80%
```

---

## 7.10.4 Replicas and redundant copies

This section introduces how replicas and redundant copies are used with IBM Spectrum Archive EE and describes how to create replicas of migrated files during the migration process.

### Overview

IBM Spectrum Archive EE enables the creation of a replica of each IBM Spectrum Scale file during the migration process. The purpose of the replica function is to enable creating multiple LTFS copies of each GPFS file during migration that can be used for disaster recovery, including across two tape libraries at two different locations.

The first replica is the primary copy, and more replicas are called redundant copies. Redundant copies must be created in tape cartridge pools that are different from the pool of the primary replica and from the pools of other redundant copies. Up to two redundant copies can be created (for a total of three copies of the file on various tapes).

The tape cartridge where the primary replica is stored and the tape cartridges that contain the redundant copies are referenced in the GPFS inode with an IBM Spectrum Archive EE DMAPI attribute. The primary replica is always listed first.

For transparent recalls such as double-clicks of a file or through application reads, IBM Spectrum Archive EE always performs the recall by using the primary copy tape. The primary copy is the first tape cartridge pool that is defined by the migration process. If the primary copy tape cannot be accessed, including recall failures, then IBM Spectrum Archive EE automatically tries the recall job again by using the remaining replicas if they are available during the initial migration process. This automatic retry operation is transparent to the transparent recall requester.

For selective recalls initiated by the `ltfsee recall` command, an available copy is selected from the available replicates and the recall job is generated against the selected tape cartridge. There are no retries. The selection is based on the available copies in the tape library, which is supplied by the `-l` option in a two-tape library environment.

When a migrated file is recalled for a write operation or truncated, the file is marked as resident and the pointers to tape are dereferenced. The remaining copies are no longer referenced and are removed during the reconciliation process. In the case of a truncate to 0 operation, it does not generate a recall from tape. The truncated 0 file is marked as resident only.

Redundant copies are written to their corresponding tape cartridges in the IBM Spectrum Archive EE format. These tape cartridges can be reconciled, exported, reclaimed, or imported by using the same commands and procedures that are used for standard migration without replica creation.

## Creating replicas and redundant copies

You can create replicas and redundant copies during automated IBM Spectrum Scale policy-based migrations or during manual migrations by running the `ltfsee migrate` (or `ltfsee premigrate`) command.

If an IBM Spectrum Scale scan is used and you use a scan policy file to specify files for migration, you must modify the OPTS line of the policy file to specify the tape cartridge pool for the primary replica and different tape cartridge pools for each redundant copy. The tape cartridge pool for the primary replica (including primary library) is listed first, followed by the tape cartridge pools for each copy (including a secondary library), as shown in Example 7-65.

A pool cannot be listed more than once in the OPTS line. If a pool is listed more than once per line, the file is not migrated. Example 7-65 shows the OPTS line in a policy file, which makes replicas of files in two tape cartridge pools in a single tape library.

*Example 7-65 Extract from IBM Spectrum Scale policy file for replicas*

---

```
OPTS '-p PrimPool@PrimLib CopyPool@PrimLib'
```

---

For more information about IBM Spectrum Scale policy files, see 7.10.2, “Threshold-based migration” on page 161.

If you are running the `ltfsee migrate` (or `ltfsee premigrate`) command, a scan list file must be passed along with the designated pools that the user wants to migrate the files to. This process can be done by calling `ltfsee migrate -s <scan list> -p <pool1> <pool2> <pool3>` (or `ltfsee premigrate`).

Example 7-66 shows what the scan list looks like when selecting files to migrate.

*Example 7-66 Example scan list file*

---

```
[root@ltfs97 /]# cat migrate.txt
-- /ibm/glues/document10.txt
-- /ibm/glues/document20.txt
```

---

Example 7-67 shows how one would run a manual migration using the **ltfsee migrate** command on the scan list from Example 7-66 to two tapes.

*Example 7-67 Creation of replicas during migration*

---

```
[root@ltfs97]# ltfsee migrate mig -p primary_ltfs@lib_lto copy_ltfs@lib_lto
GLES167I(00400): A list of files to be migrated has been sent to LTFS EE using
scan id 3190476033.
GLES1038I(00448): Migration result: 2 succeeded, 0 failed, 0 duplicate, 0
duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too
early for migration.
```

---

IBM Spectrum Archive EE attempts to create redundant copies as efficiently as possible with a minimum number of mount and unmount steps. For example, if all tape drives are loaded with tape cartridges that belong only to the primary copy tape cartridge pool, data is written to them before IBM Spectrum Archive EE begins loading the tape cartridges that belong to the redundant copy tape cartridge pools. For more information, see 3.4.1, “IBM Spectrum Archive EE metadata file system” on page 43.

By monitoring the **ltfsee info jobs** command as the migration is running, you can observe the status of the migration and migration (copy) jobs changing, as shown in Example 7-68.

*Example 7-68 Job status during migration*

---

```
[root@ltfseesrv ~]# ltfsee info jobs
Job Type  Status      Idle(sec)  Scan ID   Tape      Pool      Library   Node  File
Name or inode
Migration In-progress    20  852297473 IM1178L6 primary lto_ts4500  1
/ibm/gpfs/LTFS_EE_FILE_Fa43WzQKcAeP2_j2XT9t.bin
Migration In-progress    20  852297473 IM1178L6 primary lto_ts4500  1
/ibm/gpfs/LTFS_EE_FILE_JmkgSoqPX9FhRoD1gTdCnDttA3Ee0A24pfQfbeva2ruRvTp_mXYbKW.bin
Migration In-progress    20  852297473 IM1178L6 primary lto_ts4500  1
/ibm/gpfs/LTFS_EE_FILE_1Y08uc068bAzXBgPOTSKWHhU21Sv_LFFDAd.bin

[root@ltfseesrv ~]# ltfsee info jobs
Job Type  Status      Idle(sec)  Scan ID   Tape      Pool      Library   Node  File
Name or inode
Migration Copied         20  852297473 IM1178L6 primary lto_ts4500  1
/ibm/gpfs/LTFS_EE_FILE_Fa43WzQKcA eP2_j2XT9t.bin
Migration Copied         20  852297473 IM1178L6 primary lto_ts4500  1
/ibm/gpfs/LTFS_EE_FILE_JmkgSoqPX9 FhRoD1gTdCnDttA3Ee0A24pfQfbeva2ruRvTp_mXYbKW.bin
Migration Copied         20  852297473 IM1178L6 primary lto_ts4500  1
/ibm/gpfs/LTFS_EE_FILE_1Y08uc068b AzXBgPOTSKWHhU21Sv_LFFDAd.bin
```

---

For more information and command syntax, see the **ltfsee migrate** command in 7.10, “Migration” on page 159.

## Considerations

Consider the following points when replicas are used:

- ▶ Redundant copies must be created in different tape cartridge pools. The pool of the primary replica must be different from the pool for the first redundant copy, which, in turn, must be different from the pool for the second redundant copy.
- ▶ The migration of a premigrated file does not create replicas.

If offsite tapes are required, redundant copies can be exported out of the tape library and shipped to an offsite location after running the export offline. A second option would be to create the redundant copy in a different tape library.

## 7.10.5 Data Migration

With the release of Spectrum Archive EE v1.2.6, users have a more intuitive way of migrating their data from disk to tape, and between tape generations. This release also introduces new tape states (such as *appendable* and *used space*) and new pool attributes (such as *mediarestriction* and *mode*) to better control the flow of data.

The tape state attributes can be viewed by running the **ltfsee info tapes** command. The used space column shows how much space is being used on the tape. The unref(GiB) column has been renamed to reclaimable(GiB) to make it more understandable to the user by showing how much space is being used by unreferenced data.

The appendable column in the state status is dependent on multiple factors. An appendable tape means the tape is in a healthy state and is empty, or has data on it but is not full, and meets the pool attribute settings.

Tapes that are not appendable can be caused by multiple factors as well. Some factors include erroneous tapes that do not permit write functionality, full tapes, or pool settings that restrict migrations to specific tape format or generation type. Example 7-69 shows how the pool attribute, *mediarestriction*, and *format* can be used to isolate data flow to newer tape generations inside a pool with a mixed generation.

*Example 7-69 Controlling appendable tapes within a pool*

```
[root@ueno ~]# ltfsee info tapes
Tape ID  Status      Type      Capacity(GiB)  Used(GiB)  Free(GiB)
Reclaimable(GiB) Pool      Library  Address  Drive  Appendable
JCA393JC Valid      JC(E08)    6292      0        0
0 pool1   perf_lib  1058      -        no
P2B143JD Valid      JD(55F)   13562     0        13562
0 pool1   perf_lib  1050      -        yes

[root@ueno ~]# ltfsee pool show -p pool1
Attribute      Value
poolname       pool1
poolid         4b8324ca-c347-4153-92be-0500a3f9214f
devtype        3592
mediarestriction JD
format         55F (0x56)
worm           no (0)
nodegroup      G0
fillpolicy     Default
owner          System
mountlimit     0
```

lowspacewarningenable	yes
lowspacewarningthreshold	0
nospacewarningenable	yes
mode	normal

---

Because of the need to upgrade tape generations, **ltfsee datamigrate** is introduced into IBM Spectrum Archive EE v1.2.6. This command allows users to move data to a newer tape media generation in a different pool within the same nodegroup, or pool-to-pool migration within the same nodegroup. This is ideal to use when newer tape generations are being introduced into the user's environment and the older generations are no longer needed.

Data on older generation media can be moved as a whole within a pool or specific tapes can be chosen in chunks. This command uses pool-to-pool reclamation to move data from one pool to another. Example 7-70 shows how users can use the **ltfsee datamigrate** command to migrate their data from an older generation pool to a newer one.

*Example 7-70 Migrating data to newer generation tape pool*

---

```
[root@ltfssrv1 ~]# ltfsee datamigrate -p JC_pool -d JD_pool
```

Start the datamigrate command by reclaim operation.

Processing 2 tapes in the following list of tapes from source pool JC\_pool.

The files in the tapes are moved to the tapes in target pool JD\_pool:

JCA384JC .

JCA393JC .

Files in tape JCA384JC are copied to tape JD2406JD.

Tape JCA384JC successfully reclaimed as datamigrate process, it remains in storage pool JC\_pool.

Files in tape JCA393JC are copied to tape JD2406JD.

Tape JCA393JC successfully reclaimed as datamigrate process, it remains in storage pool JC\_pool.

Reclamation complete as datamigrate command. 2 tapes reclaimed, 0 tapes removed from the storage pool.

---

For the **ltfsee datamigrate** command syntax, see “The ltfsee datamigrate command” on page 300.

**Note:** If the **-t** option is not used to select specific tapes for data migration the tapes being reclaimed remain in the source pool until the user manually repurposes them. If the **-t** option is used on specific tapes, those tapes will be removed from the source pool after the reclamation completes.

In addition to the pool-to-pool data migration, users can perform in-pool data migration by configuring the pool settings to modify the `mediarestriction` or `format` type to make older generation media to become non-appendable. By doing so, this enables users to run reclamation on those older generation tape cartridges and have assurance that the data is reclaimed onto the new generation medias.

When changing the `mediarestriction` attribute of the pool, the `format` type is also automatically updated to the highest gen drive available to the cluster. The `format` attribute is automatically updated only after each time the `mediarestriction` is modified. If new drives are added to a cluster with `mediarestriction` already set, users are expected to update `format` or `mediarestriction` manually.

Example 7-71 on page 175 shows the automatic format update when changing the `mediarestriction` attribute from JC to JD when there are both TS1150 and TS1155 drives.

### Example 7-71 Updating pool mediarestriction

---

```
[root@ueno ~]# ltfsee pool show -p pool1
Attribute          Value
poolname           pool1
poolid             4b8324ca-c347-4153-92be-0500a3f9214f
devtype            3592
mediarestriction   JC
format             E08 (0x55)
worm               no (0)
nodegroup          G0
fillpolicy         Default
owner              System
mountlimit         0
lowspacewarningenable yes
lowspacewarningthreshold 0
nospacewarningenable yes
mode               normal

[root@ueno ~]# ltfsee pool set -p pool1 -a mediarestriction -v JD

[root@ueno ~]# ltfsee pool show -p pool1
Attribute          Value
poolname           pool1
poolid             4b8324ca-c347-4153-92be-0500a3f9214f
devtype            3592
mediarestriction   JD
format             55F (0x56)
worm               no (0)
nodegroup          G0
fillpolicy         Default
owner              System
mountlimit         0
lowspacewarningenable yes
lowspacewarningthreshold 0
nospacewarningenable yes
mode               normal
[root@ueno ~]#
```

---

All tapes within pool1 that fall under the restriction of JD media and 55F format type are all appendable tapes, and everything else is now non-appendable. The **ltfsee reclaim** command is used to perform the in pool migration. To migrate data off of older generation or format type media within a pool and onto newer generations or format type media, use the **reclaim** command with the **-t** option to specify the source tapes. The **reclaim** will move the data only to appendable tapes set by the mediarestriction or the format type of the pool.

## 7.10.6 Migration hints and tips

This section provides preferred practices for successfully managing the migration of files.

### Overlapping IBM Spectrum Scale policy rules

After a file is migrated to a tape cartridge pool and is in the migrated state, it cannot be migrated to other tape cartridge pools (unless it is first recalled).

It is preferable that you do not use overlapping IBM Spectrum Scale policy rules within different IBM Spectrum Scale policy files that can select the same files for migration to different tape cartridge pools. If a file is already migrated, a later migration fails.

In this example, an attempt is made to migrate four files to tape cartridge pool CopyPool. Before the migration attempt, tape 055AGWL5, which is defined in a different tape cartridge pool (PrimPool), already contains three of the four files. The state of the files on these tape cartridges before the migration attempt is displayed by the `ltfsee info` command as shown in Example 7-72.

*Example 7-72 Before migration*

---

```
[root@ltfs97 gpfs]# ltfsee info files *.ppt
Name: fileA.ppt
Tape id:2MA260L5@lib_lto Status: migrated
Name: fileB.ppt
Tape id:2MA260L5@lib_lto Status: migrated
Name: fileC.ppt
Tape id:2MA260L5@lib_lto Status: migrated
Name: fileD.ppt
Tape id:-          Status: resident
```

---

The attempt to migrate the files to a different tape cartridge pool produces the results that are shown in Example 7-73.

*Example 7-73 Attempted migration of already migrated files*

---

```
[root@ltfs97 gpfs]# ltfsee migrate mig -p CopyPool@lib_lto
GLES167I(00400): A list of files to be migrated has been sent to LTFS EE using
scan id 3353727489.
GLES159E(00440): Not all migration has been successful.
GLES1038I(00448): Migration result: 1 succeeded, 3 failed, 0 duplicate, 0
duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too
early for migration
```

---

If the IBM Spectrum Archive EE log is viewed, the error messages that are shown in Example 7-74 explain the reason for the failures.

*Example 7-74 Migration errors reported in the IBM Spectrum Archive EE log file*

---

```
2016-12-13T21:17:41.554526-07:00 ltfs97 mmm[3743]: GLESM148E(00538): File
/ibm/gpfs/fileA.ppt is already migrated and will be skipped.
2016-12-13T21:17:41.555037-07:00 ltfs97 mmm[3743]: GLESM148E(00538): File
/ibm/gpfs/fileB.ppt is already migrated and will be skipped.
2016-12-13T21:17:41.555533-07:00 ltfs97 mmm[3743]: GLESM148E(00538): File
/ibm/gpfs/fileC.ppt is already migrated and will be skipped.
```

---

The files on tape 2MA260L5 (fileA.ppt, fileB.ppt, and fileC.ppt) are already in storage pool PrimPool. Therefore, the attempt to migrate them to storage pool CopyPool produces a migration result of Failed. Only the attempt to migrate the resident file fileD.ppt succeeds.

If the aim of this migration was to make redundant replicas of the four PPT files in the CopyPool tape cartridge pool, the method that is described in 7.10.4, “Replicas and redundant copies” on page 170 must be followed instead.

## IBM Spectrum Scale policy for the .SPACEMAN directory

Prevent migration of the .SPACEMAN directory of an IBM Spectrum Scale by excluding the directory with an IBM Spectrum Scale policy rule. An example is shown in Example 7-60 on page 165.

## Automated IBM Spectrum Scale policy-driven migration

To ensure that a specified GPFS file system is migrated only once, run the `mmapplypolicy` command with the `--single-instance` option. The `--single-instance` option ensures that multiple `mmapplypolicy` commands are not running in parallel because it can take longer than two minutes to migrate a list of files to tape cartridges.

## Tape format

For more information about the format of tapes that are created by the migration process, see 11.2, “Formats for IBM Spectrum Scale to IBM Spectrum Archive EE migration” on page 326.

## Migration Policy

A migration policy is used to make your lives easier. When run, IBM Spectrum Scale performs a scan of all candidate files in the IBM Spectrum Archive name space to be migrated onto tape. This process saves the user lots of time because they do not need to manually search their file system and find candidate files for migrations. This feature is especially important when there are millions of files created. For use cases on migration policy, see 8.10, “Real world use cases for `mmapplypolicy`” on page 235.

# 7.11 Premigration

A premigrated file is a file that is on both disk and tape. To change a file to a premigrated state, you have two options:

- ▶ Recalling migrated files:
  - a. The file initially is only on a disk (the file state is *resident*).
  - b. The file is migrated to tape by running `1tfsee migrate`. After this migration, the file is a stub on the disk (the file state is *migrated*) and the IDs of the tapes containing the redundant copies are written to an IBM Spectrum Archive EE DMAPI attribute.
  - c. The file is recalled from tape by using a recall for read when a client attempts to read from the file. The file is on both disk and tape (the file state is *premigrated*).
- ▶ Premigrating files:
  - a. The file initially is only on disk (the file state is *resident*).
  - b. The file is premigrated to tape by running `1tfsee premigrate`. The IDs of the tapes that contain the redundant copies are written to an IBM Spectrum Archive EE DMAPI attribute.

Premigration works similar to migration:

- ▶ The premigration scan list file has the same format as the migration scan list file.
- ▶ Up to two more redundant copies are allowed (the same as with migration).
- ▶ Manual premigration is available by running either `1tfsee premigrate` or `mmapplypolicy`.
- ▶ Automatic premigration is available by running `1tfsee premigrate` through the `mmapplypolicy/mmaddcallback` command or a cron job.
- ▶ Migration hints and tips are applicable to premigration.

For the **ltfsee migrate** command, each migrate job is achieved internally by splitting the work into three steps:

1. Writing the content of the file to tapes, including redundant copies.
2. Writing the IDs of the tapes that contain the redundant copies of the file, which are written to an IBM Spectrum Archive EE DMAPI attribute.
3. Stubbing the file on disk.

For premigration, step 3 is not performed. The omission of this step is the only difference between premigration and migration.

### 7.11.1 Premigration with the **ltfsee premigrate** command

The **ltfsee premigrate** command is used to premigrate non-empty regular files to tape. The command syntax is the same as for the **ltfsee migrate** command. The following is an example of the syntax:

```
ltfsee premigrate -s <GPFS scan list file> -p <target tape cartridge pool 1>  
<target tape cartridge pool 2> <target tape cartridge pool 3>
```

The **<GPFS scan list file>** file includes the list of non-empty regular files to be premigrated. Each line of this file must end with **-- <full path filename>**. All file system objects are saved to the specified target tape cartridge pool. Optionally, the target tape cartridge pool can be followed by up to two more tape cartridge pools (for redundant copies) separated by spaces.

### 7.11.2 Premigration running the **mmapplypolicy** command

To perform premigration by running the **mmapplypolicy** command, the **THRESHOLD** clause is used to determine the files for premigration. There is no IBM Spectrum Scale **premiigrate** command, and the default behavior is to not premigrate files.

The **THRESHOLD** clause can have the following parameters to control migration and premigration:

**THRESHOLD** (high percentage, low percentage, premigrate percentage)

If no premigrate threshold is set with the **THRESHOLD** clause or a value is set greater than or equal to the low threshold, then the **mmapplypolicy** command does not premigrate files. If the premigrate threshold is set to zero, the **mmapplypolicy** command premigrates all files.

For example, the following rule premigrates all files if the storage pool occupancy is 0 - 30%. When the storage pool occupancy is 30% or higher, files are migrated until the storage pool occupancy drops below 30%. Then, it continues by premigrating all files:

```
RULE 'premig1' MIGRATE FROM POOL 'system' THRESHOLD (0,30,0) TO POOL 'ltfs'
```

The rule in the following example takes effect when the storage pool occupancy is higher than 50%. Then, it migrates files until the storage pool occupancy is lower than 30%, after which it premigrates the remaining files:

```
RULE 'premig2' MIGRATE FROM POOL 'system' THRESHOLD (50,30,0) TO POOL 'ltfs'
```

The rule in the following example is configured so that if the storage pool occupancy is below 30%, it selects all files that are larger than 5 MB for premigration. Otherwise, when the storage pool occupancy is 30% or higher, the policy migrates files that are larger than 5 MB until the storage pool occupancy drops below 30%.

Then, it continues by premigrating all files that are larger than 5 MB:

```
RULE 'premig3' MIGRATE FROM POOL 'system' THRESHOLD (0,30,0) TO POOL 'lufs' WHERE(
AND (KB_ALLOCATED > 5120))
```

The rule in the following example is the preferred rule when performing premigrations only. It requires a callback to perform the stubbing. If the storage pools occupancy is below 100%, it selects all files larger than 5 MB for premigration. If you set the threshold to 100% the storage pools occupancy will never exceed this value, and so migrations will not be performed. In this case, a callback is needed to run the stubbing. For an example of a callback, see 8.10.2, “Creating active archive system policies” on page 236.

```
RULE 'premig4' MIGRATE FROM POOL 'system' THRESHOLD (0,100,0) TO POOL 'lufs' WHERE
(FILE_SIZE > 5242880)
```

## 7.12 Preserving file system objects on tape

Symbolic links, empty regular files, and empty directories are some file system objects that do not contain data or content. When you save these types of file system objects, you cannot use **migration** and **premigration** commands. HSM is used to move data to and from tapes, that is, for space management.

Because these file system objects do not have data, they cannot be processed by migration or premigration. A new driver (called the *save driver*) was introduced to save these file system objects to tape.

The following items (data and metadata that is associated with an object) are written and read to and from tapes:

- ▶ File data for non-empty regular files
- ▶ Path and file name for all objects
- ▶ Target symbolic name only for symbolic links
- ▶ User-defined extended attributes for all objects except symbolic links

The following items are not written and read to and from tapes:

- ▶ Timestamps
- ▶ User ID and group ID
- ▶ ACLs

To save these file system objects on tape, you have two options:

- ▶ Calling the **lufseesave** command directly with a scan list file
- ▶ An IBM Spectrum Scale policy with the **mmapplypolicy** command

### 7.12.1 Saving file system objects with the **lufsee save** command

The **lufsee save** command is used to save symbolic links, empty regular files, and empty directories to tape. The command syntax is the same as the **lufsee migrate** or **lufsee premigrate** commands. The following is the syntax of the **lufsee save** command:

```
lufsee save -s <GPFS scan list file> -p <target tape cartridge pool 1> <target
tape cartridge pool 2> <target tape cartridge pool 3>
```

The **<GPFS scan list file>** file includes the list of file system objects (symbolic links, empty regular files, and empty directories) to be saved. Each line of this file must end with **-- <full path file system object name>**.

All file system objects are saved to the specified target tape cartridge pool. Optionally, the target tape cartridge pool can be followed by up to two more tape cartridge pools (for redundant copies) separated by spaces.

**Note:** This command is not applicable for non-empty regular files.

## 7.12.2 Saving file system objects with policies

Migration and premigration cannot be used for file system objects that do not occupy space for data. To save file system objects, such as symbolic links, empty regular files, and empty directories with an IBM Spectrum Scale policy, the IBM Spectrum Scale list rule must be used.

A working policy sample of IBM Spectrum Scale list rules to save these file system objects without data to tape can be found in the `/opt/ibm/ltfsee/data/sample_save.policy` file. The only change that is required to the following sample policy file is the specification of the cartridge pool (in blue colored letters).

These three list rules can be integrated into existing IBM Spectrum Scale policies. Example 7-75 shows the sample policy.

*Example 7-75 Sample policy to save file system objects without data to tape*

---

```
/*
Sample policy rules
to save
  symbolic links,
  empty directories and
  empty regular files
*/

RULE
EXTERNAL LIST 'emptyobjects'
EXEC '/opt/ibm/ltfsee/bin/ltfseesave'
OPTS '-p sample_pool'

define(DISP_XATTR,
CASE
  WHEN XATTR($1) IS NULL
  THEN '_NULL_'
  ELSE XATTR($1)
END
)

RULE 'symoliclinks'
LIST 'emptyobjects'
DIRECTORIES_PLUS
/*
SHOW ('mode=' || SUBSTR(MODE,1,1) ||
      ' stime=' || DISP_XATTR('dmapi.IBMSTIME') ||
      ' ctime=' || VARCHAR(CHANGE_TIME) ||
      ' spath=' || DISP_XATTR('dmapi.IBMSPATH'))
*/
WHERE
( /* if the object is a symbolic link */
```

```

        MISC_ATTRIBUTES LIKE '%L%'
    )
    AND
    (
        PATH_NAME NOT LIKE '%/.SpaceMan/%'
    )
    AND
    (
        ( /* if the object has not been saved yet */
          XATTR('dmapi.IBMSTIME') IS NULL
        AND
          XATTR('dmapi.IBMSPATH') IS NULL
        )
    OR
        ( /* if the object is modified or renamed after it was saved */
          TIMESTAMP(XATTR('dmapi.IBMSTIME')) < TIMESTAMP(CHANGE_TIME)
        OR
          XATTR('dmapi.IBMSPATH') != PATH_NAME
        )
    )
)

RULE 'directories'
LIST 'emptyobjects'
DIRECTORIES_PLUS
/*
SHOW ('mode=' || SUBSTR(MODE,1,1) ||
      ' stime=' || DISP_XATTR('dmapi.IBMSTIME') ||
      ' ctime=' || VARCHAR(CHANGE_TIME) ||
      ' spath=' || DISP_XATTR('dmapi.IBMSPATH'))
*/
WHERE
( /* if the object is a directory */
  MISC_ATTRIBUTES LIKE '%D%'
)
AND
(
  PATH_NAME NOT LIKE '%/.SpaceMan'
AND
  PATH_NAME NOT LIKE '%/.SpaceMan/%'
)
AND
(
  ( /* directory's emptiness is checked in the later processing */
    /* if the object has not been saved yet */
    XATTR('dmapi.IBMSTIME') IS NULL
  AND
    XATTR('dmapi.IBMSPATH') IS NULL
  )
OR
  ( /* if the object is modified or renamed after it was saved */
    TIMESTAMP(XATTR('dmapi.IBMSTIME')) < TIMESTAMP(CHANGE_TIME)
  OR
    XATTR('dmapi.IBMSPATH') != PATH_NAME
  )
)
)
)

```

```

RULE 'emptyregularfiles'
LIST 'emptyobjects'
/*
SHOW ('mode=' || SUBSTR(MODE,1,1) ||
      ' stime=' || DISP_XATTR('dmapi.IBMSTIME') ||
      ' ctime=' || VARCHAR(CHANGE_TIME) ||
      ' spath=' || DISP_XATTR('dmapi.IBMSPATH'))
*/
WHERE
( /* if the object is a regular file */
  MISC_ATTRIBUTES LIKE '%F%'
)
AND
(
  PATH_NAME NOT LIKE '%/.SpaceMan/%'
)
AND
(
  ( /* if the size = 0 and the object has not been saved yet */
    FILE_SIZE = 0
    AND
    XATTR('dmapi.IBMSTIME') IS NULL
    AND
    XATTR('dmapi.IBMSPATH') IS NULL
  )
  OR
  ( /* if the object is modified or renamed after it was saved */
    FILE_SIZE = 0
    AND
    (
      TIMESTAMP(XATTR('dmapi.IBMSTIME')) < TIMESTAMP(CHANGE_TIME)
    OR
      XATTR('dmapi.IBMSPATH') != PATH_NAME
    )
  )
)
)
)
)
)

```

---

## 7.13 Restoring non-empty regular files and file system objects from tape

The `ltfsee rebuild` command rebuilds a GPFS file system by restoring migrated files and saved file system objects (symbolic links, empty files, and empty directories) from tapes. The migrated files and saved file system objects are restored to a specified directory by using the files that are found on the specified tapes.

If multiple versions or generations of a file are found, the latest version or generation is selected. If any of the versions or generations cannot be determined, the file that most recently was imported is renamed and two (or more) versions or generations are rebuilt or recovered.

**Note:** The use of `ltfsee rebuild` is for disaster scenarios, such as if the IBM Spectrum Scale file system is lost and a rebuild is required from tape.

The `ltfsee rebuild` command has the following syntax:

```
ltfsee rebuild -P <pathName> -p <poolName> -l <libraryName> <-t tape_id_1  
tape_id_2 ... tape_id_N>
```

<pathName> specifies the directory where the IBM Spectrum Scale file system is rebuilt to, and <-t tape\_id\_1 tape\_id\_2 ... tape\_id\_N> specifies the tapes from the pool that is specified by <poolName> from the tape library that is specified by <libraryName> to search for the files to rebuild the IBM Spectrum Scale file system.

For more information, see “Restoring file system objects and files from tape” on page 251.

## 7.14 Recall

In space management solutions, there are two different types of recall possibilities: Transparent and selective recall processing. Both are possible with the current IBM Spectrum Archive EE implementation.

Transparent recalls are initiated by an application that tries to read, write, or truncate a migrated file while not being aware that it was migrated. The specific I/O request that initiated the recall of the file is fulfilled, with a possible delay because the file data is not available immediately (it is on tape).

For transparent recalls, it is difficult to do an optimization because it is not possible to predict when the next transparent recall will happen. Some optimization is already possible because within the IBM Spectrum Archive EE job queue, the requests are run in an order that is based on the tape and the starting block to which a file is migrated. This becomes effective only if requests happen close together in time.

Furthermore, with the default IBM Spectrum Archive EE settings, there is a limitation of up to only 60 transparent recalls possible on the IBM Spectrum Archive EE job queue. A 61st request appears only if one of the previous 60 transparent recall requests completes. Therefore, the ordering can happen only on this small 60 transparent recall subset. It is up to the software application to send the transparent recalls in parallel to have multiple transparent recalls to run at the same time.

Selective recalls are initiated by users that are aware that the file data is on tape and they want to transfer it back to disk before an application accesses the data. This action avoids delays within the application that is accessing the corresponding files.

Contrary to transparent recalls, the performance objective for selective recalls is to provide the best possible throughput for the complete set of files that is being recalled, disregarding the response time for any individual file.

However, to provide reasonable response times for transparent recalls in scenarios where recall of many files is in progress, the processing of transparent recalls are modified to have higher priority than selective recalls. Selective recalls are performed differently than transparent recalls, and so they do not have a limitation like transparent recalls.

Recalls have higher priority than other IBM Spectrum Archive EE operations. For example, if there is a recall request for a file on a tape cartridge being reclaimed or for a file on the tape cartridge being used as reclamation target, the reclamation job is stopped, the recall or recalls from the tape cartridge that is needed for recall are served, and then the reclamation resumes automatically.

Recalls also have higher priority over tape premigration processes. They are optimized across tapes and optimized within a tape used for premigration activities. The recalls that are in close proximity are given priority.

### 7.14.1 Transparent recall

Transparent recall processing automatically returns migrated file data to its originating local file system when you access it. After the data is recalled by reading the file, the HSM client leaves the copy of the file in the tape cartridge pool, but changes it to a premigrated file because an identical copy exists on your local file system and in the tape cartridge pool. If you do not modify the file, it remains premigrated until it again becomes eligible for migration. A transparent recall process waits for a tape drive to become available.

If you modify or truncate a recalled file, it becomes a resident file. The next time your file system is reconciled, MMM marks the stored copy for deletion.

The order of selection from the replicas is always the same. The primary copy is always selected first from which to be recalled. If this recall from the primary copy tape fails or is not accessible, then IBM Spectrum Archive EE automatically retries the transparent recall operation against the other replicas if they exist.

**Note:** Transparent recall is used most frequently because it is activated when you access a migrated file, such as opening a file.

### 7.14.2 Selective recall

Use selective recall processing if you want to return specific migrated files to your local file system. The access time (`atime`) changes to the current time when you selectively recall a migrated file.

To selectively recall files, run the `tail` command or any similar command. For example, the command that is shown in Example 7-76 recalls a file that is named `file6.img` to the `/ibm/glues` directory.

*Example 7-76 Recall a single file*

---

```
[root@l1tfs97 glues]# tail /ibm/glues/file6.img
```

---

No message is displayed to confirm the successful recall of the file. However, if there is an error message, it is logged to the `dsmerror.log` file. The `l1tfs97 info files` command can be used to verify a successful recall. After a successful recall, the file status changes from migrated to premigrated.

### 7.14.3 Recalling files with the `ltfsee recall` command

The `ltfsee recall` command performs selective recalls of migrated files to the local file system. This command performs selective recalls in multiple ways:

- ▶ Using a recall list file
- ▶ Using an IBM Spectrum Scale scan list file
- ▶ From the output of another command
- ▶ Using an IBM Spectrum Scale scan list file that is generated through an IBM Spectrum Scale policy and the `mmapplypolicy` command

With multiple tape libraries configured, the `ltfsee recall` command requires the `-l` option to specify the tape library from which to recall. When a file is recalled, the recall can occur on any of the tapes (that is, either primary or redundant copies) from the specified tape library. The following conditions are applied to determine the best replica:

- ▶ The condition of the tape
- ▶ If a tape is mounted
- ▶ If a tape is mounting
- ▶ If there are jobs that are assigned to a tape

If conditions are equal between certain tapes, the primary tape is preferred over the redundant copy tapes. The secondary tape is preferred over the third tape. These rules are necessary to make the tape selection predictive. However, there are no automatic retries likes with transparent recalls.

For example, if a primary tape is not mounted but a redundant copy is, the redundant copy tape is used for the recall job to avoid unnecessary mount operations.

If the specified tape library does not have any replicas, IBM Spectrum Archive EE automatically resubmits the request to the other tape library to process the bulk recalls:

- ▶ Three copies: TAPE1@Library1 TAPE2@Library1 TAPE3@Library2
  - If `-l Library1` → TAPE1 or TAPE2
  - If `-l Library2` → TAPE3
- ▶ Two copies: TAPE1@Library1 TAPE2@Library1
  - If `-l Library1` → TAPE1 or TAPE2
  - If `-l Library2` → TAPE1 or TAPE2

#### The `ltfsee recall` command

The `ltfsee recall` command is used to recall non-empty regular files from tape. The command syntax is the same as the `ltfsee migrate` command. Here are some examples of the syntax:

- ▶ `ltfsee recall -l <library_name> -f <recall list file>`

The `<recall list file>` file includes a list of non-empty regular files to be recalled. Each line contains the file name with an absolute path or a relative path based on the working directory.

- ▶ `ltfsee recall -l <library_name> -s <GPFS scan list file>`

The `<GPFS scan list file>` file includes the list of non-empty regular files to be recalled. Each line of this file must end with “`-- <full path filename>`”.

## The ltfsee recall command with the output of another command

The `ltfsee recall` command can take as input the output of other commands through a pipe. In Example 7-77, all files with names ending with `.bin` are recalled under the `/ibm/gpfs/production` directory, including subdirectories. Therefore, it is convenient to recall whole directories with a simple command.

*Example 7-77 The ltfsee recall command with the output of another command*

---

```
[root@ltfsm11 ~]# find /ibm/gpfs/production -name "*.bin" -print | ltfsee recall -l lib_ltfsm11
GLES277I(00318): The ltfsee recall command is called without specifying an input file waiting for standard input.
If necessary, press ^D to exit.
GLES268I(00142): 4 file name(s) have been provided to recall.
GLES263I(00191): Recall result: 4 succeeded, 0 failed, 0 duplicate, 0 not migrated, 0 not found.
```

---

Example 7-78 shows the output of the command that is run in Example 7-77.

*Example 7-78 The ltfsee info jobs command output from the ltfsee recall command in Example 7-77*

---

```
[root@ltfsm11 ~]# ltfsee info jobs
Job Type      Status      Idle(sec)  Scan ID    Tape      Pool      Library   Node  File Name or inode
Selective Recall In-progress  2  3108835585 2FC140L5  copy_ltfsm11 lib_ltfsm11  2  10374858
Selective Recall In-progress  2  3108835585 2FC141L5  copy_ltfsm11 lib_ltfsm11  2  10374888
Selective Recall Unscheduled  2  3108835585 2FC140L5  copy_ltfsm11 lib_ltfsm11  -  10374875
Selective Recall Unscheduled  2  3108835585 2FC140L5  copy_ltfsm11 lib_ltfsm11  -  10374877
```

---

### 7.14.4 Read Starts Recalls: Early trigger for recalling a migrated file

IBM Spectrum Archive EE can define a stub size for migrated files so that the stub size initial bytes of a migrated file are kept on disk while the entire file is migrated to tape. The migrated file bytes that are kept on the disk are called the *stub*. Reading from the stub does not trigger a recall of the rest of the file. After the file is read beyond the stub, the recall is triggered. The recall might take a long time while the entire file is read from tape because a tape mount might be required, and it takes time to position the tape before data can be recalled from tape.

When Read Start Recalls (RSR) is enabled for a file, the first read from the stub file triggers a recall of the complete file in the background (asynchronous). Reads from the stubs are still possible while the rest of the file is being recalled. After the rest of the file is recalled to disks, reads from any file part are possible.

With the Preview Size (PS) value, a preview size can be set to define the initial file part size for which any reads from the resident file part does not trigger a recall. Typically, the PS value is large enough to see whether a recall of the rest of the file is required without triggering a recall for reading from every stub. This process is important to prevent unintended massive recalls. The PS value can be set only smaller than or equal to the stub size.

This feature is useful, for example, when playing migrated video files. While the initial stub size part of a video file is played, the rest of the video file can be recalled to prevent a pause when it plays beyond the stub size. You must set the stub size and preview size to be large enough to buffer the time that is required to recall the file from tape without triggering recall storms.

For more usage details, see “The ltfsee failover command” on page 303.

## 7.15 Repairing files to their resident state

This section describes the **ltfsee repair** command. However, this command should rarely be used by itself. The **recover** command in 7.16, “Handling of write-failure tapes” on page 189 internally uses this command to fix the file states when recovering from a write failure. The **ltfsee repair** command is used to repair a file or object by changing the state to Resident when the tape (or tapes) that are used for migration, premigration, or save are not available.

The **ltfsee repair** command is available only for files in the Premigrated state. This option removes metadata on IBM Spectrum Scale, which is used for keeping the file/object state. Here is an example of the syntax:

```
ltfsee repair <pathName>
```

<pathName> specifies the path name of the file to be repaired to the Resident state.

A typical usage of the **ltfsee repair** command is when a tape malfunctions and the files on the tape must be marked Resident again to allow for migration, premigration, or save. After they are migrated, premigrated, or saved again, the files are on a primary tape and redundant copy tapes.

When a tape goes to the Critical, Write Fenced, or Warning state, the number of tapes that can be used for recalls is reduced. For example, if a file is migrated to two tapes (one primary and one redundant copy tape) and one of those tapes malfunctions, the file now can only be recalled from the remaining tape (redundancy is reduced).

**Note:** To prevent further damage to tape if redundant copies exist, recalls favor valid tapes rather than the Critical, Write Fenced, or Warning tapes. In the case where there is only one copy, then recalls occur on that tape.

Use the following procedure:

1. Recall the migrated files on the malfunctioning tape by using the remaining tape.
2. Mark those recalled files as Resident.
3. Migrate the files to tape again, which regains the two copy redundancy.

After recalling files, the files are in the Premigrated state (data content is on disk and on tape). Now that the data content is on disk, you can mark the files to be in the Resident state by removing all metadata from the IBM Spectrum Scale files by running the **ltfsee repair** command. This task can be easily done by using an IBM Spectrum Scale policy that selects files that are in the Premigrated state and are on the bad tape, and then calling an external script by running the **ltfsee repair** command.

A sample IBM Spectrum Scale policy to select all premigrated files that are on the malfunctioning tape is shown in Example 7-79. In this sample policy, replace the tape VOLSER ID with your tape VOLSER ID.

*Example 7-79 Sample IBM Spectrum Scale policy*

```
/*  
 Sample policy rules to make premigrated files to resident  
*/  
  
define(is_premigrated,(MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE '%V%'))  
  
RULE
```

```

EXTERNAL LIST 'premigrated_files'
EXEC './make_resident.sh'
OPTS 'premig_to_resident.list'

RULE 'rule_premigrated_files'
LIST 'premigrated_files'
WHERE PATH_NAME NOT LIKE '%/.SpaceMan/%'
AND is_premigrated
AND XATTR('dmapi.IBMTPS') LIKE '%VOLSER ID%'
AND XATTR('dmapi.IBMTPS') LIKE ':%:%'

```

---

Example 7-80 shows the `make_resident.sh` script that is referenced in Example 7-79 on page 187. The script calls the `ltfsee repair` command.

*Example 7-80 The `make_resident.sh` script*

---

```

#!/bin/bash
# $1: TEST or LIST command
# $2: GPFS policy scan result file
# $3: File name to backup GPFS policy scan result file

rc_latch=0

# Do nothing for TEST command
if [ "$1" = "TEST" ]
then
    exit 0
elif [ "$1" != "LIST" ]
then
    echo "usage $0 <TEST|LIST>"
    exit 1
fi

#Save GPFS policy scan result if $3 is specified
if [ "$3" != "" ]
then
    cp $2 $3
fi

# Obtain premig file name from GPFS scan result file, and make it resident
cat $2 | sed -e "s/.*-- //" | while read premig_file
do
    /opt/ibm/ltfsee/bin/ltfsee repair $premig_file
    if [ $? != 0 ]
    then
        rc_latch = 1
    fi
done

exit ${rc_latch}

```

---

At this stage, all previously migrated files on the bad tape are in the Resident state. Because these files are now resident again, the currently set IBM Spectrum Scale policy or manual operation can migrate these files to tape again (allowing for redundancy on multiple tapes).

## 7.16 Handling of write-failure tapes

When a tape has a write failure, the two possible outcomes for the tape are `Critical` and `Write Fenced`. In the first outcome, the tape fails to sync and the index is not written to the tape, so the tape goes into the `Critical` state. In the second outcome, starting with IBM Spectrum Archive EE V1.2.2.0, write failures to tape might have a successful sync and the index successfully written for all the files migrated to tape right before the write failure.

This scenario makes the tape go into a `Write Fenced` state. If the tape has a write failure and the tape is in the `Write Fenced` state, the drive does not go into a `Locked` state. This state allows the `Write Fenced` tape to be unloaded, where it remains as `Write Fenced` for other IBM Spectrum Archive jobs to access that drive.

For both `Critical` and `Write Fenced` tapes, the preferred action item is to copy the data off the write failure tape and onto another tape to avoid future failures. “Recover the data” means recalling migrated files and saved objects from the tape to IBM Spectrum Scale, and copying the other files from the tape to IBM Spectrum Scale or local disk, as specified by the user.

To recover the data from the `Critical` and `Write Fenced` tapes and remove the `Critical` tape from IBM Spectrum Archive LE+ to make the tape drive available again, run the `ltfsee recover` command. The `ltfsee recover` command can be used only when the tape is in the `Critical` state or the `Write Fenced` state.

For steps to recover from a write failure tape, see 10.3, “Recovering data from a write failure tape” on page 284.

## 7.17 Handling of read-failure tapes

When a tape has a read failure, its state is changed to `Warning`. Starting with IBM Spectrum Archive EE V1.2.2, tapes that have a read failure are no longer candidates for future migration jobs, due to the poor read quality on one or more `Warning` tapes. It is important to copy the data on the `Warning` tape with the read failure over to a valid tape within the same pool to prevent further read failures or permanent damage to the `Warning` tape. After the data is copied over to a new tape, you can discard the `Warning` tape.

For steps to recover data from a read failure, see 10.4, “Recovering data from a read failure tape” on page 286.

## 7.18 Reconciliation

This section describes file reconciliation with IBM Spectrum Archive EE and presents considerations for the reconciliation process.

HSM is not notified upon moves, renames, or deletions of files in IBM Spectrum Scale. Therefore, over time the metadata of migrated files on IBM Spectrum Scale can diverge from their equivalents on LTFS. The goal of the reconciliation function is to synchronize the IBM Spectrum Scale namespace with the corresponding LTFS namespace (per tape cartridge) and the corresponding LTFS attributes (per tape cartridge).

The reconciliation process resolves any inconsistencies that develop between files in the IBM Spectrum Scale and their equivalents in IBM Spectrum Archive EE. When files are deleted, moved, or renamed in IBM Spectrum Scale, the metadata of those files becomes out of sync with their copies in LTFS.

By performing file reconciliation, it is possible to synchronize the IBM Spectrum Scale namespace and attributes that are stored in LTFS (on tape cartridges) with the current IBM Spectrum Scale namespace and attributes. However, reconciliation works on only tape cartridges that were used in IBM Spectrum Archive EE. Tapes that were not used in LTFS Library Edition (LE) cannot be reconciled.

For each file that was deleted in IBM Spectrum Scale, the reconciliation process deletes the corresponding LTFS files and symbolic links. If the parent directory of the deleted symbolic link is empty, the parent directory is also deleted. This process frees memory resources that were needed for storing the LTFS index entries of those deleted files.

For each IBM Spectrum Scale file that was moved or renamed in IBM Spectrum Scale, the reconciliation process updates for each LTFS instance (replica) of that IBM Spectrum Scale file the LTFS extended attribute that contains the IBM Spectrum Scale path and the LTFS symbolic link.

Reconciliation can be performed on one or more GPFS file systems, one or more tape cartridge pools, or a set of tape cartridges. When the reconciliation process involves multiple tape cartridges, multiple IBM Spectrum Archive EE nodes and tape drives can be used in parallel. However, because recall jobs have priority, only available tape drives are used for reconciliation. After reconciliation is started, the tape cartridge cannot be unmounted until the process completes.

The following list presents limitations of the reconciliation process:

1. Only one reconciliation process can be started at a time. If an attempt is made to start a reconciliation process while another process is running, the attempt fails. The **ltfsee reconcile** command fails and the following failure message appears:  

```
GLESLO98E(00774): Another reconciliation, reclamation or export job is currently executing. Wait for completion of the executing process and try again.
```
2. After a reconciliation process is started, new migration jobs are prevented until the reconciliation process completes on the reconciling tapes. However, if any migration jobs are running, the reconciliation process does not begin until all migration jobs complete.
3. Recalls from a tape cartridge being reconciled are not available while the reconciliation process is updating the index for that tape cartridge, which is a short step in the overall reconciliation process.

The command outputs in the following examples show the effect that reconciliation has on a file after that file is renamed in the GPFS file system. Example 7-81 shows the initial state with a single file that is called `file1.img` on tape.

*Example 7-81 List the file on LTFS tape*

---

```
[root@ltfs97 /]# ls -la /ltfs/058AGWL5/ibm/glues/file1.img*  
lrwxrwxrwx 1 root root 87 Apr  9 13:40 /ltfs/058AGWL5/ibm/glues/file1.img ->  
/ltfs/058AGWL5/.LTFSEE_DATA/1066406549503693876-17750464391302654144-190660251-134  
206-0
```

---

The file is also present on the GPFS file system, as shown in Example 7-82 on page 191.

*Example 7-82 List the file on the GPFS file system*

---

```
[root@ltfs97 /]# ls -la /ibm/glues/file1.img*  
-r--r--r-- 1 root root 137248768 Mar 27 16:28 /ibm/glues/file1.img
```

---

IBM Spectrum Archive EE considers the file to be in a Premigrated state (Example 7-83).

*Example 7-83 List the file in IBM Spectrum Archive EE*

---

```
[root@ltfs97 /]# ltfsee info files /ibm/glues/file1.img*  
File name: /ibm/glues/file1.img  
Tape id:058AGWL5@lib_ltfsml1 Status: premigrated
```

---

The file is renamed to .old, as shown in Example 7-84.

*Example 7-84 Rename the file*

---

```
[root@ltfs97 /]# mv /ibm/glues/file1.img /ibm/glues/file1.img.old  
  
[root@ltfs97 /]# ls -la /ibm/glues/file1.img*  
-r--r--r-- 1 root root 137248768 Mar 27 16:28 /ibm/glues/file1.img.old
```

---

However, the file on the tape cartridge is not immediately and automatically renamed as a result of the previous change to the file in IBM Spectrum Scale. This rename can be confirmed by running the command that is shown in Example 7-85, which still shows the original file name.

*Example 7-85 List the file on the LTFS tape cartridge*

---

```
[root@ltfs97 /]# ls -la /ltfs/058AGWL5/ibm/glues/file1.img*  
lrwxrwxrwx 1 root root 87 Apr 9 13:40 /ltfs/058AGWL5/ibm/glues/file1.img ->  
/ltfs/058AGWL5/.LTFSEE_DATA/1066406549503693876-17750464391302654144-190660251-134  
206-0
```

---

If you perform a reconciliation of the tape now, IBM Spectrum Archive EE synchronizes the file in IBM Spectrum Scale with the file on tape, as shown in Example 7-86.

*Example 7-86 Reconcile the LTFS tape*

---

```
[root@ltfs97 /]# ltfsee reconcile -t 058AGWL5 -p copy_ltfsml1 -l lib_ltfsml1  
GLESS016I(00109): Reconciliation requested  
GLESS049I(00610): Tapes to reconcile: 058AGWL5  
GLESS050I(00619): GPFS filesystems involved: /ibm/glues  
GLESS053I(00647): Number of pending migrations: 0  
GLESS054I(00651): Creating GPFS snapshots:  
GLESS055I(00656): Creating GPFS snapshot for /ibm/glues ( /dev/gpfs )  
GLESS056I(00724): Scanning GPFS snapshots:  
GLESS057I(00728): Scanning GPFS snapshot of /ibm/glues ( /dev/gpfs )  
GLESS058I(00738): Removing GPFS snapshots:  
GLESS059I(00742): Removing GPFS snapshot of /ibm/glues ( /dev/gpfs )  
GLESS060I(00760): Processing scan results:  
GLESS061I(00764): Processing scan results for /ibm/glues ( /dev/gpfs )  
GLESS063I(00789): Reconciling the tapes:  
GLESS001I(00815): Reconciling tape 058AGWL5 has been requested  
GLESS002I(00835): Reconciling tape 058AGWL5 complete  
GLES172I(02984): Synchronizing LTFS EE tapes information
```

---

If you list the files on the tape cartridge, you can see that the file name changed. Compare Example 7-87 with the output from Example 7-85 on page 191. Although files are moved or renamed, recalls on those files are fine.

*Example 7-87 List the files on the LTFS tape cartridge*

---

```
[root@ltfs97 /]# ls -la /ltfs/058AGWL5/ibm/glues/file1.img*
lrwxrwxrwx 1 root root 87 Apr  9 13:44 /ltfs/058AGWL5/ibm/glues/file1.img.old ->
/ltfs/058AGWL5/.LTFSEE_DATA/1066406549503693876-17750464391302654144-190660251-134
206-0
```

---

## 7.19 Reclamation

The space on tape that is occupied by deleted files is not reused during normal IBM Spectrum Archive EE operations. New data is always written after the last index on tape. The process of reclamation is similar to the same named process in IBM Spectrum Protect (formerly Tivoli Storage Manager) environment because all active files are consolidated onto a new, empty, second tape cartridge. This process improves overall tape usage and utilization.

When files are deleted, overwritten, or edited on IBM Spectrum Archive EE tape cartridges, it is possible to reclaim the space. The reclamation function of IBM Spectrum Archive EE frees tape space that is occupied by non-referenced files and non-referenced content that is present on the tape. The reclamation process copies the files that are referenced by the LTFS index of the tape cartridge being reclaimed to another tape cartridge, updates the GPFS/IBM Spectrum Scale inode information, and then reformats the tape cartridge that is being reclaimed.

### 7.19.1 Reclamation considerations

The following considerations should be reviewed before the reclamation function is used:

- ▶ Reconcile before reclaiming tape cartridges

It is preferable to perform a reconciliation of the set of tape cartridges that are being reclaimed before the reclamation process is initiated. For more information, see 7.15, “Repairing files to their resident state” on page 187. If this is not performed, the reclamation might fail with the following message:

```
GLSL086I(01990): Reclamation has not completed since at least tape 058AGWL5
needs to be reconciled.
```

- ▶ Scheduled reclamation

It is preferable to schedule periodically reclamation for the IBM Spectrum Archive EE tape cartridge pools.

- ▶ Recall priority

Recalls are prioritized over reclamation. If there is a recall request for a file on a tape cartridge that is being reclaimed or for a file on the tape cartridge being used as the reclamation target, the reclamation job is stopped for the recall. After the recall is complete, the reclamation resumes automatically.

- ▶ One tape cartridge at a time

Only one tape cartridge is reclaimed at a time. The reclamation function does not support parallel use of drives for reclaiming multiple tape cartridges simultaneously.

Use the `reclaim` option of the `ltfsee` command to start reclamation of a specified tape cartridge pool or of certain tape cartridges within a specified tape cartridge pool. The `ltfsee reclaim` command is also used to specify thresholds that indicate when reclamation is performed by the percentage of the available capacity on a tape cartridge.

Example 7-88 shows the results of reclaiming a single tape cartridge 058AGWL5.

*Example 7-88 Reclamation of a single tape cartridge*

---

```
[root@ltfs97 glues]# ltfsee reclaim -p myfirstpool -l myfirstlib -t 058AGWL5
Start reclaiming the following 1 tapes:
058AGWL5
Tape 058AGWL5 successfully reclaimed, formatted, and removed from storage pool
myfirstpool.
Reclamation complete. 1 tapes reclaimed, 1 tapes removed from the storage pool.
```

---

At the end of the process, the tape cartridge is reformatted and removed from the tape cartridge pool only if the `-t` or `-n` options are used. For more information, see “The `ltfsee premigrate` command” on page 316.

## 7.19.2 Reclamation performance

Reclaim operation might take a long time to copy all required files from one tape to another, especially when the number of files to be copied is large. In a lab demonstration, it was observed that reclamation of a tape that has only 2,000 of 1 MiB files that are left takes almost 2,000 seconds. As a rough estimate, 2.5 TB of LTO tape has 2,500,000 of 1 MiB files. Assuming 10% of them remain in the tape when a user requests reclaim, it takes 250,000 seconds, that is, around 70 hours.

Preparing this time window to fit into the schedule of production system by using IBM Spectrum Archive EE is difficult for common use cases. From this point of view, improvement of the performance of the execution time of the reclaim operation was introduced with IBM Spectrum Archive V1R1.1.3.

A new command-line option was introduced to limit the number of files to be moved within a reclaim operation. With this option, a user can manage the length of time of a reclaim operation to fit the specified time window in the production system operations.

Another new function called *quick reconcile* was added so that a user does not need to run `reconcile` before the reclamation task. The quick reconcile operation is performed implicitly at the start of the reclaim operation, and the reclaim operation now can run without a reconcile operation before it.

**Note:** Even if the quick reconcile function is introduced, a full reconcile might be required in some rare cases before running the reclaim operation.

The chart that is shown in Figure 7-2 indicates how many seconds are required to reclaim all files in the source tape. With prior LTFS EE versions, this process takes around 1 second per file. Starting with IBM Spectrum Archive EE V1R1.1.3 (PGA2.2), it is almost 10 times faster when the source and the target tape cartridge are on the same node.

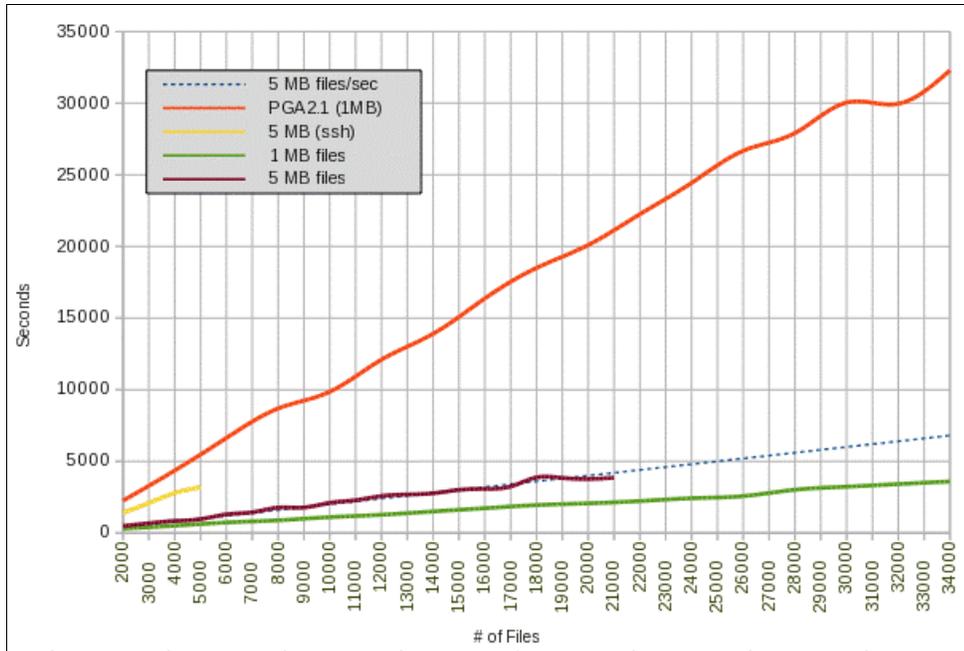


Figure 7-2 Reclamation performance enhancement by using “quick reconcile”

The yellow line shows performance just for a reference when `scp` and `ssh rm` commands are used on the same node by IBM Spectrum Archive EE. When the source and the target tapes are on a different node, the performance is worse than this reference because it requires metadata access to get the list of files in the source tape. In this case for the quick reconcile operation, the total time of reclamation might become longer than the case without the quick reconcile.

The time required to have the quick reconcile varies by the number of files to be reconciled. The quick reconcile includes removing files that are not required anymore from an IBM Spectrum Archive EE tape.

New messages are added to log so that a user can see the progress of the operation. Some examples are shown in Example 7-89.

*Example 7-89 New log message for reclamation enhancement*

```

GLESR051I(01439): 17608 of 22000 files have been processed.
GLESR052I(01505): Removing files from source tape TY0118L5.
GLESR053I(01528): 1000 of 22000 files have been removed.

```

When the option to limit the number of files for reclamation is used and one or more files remain in the source tape, the following message is shown in the log file:

```

GLESL332I(1tfsee:4404): Reclamation has been partially performed. Run the
reclamation again for tape JCA811JC.

```

Two option switches were added to the **reclaim** command starting with IBM Spectrum Archive EE V1R1.1.3 and later:

- ▶ **-L**: Limit the number of files to be processed by the **reclaim** command.

This **-L** option can be set with the **-t** option when the **-t** option has only one `tape_id`. When the number of files that are stored on the tape to be reclaimed is more than the value provided with the **-L** option, the **reclaim** command completes its operation when the number of files that are provided by the **-L** option is processed. You might not gain any capacity from unreferenced capacity at this point because another reclaim operation might be required.

However, this option is useful to limit the time of operation taken by one **reclaim** command. For example, when there is only a limited time window available for the **reclaim** operation, a user can limit the number of files by using the **-L** option with a small number. After the completion of the command with the **-L** option, the user has the choice to run the command again if there is time available.

When the number of files that are stored on the tape is less than the number provided by the **-L** option, the **reclaim** operation completes as though the **reclaim** option was run without the **-L** option. If this option is specified without any value, the default value is 100,000.

- ▶ **-q**: Quick Reconcile.

This option enables the quick reconcile feature of the **reclaim** command. The quick reconcile feature handles basic inconsistencies between the files on LTFSS and IBM Spectrum Scale. The basic file inconsistencies have the following conditions:

- A file on IBM Spectrum Scale was removed.
- A file on IBM Spectrum Scale was renamed.

When a file with these conditions is found, the **reclaim** operation without the **-q** option does not touch those files and reports to the user that the **reconcile** command must be run. With the **-q** option, those files are processed without running the **reconcile** operation in advance.

When there is a file that cannot be processed by the quick reconcile feature, the **reclaim** command continues the operation for the other files and reports at the end of operation that a reconcile is required. Only the files that are not handled by the **reclaim** command are left on the tape.

To use the quick reconcile operation, all of the DMAPI enabled file systems should be mounted. Otherwise, the **reclaim** operation fails. Files that are imported from outside of IBM Spectrum Archive EE by the **import** command are not handled by the quick reconcile feature.

The **-q** option can be used in any combination of the options, but it should be put after the **-t** or **-n** option when the **-t** or **-n** option is used.

## 7.20 Checking and repairing

You can run the **1tfsee pool add** command with a check option or a deep recovery option to check the medium when one or more tape cartridges in the library are inconsistent and it becomes necessary to perform a check and recover operation.

Example 7-90 shows the output of the check and recovery on a single tape cartridge.

*Example 7-90 Check and recovery of a tape*

---

```
[root@l1fs97 glues]# ltfsee pool add -p CopyPool -t 058AGWL5 -c
Tape 058AGWL5 successfully checked.
Adding tape 058AGWL5 to storage pool CopyPool
```

---

Most corrupted tapes can be repaired by using the **--check** (or **-c**) option. If the command fails, it might be because the tape is missing an end-of-data (EOD) mark. Try to repair the tape again by running the **ltfsee pool add** command with the **--deep\_recovery** (or **-d**) option, as shown in Example 7-91.

*Example 7-91 Deep recovery of a tape*

---

```
[root@l1fs97 glues]# ltfsee pool add -p CopyPool -t 058AGWL5 -d
Tape 058AGWL5 successfully checked.
Adding tape 058AGWL5 to storage pool CopyPool
```

---

## 7.21 Importing and exporting

The import and export processes are the mechanisms for moving existing data on the LTFS written tape cartridges into or out of the IBM Spectrum Archive EE environment.

### 7.21.1 Importing

Import tape cartridges to your IBM Spectrum Archive EE system by running the **ltfsee import -p <pool> -t <tape>** command.

When you import a tape cartridge, the **ltfsee import -p <pool> -t <tape>** command performs the following actions:

1. Adds the specified tape cartridge to the IBM Spectrum Archive EE library
2. Assigns the tape to the designated pool

Imports the files on that tape cartridge into the IBM Spectrum Scale namespace process puts the stub file back in the GPFS file system and the imported files are in a migrated state, which means that the data remains on tape cartridge. The data portion of the file is not copied to disk during the import.

Example 7-92 shows the import of an LTFS tape cartridge that was created on a different LTFS system into a directory that is called 075AGWL5 in the `/ibm/glues` file system.

*Example 7-92 Import an LTFS tape cartridge*

---

```
[root@l1fs97 /]# ltfsee import -p myPool -t 075AGWL5 -P /ibm/glues
Import of tape 075AGWL5 has been requested...
Import of tape 075AGWL5 complete.
```

---

#### Importing file paths

The default import file path for the **ltfsee import** command is `/{GPFS file system}/IMPORT`. As shown in Example 7-93 on page 197, if no other parameters are specified on the command line, all files are restored to the `../IMPORT/{VOLSER}` directory under the GPFS file system.

*Example 7-93 Import by using default parameters*

---

```
[root@ltfs97 glues]# ltfsee import -p myPool -t 037AGWL5
Import of tape 037AGWL5 has been requested...
Import of tape 037AGWL5 complete.
```

```
[root@ltfs97 glues]# ls -las /ibm/glues/IMPORT/037AGWL5
total 0
0 drwxr-xr-x 2 root root      512 Apr 18 16:22 .
0 drwxr-xr-x 3 root root      512 Apr 18 16:22 ..
0 -rw----- 1 root root 104857600 Apr 18 16:22 file10.img
0 -rw----- 1 root root 104857600 Apr 18 16:22 file9.img
0 -rw----- 1 root root 104857600 Apr 18 16:22 fileA.ppt
0 -rw----- 1 root root 104857600 Apr 18 16:22 fileB.ppt
0 -rw----- 1 root root 104857600 Apr 18 16:22 fileC.ppt
0 -rw----- 1 root root 104857600 Apr 18 16:22 offsite1.mpeg
0 -rw----- 1 root root 104857600 Apr 18 16:22 offsite2.mpeg
0 -rw----- 1 root root 104857600 Apr 18 16:22 offsite3.mpeg
0 -rw----- 1 root root 104857600 Apr 18 16:22 offsite4.mpeg
0 -rw----- 1 root root 104857600 Apr 18 16:22 offsite5.mpeg
```

---

Example 7-94 shows the use of the `-P` parameter, which can be used to redirect the imported files to an alternative directory. The VOLSER is still used in the directory name, but you can now specify a custom import file path by using the `-P` option. If the specified path does not exist, it is created.

*Example 7-94 Import by using the -P parameter*

---

```
[root@ltfs97 glues]# ltfsee import -p myPool -t 037AGWL5 -P /ibm/glues/alternate
Import of tape 037AGWL5 has been requested...
Import of tape 037AGWL5 complete.
```

```
[root@ltfs97 glues]# ls -las /ibm/glues/alternate/037AGWL5
total 32
 0 drwxr-xr-x  2 root root      512 Apr 18 16:24 .
32 drwxr-xr-x 11 root root    32768 Apr 18 16:24 ..
 0 -rw-----  1 root root 104857600 Apr 18 16:24 file10.img
 0 -rw-----  1 root root 104857600 Apr 18 16:24 file9.img
 0 -rw-----  1 root root 104857600 Apr 18 16:24 fileA.ppt
 0 -rw-----  1 root root 104857600 Apr 18 16:24 fileB.ppt
 0 -rw-----  1 root root 104857600 Apr 18 16:24 fileC.ppt
 0 -rw-----  1 root root 104857600 Apr 18 16:24 offsite1.mpeg
 0 -rw-----  1 root root 104857600 Apr 18 16:24 offsite2.mpeg
 0 -rw-----  1 root root 104857600 Apr 18 16:24 offsite3.mpeg
 0 -rw-----  1 root root 104857600 Apr 18 16:24 offsite4.mpeg
 0 -rw-----  1 root root 104857600 Apr 18 16:24 offsite5.mpeg
```

---

Example 7-95 shows the use of the `-R` parameter during the import, which has the effect of importing files to the root of the file system that is specified and not creating a VOLSER directory.

*Example 7-95 Import by using the -R parameter*

---

```
[root@ltfs97 glues]# ltfsee import -p myPool -t 037AGWL5 -P /ibm/glues -R
Import of tape 075AGWL5 has been requested...
Import of tape 075AGWL5 complete.
```

```
[root@ltfs97 glues]# ls -las /ibm/glues
total 1225525
    32 drwxr-xr-x 10 root root    32768 Apr 18 18:03 .
     4 drwxr-xr-x  4 root root     4096 Apr  3 15:00 ..
102400 -rw-----  1 root root  104857600 Apr 18 18:02 file10.img
102400 -rw-----  1 root root  104857600 Apr 16 14:51 file1.img
102400 -rw-----  1 root root  104857600 Apr 16 14:51 file3.img
102400 -rw-----  1 root root  104857600 Apr 18 18:02 file9.img
102400 -rw-----  1 root root  104857600 Apr 18 18:02 fileA.ppt
102400 -rw-----  1 root root  104857600 Apr 18 18:02 fileB.ppt
102400 -rw-----  1 root root  104857600 Apr 18 18:02 fileC.ppt
102400 -rw-----  1 root root  104857600 Apr 18 18:03 offsite1.mpeg
102400 -rw-----  1 root root  104857600 Apr 18 18:03 offsite2.mpeg
102400 -rw-----  1 root root  104857600 Apr 18 18:03 offsite3.mpeg
100864 -rw-----  1 root root  104857600 Apr 18 18:03 offsite4.mpeg
100608 -rw-----  1 root root  104857600 Apr 18 18:03 offsite5.mpeg
     1 dr-xr-xr-x  2 root root     8192 Apr 18 17:58 .snapshots
     6 drwxrwsr-x  6 bin  bin     8192 Apr 18 17:58 .SpaceMan
```

---

With each of these parameters, you have the option of overwriting, ignoring, or renaming existing files by using the **-o**, **-i**, or **-r** parameters.

## Handling import file name conflicts

This section describes how you can use the **rename**, **overwrite**, and **ignore** options of the **ltfsee import** command to handle import file name conflicts. The default behavior is to rename a file being imported if that file name exists in the full target import path. When you import files by running the **ltfsee import** command, you can use the following **ltfsee import** command options to affect command processing when the file name of an import file exists in the full target path:

► **-r, --rename**

This is the default setting. All existing files are kept and any import files with conflicting names are renamed. Files are renamed by appending the suffix **\_i**, where **i** is a number 1 - n. For example, a file that is named **file1.txt** is renamed **files1.txt\_1**.

► **-o, --overwrite**

All import files are imported and any existing files with conflicting names are overwritten.

► **-i, --ignore**

All existing files are kept and any import files with conflicting names are ignored.

## Importing offline tape cartridges

For more information about offline tape cartridges, see 7.21.2, “Exporting” on page 199. Offline tape cartridges can be reimported to the IBM Spectrum Scale namespace by running the **ltfsee import** command with the **-offline** option. If an offline exported tape cartridge was modified by the user while it was outside the library, the **ltfsee import -offline** command fails and the user must reimport the tape cartridge again by using another **ltfsee import** option.

When the tape cartridge is offline and outside the library, the IBM Spectrum Scale offline files on disk or the files on tape cartridge should not be modified.

Problems that are caused by trying to import a tape cartridge that was exported by using the option `--offline` can be solved by reimporting the tape cartridge by using the other options available. By using `import --recreate` and `-overwrite` (assuming the original path is used), some of the inodes on IBM Spectrum Scale are overwritten and new inodes are created for the files that did not have one. At the end of this process, all the files on the tape cartridge have an inode on IBM Spectrum Scale.

Example 7-96 shows an example of importing an offline tape cartridge.

*Example 7-96 Import an offline tape cartridge*

---

```
[root@l1tfssn1 ~]# l1tfsee import -p JZJ5W0RM -t JZ0072JZ --offline
Import of tape JZ0072JZ has been requested.
Import of tape JZ0072JZ complete.
Updated offline state of tape JZ0072JZ.
```

---

## 7.21.2 Exporting

Export tape cartridges from your IBM Spectrum Archive EE system by running the `l1tfsee export` command. When you normal export a tape cartridge, the `l1tfsee export` command removes the tape cartridge from the IBM Spectrum Archive EE library. The tape cartridge is reserved so that it is no longer a target for file migrations. It is then reconciled to remove any inconsistencies between it and IBM Spectrum Scale.

The export process then removes all files from the IBM Spectrum Scale file system that exist on the exported tape cartridge. The files on tape cartridges are unchanged by the export, and are accessible by other LTFS systems.

If the `--offline` option is specified, all files from the tape cartridges or tape cartridge pool that are specified are moved to an offline status and those files cannot be accessed. However, the corresponding inode of each file is kept in IBM Spectrum Scale. Those files can be brought back to the IBM Spectrum Scale namespace by reimporting the tape cartridge by using the `--offline` option.

This `--offline` option can be used, for example, when exporting tape cartridges containing redundant copies to an off-site storage location for disaster recovery purposes.

**Important:** If the `--offline` option is omitted in the `export` command, all files on the exported tape cartridge are removed from the GPFS file system.

### Export considerations

Consider the following information when planning IBM Spectrum Archive EE export activities:

- ▶ If you put different logical parts of an IBM Spectrum Scale namespace (such as the project directory) into different LTFS tape cartridge pools, you can export tape cartridges that contain the entire IBM Spectrum Scale namespace or only the files from a specific directory within the namespace.

Otherwise, you must first recall all the files from the namespace of interest (such as the project directory), then migrate the recalled files to an empty tape cartridge pool, and then export that tape cartridge pool.

- ▶ Reconcile occurs automatically before the export is processed.

Although the practice is not preferable, tape cartridges can be physically removed from IBM Spectrum Archive EE without exporting them. In this case, no changes are made to the IBM Spectrum Scale inode. The following results can occur:

- ▶ Causes a file operation that requires access to the removed tape cartridge to fail. No information as to where the tape cartridges are is available.
- ▶ Files on an LTFS tape cartridge can be replaced in IBM Spectrum Archive EE without reimporting (that is, without updating anything in IBM Spectrum Scale). This process is equivalent to a library going offline and then being brought back online without taking any action in the IBM Spectrum Scale namespace or management.

**Important:** If a tape cartridge is removed from the library without the use of the export utility, modified, and then reinserted in the library, the behavior can be unpredictable.

## Exporting tape cartridges

The normal export of an IBM Spectrum Archive EE tape cartridge first reconciles the tape cartridge to correct any inconsistencies between it and IBM Spectrum Scale. Then, it removes all files from the IBM Spectrum Scale file system that exist on the exported tape cartridge.

Example 7-97 shows the typical output from the **export** command.

*Example 7-97 Export a tape cartridge*

---

```
# ltfs export -p PrimPool1 lib_lto -t 2MA262L5
GLESS016I(00184): Reconciliation requested.
GLESM401I(00194): Loaded the global configuration.
GLESS049I(00637): Tapes to reconcile: 2MA262L5 .
GLESS050I(00644): GPFS file systems involved: /ibm/gpfs .
GLESS134I(00666): Reserving tapes for reconciliation.
GLESS135I(00699): Reserved tapes: 2MA262L5 .
GLESS054I(00737): Creating GPFS snapshots:
GLESS055I(00742): Deleting the previous reconcile snapshot and creating a new one
for /ibm/gpfs ( gpfs ).
GLESS056I(00763): Scanning GPFS snapshots:
GLESS057I(00768): Scanning GPFS snapshot of /ibm/gpfs ( gpfs ).
GLESS060I(00844): Processing scan results:
GLESS061I(00849): Processing scan results for /ibm/gpfs ( gpfs ).
GLESS141I(00862): Removing stale DMAPI attributes:
GLESS142I(00867): Removing stale DMAPI attributes for /ibm/gpfs ( gpfs ).
GLESS063I(00900): Reconciling the tapes:
GLESS001I(00994): Reconciling tape 2MA262L5 has been requested.
GLESS002I(01013): Reconciling tape 2MA262L5 complete.
GLESS137I(01134): Removing tape reservations.
GLESS058I(02320): Removing GPFS snapshots:
GLESS059I(02327): Removing GPFS snapshot of /ibm/gpfs ( gpfs ).
Export of tape 2MA262L5 has been requested...
GLESLO74I(00649): Export of tape 2MA262L5 complete.
GLESLO373I(00805): Moving tape 2MA262L5.
Tape 2MA262L5 is unmounted because it is inserted into the drive.
GLESLO43I(00151): Removing tape 2MA262L5 from storage pool PrimPool.
GLESLO490I(00179): Export command completed successfully.
```

---

Example 7-98 shows how the normal exported tape is displayed as exported by running the `ltfsee info tapes` command.

*Example 7-98 Display status of normal export tape cartridge*

---

# ltfsee info tapes											
Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)	Pool	Library	Address	Drive	Appendable
11A105L5	Valid	LT05	1327	0	1327		0 PrimPool	lib_lto 4141	-	yes	
D00397L5	Valid	LT05	1327	0	1327		0 PrimPool	lib_lto 4143	-	yes	
D00454L5	Valid	LT05	1327	0	1327		0 PrimPool	lib_lto 4136	-	yes	
2MA262L5	Exported	LT05	0	0	0		0	lib_lto 4138	-	-	

---

If errors occur during the export phase, the tape goes to the export state. However, some of the files that belong to that tape might still remain in the file system and still have a reference to that tape. Such an error can occur when an export is happening and while reconciliation occurs one starts to modify the files belonging to the exporting tape. In such a scenario, see 10.5, "Handling export errors" on page 287 on how to clean up the remaining files on the IBM Spectrum Scale file system.

**Exporting offline tape cartridges**

If you want to move tape cartridges to an off-site location for DR purposes but still retain files in the IBM Spectrum Scale file system, follow the procedure that is described here. In Example 7-99, tape 051AGWL5 contains redundant copies of five MPEG files that must be moved off-site.

*Example 7-99 Export an offline tape cartridge*

```
[root@ltfssn1 ~]# ltfsee export -p JZJ5WORM -t JZ0072JZ -o "Export offline"
```

```
Export of tape JZ0072JZ has been requested...
GLESLO74I(00649): Export of tape JZ0072JZ complete.
Updated offline state of tape JZ0072JZ to OFFLINE.
GLESL487I(00139): Tape JZ0072JZ stays in pool JZJ5WORM while it is offline
exported.
GLESL373I(00805): Moving tape JZ0072JZ.
Tape JZ0072JZ is unmounted because it is inserted into the drive.
```

---

Offline export is a much quicker export function than normal export because no reconciliation is run. If you run the `ltfsee info tapes` command, you can see the offline status of the tape cartridge as shown in Example 7-100.

*Example 7-100 Display status of offline tape cartridges*

---

```
[root@ltfssn1 ~]# ltfsee info tapes
```

Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)	Pool	Library	Address	Drive	Appendable
JYA825JY	Valid	TS1140(J5)	6292	0	6292	0	JY	TS4500_3592	1035	-	yes
JCC541JC	Unknown	TS1140	0	0	0	0	J5	TS4500_3592		0	-
JZ0072JZ	Offline	TS1150(J5)	9022	0	9022	0	JZJ5WORM	TS4500_3592	258	0000078D8320	-

---

```
no
```

It is now possible to physically remove tape JZ0072JZ from the tape library so that it can be sent to the off-site storage location.

### 7.21.3 Import/Export enhancements

Starting with IBM Spectrum Archive EE version V1R1.1.3 and later, the Import/Export function was enhanced and improved. These changes address the following main aspects from prior version of LTFs EE releases beside general performance improvements for Import/Export:

- ▶ Offline export against a redundant (secondary or third tape) copy can be used to move a tape or a set of tapes to a different location. In this case, even if the redundant copy is offline exported, the file can be only recalled from the primary copy.
- ▶ Only when all replicas are exported/offline exported is the file reported as exported/offline exported.
- ▶ It might take a longer time to export a tape because LTFs EE runs reconcile before exporting a tape.

Therefore, the following Import/Export enhancements were introduced to IBM Spectrum Archive EE starting with version V1R1.1.3 and later:

- ▶ Allow a stub of a premigrated file even if a redundant copy is offline exported.
- ▶ Full replica support for ease of use.
- ▶ The `ltfsee info tapes` command shows an offline tape list.
- ▶ To view the offline message, run `ltfsee tape show` with the `-a offline` option.
- ▶ The `ltfsee info files` command does not show an offline message for the files migrated to the offline exported tape.
- ▶ You can skip reconcile if you run offline export to improve performance.
- ▶ There is no update on a stub for offline export/import to improve performance.
- ▶ Keep offline exported information (Tape ID and offline message) in MMM to avoid updating the stub.

Regarding full replica support, Export/Import does not depend on the primary/redundant copy. When *all* copies are exported, the file is exported.

Table 7-1 shows a use case example where a file was migrated to three physical tapes: TAPE1, TAPE2, TAPE3. The file behaves as shown by export operations.

*Table 7-1 Export operations use case scenario of file with three tapes*

Operation	File
TAPE1 is offline exported.	File is available (can recall).
TAPE1/TAPE2 are offline exported.	File is available (can recall).
TAPE1/TAPE2/TAPE3 are offline exported.	File is offline (cannot recall).
TAPE1 is exported.	File is available (IBMTPS has TAPE2/TAPE3).
TAPE1/TAPE2 is exported.	File is available (IBMTPS has TAPE3).
TAPE1/TAPE2/TAPE3 is exported.	File is removed from GPFS.
TAPE1/TAPE2 are offline exported, and then TAPE3 is exported.	File is offline (IBMTPS has TAPE1/TAPE2).

The `ltfsee info files` command does not show the offline message even if the file is exported. It just shows offline, as shown in Example 7-101.

*Example 7-101 The ltfsee info files command not showing the offline message for exported file*

```
# ltfsee info files /mnt/gpfs/test/1MData.001
Name: /mnt/gpfs/test/1MData.001
Tape id:D00600L5 Status: offline
```

## 7.22 Drive Role settings for job assignment control

Before LTF/IBM Spectrum Archive V1 R1.1.3, the physical drive allocation logic was simple. IBM Spectrum Archive EE picked up the first found idle tape drive from the list of available drives. If the tape to be mounted was already in a drive, IBM Spectrum Archive EE picked it preferentially. This simple drive allocation logic is sometimes not good enough for a client use cases who require the following capabilities:

- ▶ Reserve drives for potential recalls, for example, not to be used for purposes other than for a quick recall job dispatch.
- ▶ Assign dedicated drives for migration jobs. Without this kind of drive assignment, the control node scheduler might change tapes from primary tape to copy one, or go in the opposite direction.

Therefore, it might be beneficial for users to configure the tape drive's capabilities to allow or disallow each type of jobs on a per-drive basis.

Configurable tape drive attributes were introduced in IBM Spectrum Archive EE version V1R1.1.3 and later. Each of the attributes corresponds to the tape drive's capability to perform a specific type of job. Here are the attributes:

- ▶ Migration
- ▶ Recall
- ▶ Generic

Table 7-2 describes each of the available IBM Spectrum Archive EE drive attributes for the attached physical tape drives.

*Table 7-2 IBM Spectrum Archive EE drive attributes*

Attributes	Description
Migration	If the Migration attribute is set for a drive, that drive can process migration jobs. If not, IBM Spectrum Archive EE never runs migration jobs by using that drive. Save jobs are also allowed/disallowed through this attribute setting. It is preferable that there be at least one tape drive that has this attribute set to Migration.
Recall	If the Recall attribute is set for a drive, that drive can process recall jobs. If not, IBM Spectrum Archive EE never runs recall jobs by using that drive. Both automatic file recall and selective file recall are enabled/disabled by using this single attribute. There is no way to enable/disable one of these two recall types selectively. It is preferable that there be at least one tape drive that has this attribute set to Recall.

Attributes	Description
Generic	If the Generic attribute is set for a drive, that drive can process generic jobs. If not, IBM Spectrum Archive EE never runs generic jobs by using that drive. IBM Spectrum Archive EE creates and runs miscellaneous generic jobs for administrative purposes, such as formatting tape, checking tape, reconciling tape, reclaiming a tape, and validating a tape. Some of those jobs are internally run with any of the user operations. It is preferable that there be at least one tape drive that has this attribute set to Generic. For reclaiming tape, at least two tape drives are required, so at least two drives need the Generic attribute.

To set these attributes for a tape drive, the attributes can be specified when adding a tape drive to IBM Spectrum Archive EE. Use the following command syntax:

```
ltfsee drive add -d <drive serial[:attr]> [-n <node_id>]
```

The **attr** option is a decimal numeric parameter that is optional for this command and can be specified after “:”. A logical **OR** applies to set the three attributes: **Migrate (4)**, **Recall (2)**, and **Generic (1)**. For example, a number of 6 for **attr** allows migration and recall job while copy and generic job are disallowed. All of the attributes are set by default. If the tape drive to update already has attributes set by IBM Spectrum Archive EE, you must remove them before adding new ones by using the **ltfsee drive remove** command.

To check the current active drive attributes, the **ltfsee info drives** command is useful. This command shows each tape drive’s attributes, as shown in Example 7-102.

*Example 7-102 Check current IBM Spectrum Archive EE drive attributes*

```
[root@ltfsm11 ~]# ltfsee info drives
```

Drive S/N	Status	Type	Role	Library	Address	Node ID	Tape	Node Group
1013000667	Not mounted	LT06	mrg	lib_ltfsm12	256	1	-	G0
1013000110	Not mounted	LT06	mrg	lib_ltfsm12	257	1	-	G0
1013000692	Not mounted	LT06	mrg	lib_ltfsm12	258	1	-	G0
1013000694	Mounted	LT06	mrg	lib_ltfsm11	256	2	2FC140L5	G0
1013000688	Mounted	LT06	mrg	lib_ltfsm11	257	2	2FC141L5	G0
1013000655	Mounted	LT06	mrg	lib_ltfsm11	258	2	2FC147L5	G0

The letters m, r, and g are shown when the corresponding attribute Migration, Recall, and Generic are set to on. If an attribute is not set, “-” is shown instead.

**Hint for the drive attributes setting:** In a multiple nodes environment, it is expected that the reclaim driver works faster if two tape drives that are used for reclaim are assigned to a single node. For that purpose, tape drives with the Generic attribute should be assigned to a single node and all of other drives of the remaining nodes should not have the Generic attribute.

## 7.23 Tape drive intermix support

This section describes the physical tape drive intermix support that was released with IBM Spectrum Archive EE V1R2.0.0 or later.

This enhancement has these objectives:

- ▶ Use IBM LTO-7 tapes and drives in mixed configuration with older IBM LTO (LTO-6 and 5) generations
- ▶ Use 3592 JC/JD cartridges along with IBM TS1155, TS1150, and TS1140 drives in mixed environments

**Note:** An intermix of LTO and TS11xx tape drive technology and media is not supported by IBM Spectrum Archive EE.

The following main use cases are expected to be used by this new feature:

- ▶ Tape media and technology migration (from old to new generation tapes)
- ▶ Continue using prior generation formatted tapes (read or write) with the current technology tape drive generation

To generate and use a mixed tape drive environment, you must define the different LTO or TS11xx drive types with the creation of the logical library partition (within your tape library) to be used along with your IBM Spectrum Archive EE setup.

When LTO-7, 6, and 5 tapes are used in a tape library, correct cartridges and drives are selected by IBM Spectrum Archive EE to read or write the required data, which includes the usage of an LTO7 drive for recall from LTO5 tapes.

When 3592 JC or 3592 JD tapes are used in a tape library and IBM TS1155, TS1150, and TS1140 drives are used for the JC tapes, correct tapes and drives are selected by IBM Spectrum Archive EE to read or write the required data.

With this new function, a data migration between different generations of tape cartridges can be achieved. You can select and configure which TS11xx format (TS1155, TS1150, or TS1140) is used by IBM Spectrum Archive EE for operating 3592 JC tapes. The default for IBM Spectrum Archive EE is always to use and format to the highest available capacity. Also, when you import 3592 JC media, you can tell those tapes in which format IBM Spectrum Archive EE uses them.

The `ltfsee pool add` command can be used for configuration when new physical tape cartridges are added to your IBM Spectrum Archive EE setup:

```
ltfsee pool add -p poolname -t tapename <-f | -F>
```

WORM support for the IBM TS1155, TS1150, and TS1140 tape drives

From the long-term archive perspective, there is sometimes a requirement to store files without any modification that is ensured by the system. Starting with IBM Spectrum Archive EE V1R2.0.0 and later, you can deploy Write Once Read Many (WORM) tape cartridges in your IBM Spectrum Archive EE setup. Only 3592 WORM tapes that can be used with IBM TS1155, TS1150, or TS1140 drives are supported.

**Note:** LTO WORM tapes are not supported for IBM Spectrum Archive EE.

For more information about IBM tape media and WORM tapes, see the following website:

<http://www.ibm.com/systems/storage/media>

## 7.23.1 Objective for WORM tape support

The IBM Spectrum Archive EE objective for WORM tapes is to store files without any modifications, which is ensured by the system, but with the following limitations:

- ▶ Only ensure that the file on tape is immutable if the user uses only IBM Spectrum Archive EE:
  - Does not detect the case where an appended modified index is at the end of tape by using a direct SCSI command.
  - From LTFS format perspective, this case can be detected but it needs time to scan every index on the tape. This feature is not provided in the release of IBM Spectrum Archive EE on which this book is based.
- ▶ Does not ensure that the file cannot be modified through GPFS in the following ways:
  - Migrate the immutable files to tape.
  - Recall the immutable files to disk.
  - Change the immutable attribute of the file on disk and modify.

## 7.23.2 Function overview for WORM tape support

To support 3592 WORM tapes, IBM Spectrum Archive EE starting with release V1R2.0.0. provides the following new features:

- ▶ Introduces a WORM attribute to the IBM Spectrum Archive EE pool attributes.
- ▶ A WORM pool can have only WORM cartridges.
- ▶ Files that have GPFS immutable attributes can still be migrated to normal pools.

Example 7-103 shows how to set the WORM attribute to an IBM Spectrum Archive EE pool by using the `ltfsee pool` command.

*Example 7-103 Set the WORM attribute to an IBM Spectrum Archive EE pool*

---

```
[root@ltfssn1 ~]# ltfsee pool create -p myWORM --worm physical
```

---

There is also an IBM Spectrum Scale layer that can provide a certain immutability for files within the GPFS file system. You can apply `immutable` and `appendOnly` restrictions either to individual files within a file set or to a directory. An `immutable` file cannot be changed or renamed. An `appendOnly` file allows append operations, but not delete, modify, or rename operations.

An `immutable` directory cannot be deleted or renamed, and files cannot be added or deleted under such a directory. An `appendOnly` directory allows new files or subdirectories to be created with 0-byte length. All such new created files and subdirectories are marked as `appendOnly` automatically.

The `immutable` flag and the `appendOnly` flag can be set independently. If both `immutability` and `appendOnly` are set on a file, `immutability` restrictions are in effect.

To set or unset these attributes, use the following IBM Spectrum Scale command options:

- ▶ `mmchattr -i yes|no`

This command sets or unsets a file to or from an immutable state:

- `-i yes`

Sets the `immutable` attribute of the file to `yes`.

- **-i no**  
Sets the immutable attribute of the file to no.

► **mmchattr -a yes|no**

This command sets or unsets a file to or from an appendOnly state:

- **-a yes**  
Sets the appendOnly attribute of the file to yes.
- **-a no**  
Sets the appendOnly attribute of the file to no.

**Note:** Before an immutable or appendOnly file can be deleted, you must change it to mutable or set appendOnly to no (by using the **mmchattr** command).

Storage pool assignment of an immutable or appendOnly file can be changed. An immutable or appendOnly file is allowed to transfer from one storage pool to another.

To display whether a file is immutable or appendOnly, run this command:

```
mmlsattr -L myfile
```

The system displays information similar to the output shown in Example 7-104.

*Example 7-104 Output of the mmlsattr -L myfile command*

---

```
file name:          myfile
metadata replication: 2 max 2
data replication:   1 max 2
immutable:         no
appendOnly:        no
flags:
storage pool name:  sp1
fileset name:      root
snapshot name:
creation Time:     Wed Feb 22 15:16:29 2012
Windows attributes: ARCHIVE
```

---

### 7.23.3 The effects of file operations on immutable and appendOnly files

After a file is set as immutable or appendOnly, the following file operations and attributes work differently from the way they work on regular files:

► **delete**

An immutable or appendOnly file cannot be deleted.

► **modify/append**

An immutable file cannot be modified or appended. An appendOnly file cannot be modified, but it can be appended.

**Note:** The immutable and appendOnly flag check takes effect after the file is closed. Therefore, the file can be modified if it is opened before the file is changed to immutable.

► **mode**

An immutable or appendOnly file's mode cannot be changed.

► **ownership, acl**

These attributes cannot be changed for an immutable or appendOnly file.

► **extended attributes**

These attributes cannot be added, deleted, or modified for an immutable or appendOnly file.

► **timestamp**

The time stamp of an immutable or appendOnly file can be changed.

► **directory**

If a directory is marked as immutable, no files can be created, renamed, or deleted under that directory. However, a subdirectory under an immutable directory remains mutable unless it is explicitly changed by the **mmchattr** command.

If a directory is marked as appendOnly, no files can be renamed or deleted under that directory. However, 0-byte length files can be created.

For more information about IBM Spectrum Scale V4.1.1 immutable and appendOnly limitations, see IBM Knowledge Center at:

[https://www.ibm.com/support/knowledgecenter/en/STXKQY\\_4.2.0/com.ibm.spectrum.scale.v4r2.adv.doc/b11adv\\_integratedarchiveplatform.htm](https://www.ibm.com/support/knowledgecenter/en/STXKQY_4.2.0/com.ibm.spectrum.scale.v4r2.adv.doc/b11adv_integratedarchiveplatform.htm)

Example 7-105 shows the output that you receive while working, showing, and changing the IBM Spectrum Scale immutable or appendOnly file attributes.

*Example 7-105 Set or change an IBM Spectrum Scale file immutable file attribute*

---

```
[root@l1tfsee_node0]# echo "Jan" > jan_jonas.out
[root@l1tfsee_node0]# mmlsattr -L -d jan_jonas.out
file name:          jan_jonas.out
metadata replication: 1 max 2
data replication:   1 max 2
immutable:       no
appendOnly:        no
flags:
storage pool name:  system
fileset name:       root
snapshot name:
creation time:      Mon Aug 31 15:40:54 2015
Windows attributes: ARCHIVE
Encrypted:          yes
gpfs.Encryption:
0x454147430001008C525B9D47000000000010001000200200008000254E60BA4024AC1D500010001
00010003000300012008921539C65F5614BA58F71FC97A46771B9195846A9A90F394DE67C4B9052052
303A82494546897FA229074B45592D61363532323261642D653862632D346663632D383961332D3461
37633534643431383163004D495A554E4F00
EncPar 'AES:256:XTS:FEK:HMACSHA512'
      type: wrapped FEK  WrpPar 'AES:KWRAP'  CmbPar 'XORHMACSHA512'
      KEY-a65222ad-e8bc-4fcc-89a3-4a7c54d4181c:l1tfssn2

[root@l1tfsee_node0]# mmchattr -i yes jan_jonas.out

[root@l1tfsee_node0]# mmlsattr -L -d jan_jonas.out
```

```

file name:          jan_jonas.out
metadata replication: 1 max 2
data replication:   1 max 2
immutable:       yes
appendOnly:        no
flags:
storage pool name:  system
fileset name:       root
snapshot name:
creation time:      Mon Aug 31 15:40:54 2015
Windows attributes: ARCHIVE READONLY
Encrypted:          yes
gpfs.Encryption:
0x454147430001008C525B9D47000000000010001000200200008000254E60BA4024AC1D500010001
00010003000300012008921539C65F5614BA58F71FC97A46771B9195846A9A90F394DE67C4B9052052
303A82494546897FA229074B45592D61363532323261642D653862632D346663632D383961332D3461
37633534643431383163004D495A554E4F00
EncPar 'AES:256:XTS:FEK:HMASHA512'
  type: wrapped FEK  WrpPar 'AES:KWRAP'  CmbPar 'XORHMASHA512'
  KEY-a65222ad-e8bc-4fcc-89a3-4a7c54d4181c:1tfssn2

[root@l1fsee_node0]# echo "Jonas" >> jan_jonas.out
-bash: jan_jonas.out: Read-only file system
[root@l1fsee_node0]#

```

---

These `immutable` or `appendOnly` file attributes can be changed at any time by the IBM Spectrum Scale administrator. They cannot provide an ultimate immutability.

If you are working with IBM Spectrum Archive EE and IBM Spectrum Scale and you plan implementing a WORM solution along with WORM tape cartridges, these two main assumptions apply:

- ▶ Only files that have IBM Spectrum Scale with the `immutable` attribute ensure no modification.
- ▶ The IBM Spectrum Scale `immutable` attribute is not changed after it is set unless it is changed by an administrator.

Consider the following limitations when using WORM tapes together with IBM Spectrum Archive EE:

- ▶ WORM tapes are supported only with IBM TS1155, TS1150, and TS1140 tape drives (3592 JY, JZ).
- ▶ If IBM Spectrum Scale `immutable` attributes are changed *after* migration, the next migration fails against the same WORM pool.
- ▶ IBM Spectrum Archive EE supports the following operations with WORM media:
  - Migrate
  - Recall
  - Offline export and offline import
- ▶ IBM Spectrum Archive EE does not support the following operations with WORM media:
  - Reclaim
  - Reconcile
  - Export and Import

For more information about the IBM Spectrum Archive EE commands, see 11.1, “Command-line reference” on page 300.

## 7.24 Obtaining the location of files and data

This section describes how to obtain information about the location of files and data by using IBM Spectrum Archive EE.

You can use the `files` option of the `ltfsee info` command to discover the physical location of files. To help with the management of replicas, this command also indicates which tape cartridges are used by a particular file. To quickly determine the file states for many files, see the IBM Spectrum Scale wiki:

<https://ibm.biz/BdsgqT>

Example 7-106 shows the typical output of the `ltfs info files` command. Some files are on multiple tape cartridges, some are in a migrated state, and others are premigrated only.

*Example 7-106 Files location*

---

```
[root@ltfs97 glues]# ltfsee info files offsite*
File name: offsite1.mpeg
Tape id:037AGWL5 Status: premigrated
File name: offsite2.mpeg
Tape id:037AGWL5:051AGWL5 Status: migrated
File name: offsite3.mpeg
Tape id:037AGWL5:051AGWL5 Status: migrated
File name: offsite4.mpeg
Tape id:037AGWL5 Status: migrated
File name: offsite5.mpeg
Tape id:037AGWL5:051AGWL5 Status: migrated
```

---

The list of supported characters for file names and directory path names can be found at:

[https://www.ibm.com/support/knowledgecenter/en/STZMZN\\_2.4.0/ltfs\\_restricted\\_characters\\_library\\_edition.html](https://www.ibm.com/support/knowledgecenter/en/STZMZN_2.4.0/ltfs_restricted_characters_library_edition.html)

**Note:** Migration of files with special characters in the file name or directory path name might succeed, but reconciliation and normal export will fail due to symlink creation failure.

## 7.25 Obtaining inventory, job, and scan status

This section describes how to obtain resource inventory information or information about ongoing migration and recall jobs or scans with IBM Spectrum Archive EE. You can use the `info` option of the `ltfsee` command to obtain information about current jobs and scans and resource inventory information about tape cartridges, drives, nodes, tape cartridge pools, and files. Example 7-107 shows the command that is used to display all IBM Spectrum Archive EE tape cartridge pools.

*Example 7-107 Tape cartridge pools*

---

```
[root@ltfsm11 ~]# ltfsee info pools
Pool Name      Total(TiB) Free(TiB) Unref(TiB) Tapes Type Library Node
Group
```

primary_ltfsml2	75.2	1.1	0.1	58	LTO	lib_ltfsml2	GO
primary_ltfsml1	70.0	4.1	0.1	54	LTO	lib_ltfsml1	GO
copy_ltfsml1	5.2	5.2	0.0	4	LTO	lib_ltfsml1	GO

Example 7-108 shows the serial numbers and status of the tape drives that are used by IBM Spectrum Archive EE.

*Example 7-108 Drives*

```
[root@ltfsml1 ~]# ltfsee info drives
```

Drive S/N	Status	Type	Role	Library	Address	Node ID	Tape	Node Group
1013000667	Not mounted	LT06	mrg	lib_ltfsml2	256	1	-	GO
1013000110	Not mounted	LT06	mrg	lib_ltfsml2	257	1	-	GO
1013000692	Not mounted	LT06	mrg	lib_ltfsml2	258	1	-	GO
1013000694	Mounted	LT06	mrg	lib_ltfsml1	256	2	2FC140L5	GO
1013000688	Mounted	LT06	mrg	lib_ltfsml1	257	2	2FC141L5	GO
1013000655	Mounted	LT06	mrg	lib_ltfsml1	258	2	2FC147L5	GO

To view all the IBM Spectrum Archive EE tape cartridges, run the command that is shown in Example 7-109.

*Example 7-109 Tape cartridges*

```
root@ltfsml1 ~]# ltfsee info tapes
```

Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)	Pool	Library	Address	Drive	Appendable
TST200L6	Valid	L6	2242	0	2242	0	test1	TS3310	4117	-	yes
LTF01L7	Valid	L7	5338	0	5338	0	test	TS3310	4119	-	yes
PFE001L6	Valid	L6	2242	95	2146	0	test1	TS3310	4118	-	yes

Regularly monitor the output of IBM Spectrum Archive EE **info** jobs and IBM Spectrum Archive EE **info** scans to ensure that jobs and scans are progressing as expected. Use the command syntax that is shown in Example 7-110.

*Example 7-110 Jobs*

```
[root@lufs97] ltfsee info jobs
```

Job Type	Status	Idle(sec)	Scan ID	Tape	Pool	Library	Node	File Name or inode
Selective Recall	In-progress	0	3108835585	2FC140L5	copy_ltfsml1	lib_ltfsml1	2	10374858
Selective Recall	In-progress	0	3108835585	2FC141L5	copy_ltfsml1	lib_ltfsml1	2	10374888
Selective Recall	Unscheduled	0	3108835585	2FC140L5	copy_ltfsml1	lib_ltfsml1	-	10374877
Selective Recall	Unscheduled	0	3108835585	2FC140L5	copy_ltfsml1	lib_ltfsml1	-	10374877

The **ltfsee info nodes** command (see Example 7-111) provides a summary of the state of each IBM Spectrum Archive LE+ component node.

*Example 7-111 lufs info nodes*

```
[root@ltfsml1 ~]# ltfsee info nodes
```

Node ID	Status	Node IP	Drives	Ctrl Node	Library	Node Group	Host Name
2	Available	9.11.120.224	2	Yes	lib_lufs2	GO	ltfsml2
1	Available	9.11.120.207	2	Yes(Active)	lib_lufs1	GO	ltfsml1

The `ltfsee info scans` command (see Example 7-112) provides a summary of the current active scans in IBM Spectrum Archive EE. The scan number corresponds to the job number that is reported by the `ltfsee info jobs` command.

*Example 7-112 ltfsee info scans*

---

Scan ID	Parent	Queued	Succeed	Failed	Canceled	Library	Ctrl Node
3587183105	-	1	0	0	0	lib_ltfsm11	ltfsm11.tuc.stglabs.ibm.com
3587183361	3587183105	1	0	0	0	lib_ltfsm11	ltfsm11.tuc.stglabs.ibm.com

---

## 7.26 Cleaning up a scan or session

This section describes how to perform a cleanup if a scan or session is complete, but was not removed by the calling process. An example is when a window was closed while the `migrate` command was running.

The `cleanup` option of the `ltfsee` command can clean up a scan or session that finished but cannot be removed by its calling process. If the scan is not finished, you cannot “clean it up”, as shown in Example 7-113.

*Example 7-113 Clean up a running scan*

---

```
[root@ltfs97 ~]# ltfsee cleanup 2282356739
Session or scan with id 2282356739 is not finished.
```

---

## 7.27 Monitoring the system with SNMP

You can use SNMP traps to receive notifications about system events. There are many processes that should be reported through SNMP. Starting with IBM Spectrum Archive EE v1.2.4.0, IBM Spectrum Archive EE uses SNMP to monitor the system and send alerts when the following events occur:

- ▶ IBM Spectrum Archive EE component errors are detected.
- ▶ Recovery actions are performed on failed components.
- ▶ At the end of an `ltfsee start` and `ltfsee stop` command, and at the successful or unsuccessful start or stop of each component of an IBM Spectrum Archive EE node.
- ▶ When the remaining space threshold for a pool is reached.

The MIB file is installed in `/opt/ibm/ltfsee/share/IBMSA-MIB.txt` on each node. It should be copied to the `/usr/share/snmp/mibs/` directory on each node.

Table 7-3 lists SNMP traps that can be issued, showing the error message that is generated, along with the trap name, severity code, and the OID for each trap.

*Table 7-3 SNMP error message information*

Description	Name	Severity	OID
GLESV100: IBM Spectrum Archive EE successfully started or restarted.	ibmsaInfoV100StartSuccess	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.100
GLESV101E: IBM Spectrum Archive EE failed to start.	ibmsaErrV101StartFail	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.101

Description	Name	Severity	OID
GLESV102W: Part of IBM Spectrum Archive EE nodes failed to start.	ibmsaDegradeV102PartialStart	Degraded	1.3.6.1.4.1.2.6.246.1.2.31.0.102
GLESV103I: IBM Spectrum Archive EE successfully stopped.	ibmsaInfoV103StopSuccess	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.103
GLESV104E: IBM Spectrum Archive EE failed to stop.	ibmsaErrV104StopFail	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.104
GLESV300E: GPFS error has been detected.	ibmsaErrV300GPFSError	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.300
GLESV301I: GPFS becomes operational.	ibmsaInfoV301GPFSOperational	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.301
GLESV302I: The IBM Spectrum Archive LE successfully started.	ibmsaInfoV302LEStartSuccess	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.302
GLESV303E: The IBM Spectrum Archive LE failed to start.	ibmsaErrV303LEStartFail	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.303
GLESV304I: The IBM Spectrum Archive LE successfully stopped.	ibmsaInfoV304LEStopSuccess	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.304
GLESV305I: IBM Spectrum Archive LE is detected.	ibmsaInfoV305LEDetected	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.305
GLESV306E: IBM Spectrum Archive LE process does not exist.	ibmsaErrV306LENotExist	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.306
GLESV307E: IBM Spectrum Archive LE process is not responding.	ibmsaErrV307LENotRespond	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.307
GLESV308I: IBM Spectrum Archive LE process is now responding.	ibmsaInfoV308LERespond	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.308
GLESV309I: The process 'rpcbind' started up.	ibmsaInfoV309RpcbindStart	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.309
GLESV310E: The process 'rpcbind' does not exist.	ibmsaErrV310RpcbindNotExist	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.310
GLESV311I: The process 'rsyslogd' started up.	ibmsaInfoV311RsyslogdStart	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.311
GLESV312E: The process 'rsyslogd' does not exist.	ibmsaErrV312RsyslogdNotExist	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.312
GLESV313I: The process 'sshd' started up.	ibmsaInfoV313SshdStart	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.313
GLESV314E: The process 'sshd' does not exist.	ibmsaErrV314SshdNotExist	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.314
GLESV315I: The IBM Spectrum Archive EE service (MMM) successfully started	ibmsaInfoV315MMMStartSuccess	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.315
GLESV316E: The IBM Spectrum Archive EE service (MMM) failed to start.	ibmsaErrV316MMMStartFail	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.316
GLESV317I: The IBM Spectrum Archive EE service (MMM) successfully stopped.	ibmsaInfoV317MMMStopSuccess	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.317
GLESV318I: The IBM Spectrum Archive EE service (MMM) is detected.	ibmsaInfoV318MMMDetected	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.318
GLESV319E: The IBM Spectrum Archive EE service (MMM) does not exist.	ibmsaErrV319MMMNotExist	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.319
GLESV320E: The IBM Spectrum Archive EE service (MMM) is not responding.	ibmsaErrV320MMMNotRespond	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.320

Description	Name	Severity	OID
GLESV321I: The IBM Spectrum Archive EE service (MMM) now responding.	ibmsaInfoV321MMMRespond	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.321
GLESV322I: The IBM Spectrum Archive EE service (MD) successfully started.	ibmsaInfoV322MDStartSuccess	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.322
GLESV323E: The IBM Spectrum Archive EE service (MD) failed to start.	ibmsaErrV323MDStartFail	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.323
GLESV324I: The IBM Spectrum Archive EE service (MD) successfully stopped.	ibmsaInfoV324MDStopSuccess	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.324
GLESV325I: The IBM Spectrum Archive EE service (MD) is detected.	ibmsaInfoV325MDDetected	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.325
GLESV326E: The IBM Spectrum Archive EE service (MD) does not exist.	ibmsaErrV326MDNotExist	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.326
GLESV327E: The IBM Spectrum Archive EE service (MD) is not responding.	ibmsaErrV327MDNotRespond	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.327
GLESV328I: The IBM Spectrum Archive EE service (MD) is now responding.	ibmsaInfoV328MDRespond	Information	1.3.6.1.4.1.2.6.246.1.2.31.0.328

## 7.28 Configuring Net-SNMP

It is necessary to modify the `/etc/snmp/snmpd.conf` and `/etc/snmp/snmptrapd.conf` configuration file to receive SNMP traps. These files should be modified on each node that has IBM Spectrum Archive EE installed and running.

To configure Net-SNMP, complete the following steps on each IBM Spectrum Archive EE node:

1. Open the `/etc/snmp/snmpd.conf` configuration file.
2. Add the following entry to the file:

```
master agentx
trap2sink <managementhost>
```

The variable `<managementhost>` is the host name or IP address of the host to which the SNMP traps are sent.

3. Open the `/etc/snmp/snmptrapd.conf` configuration file.
4. Add the following entry to the file:

```
disableauthorization yes
```

5. Restart the SNMP daemon by running the following command:

```
[root@ltfs97 ~]# systemctl restart snmpd.service
```

### 7.28.1 Starting and stopping the snmpd daemon

Before IBM Spectrum Archive EE is started, you must start the `snmpd` daemon on all nodes where IBM Spectrum Archive EE is running.

To start the `snmpd` daemon, run the following command:

```
[root@ltfs97 ~]# systemctl start snmpd.service
```

To stop the snmpd daemon, run the following command:

```
[root@l1tfs97 ~]# systemctl stop snmpd.service
```

To restart the snmpd daemon, run the following command:

```
[root@l1tfs97 ~]# systemctl restart snmpd.service
```

## 7.28.2 Example of an SNMP trap

Example 7-114 shows the type of trap information that is received by the SNMP server.

*Example 7-114 SNMP trap example of an IBM Spectrum Archive EE node that failed to start*

---

```
2017-11-08 13:44:40 l1tfsn1 [UDP: [9.11.120.131]:60096->[9.11.120.131]:162]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (41605) 0:06:56.05
SNMPv2-MIB::snmpTrapOID.0 = OID: IBMSA-MIB::ibmsaErrV101StartFail
IBMSA-MIB::ibmsaMessageSeverity.0 = INTEGER: error(20) IBMSA-MIB::ibmsaJob.0 = INTEGER:
other(7) IBMSA-MIB::ibmsaEventNode.0 = STRING: "l1tfsn1" IBMSA-MIB::ibmsaMessageText.0 =
STRING: "GLESV101E: The Spectrum Archive Enterprise Edition node failed to
start."
```

---

## 7.29 IBM Spectrum Archive REST API

The IBM Spectrum Archive EE REST API gives users another interface to interact with the IBM Spectrum Archive EE product. REST uses http GET operations to return status information about IBM Spectrum Archive EE. This section covers the GET operations for the IBM Spectrum Archive EE Rest API.

The IBM Spectrum Archive EE REST API can be accessed in two ways. The first is through a terminal window with the **curl** command and the second way is through a web browser. Both ways output the same data. In this section, the **curl** command is used.

The following are supported GET operations when executing rest commands:

- ▶ **pretty**  
Specify for pretty-printing. The default value is `false`.
- ▶ **sort: <string>(,<string>...)**  
Specify field name or names to use as sort key. The default sort order is ascending. Use the "-" sign to sort in descending order.
- ▶ **fields: <string>(,<string>...)**  
Specify field names that are to be included in the response.

**Note:** The examples in this section are all performed on the server with the REST RPM installed and uses `localhost` to request resources. When accessing the REST API over a remote server, replace `localhost` with the server ip with the rest rpms installed.

### 7.29.1 Pools endpoint

IBM Spectrum Archive EE rest **pools** endpoint returns all information in JSON objects about each pool created thus far within the environment.

The following is an example command of calling the **pools** endpoint using localhost:

```
curl -X GET 'http://localhost:7100/ibmsa/v1/pools/'
```

The following is the response data returned when requesting for the **pools** endpoint:

- ▶ **id: <string>**  
UUID of Pool, assigned by system at the creation of pool.
- ▶ **name: <string>**  
User-specified name of pool.
- ▶ **capacity: <number>**  
Total capacity of tapes assigned to the pool, in bytes. The capacity = used\_space + free\_space.
- ▶ **used\_space: <number>**  
Used space of the pool, in bytes. The used\_space = active\_space + reclaimable\_space.
- ▶ **free\_space: <number>**  
Free space of the pool, in bytes.
- ▶ **active\_space: <number>**  
Active space (used space consumed by active-referred files) of the pool, in bytes.
- ▶ **reclaimable\_space: <number>**  
The reclaimable space (used space consumed by unreferred files) of the pool, in bytes. Note that this is the amount of estimated size of the unreferenced space that is available for reclamation on the assigned tapes.
- ▶ **num\_of\_tapes: <number>**  
Number of tapes assigned to the pool.
- ▶ **format\_class: <string>**  
The format class of the pool.
- ▶ **library\_name: <string>**  
Library name that the pool belongs to.
- ▶ **library\_id: <string>**  
Library ID (serial number) that the pool belongs to.
- ▶ **node\_group: <string>**  
Node group name that the pool belongs to.
- ▶ **device\_type: <string>**  
Tape device type that can be added to the pool. Can be either 3592, LT0, or left blank.
- ▶ **worm: <string>**  
WORM type. Can be either physical, logical, no, or unknown.
- ▶ **fill\_policy: <string>**  
Tape fill policy.
- ▶ **owner: <string>**  
Owner.
- ▶ **mount\_limit: <number>**  
Maximum number of drives that can be used for migration. 0 means unlimited.

- ▶ **low\_space\_warning\_enable:** <bool>  
Whether monitoring thread sends SNMP trap for low space pool or not.
- ▶ **low\_space\_warning\_threshold:** <number>  
SNMP notification threshold value for free pool size in bytes. 0 when no threshold is set.
- ▶ **no\_space\_warning\_enable:** <bool>  
Whether monitoring thread sends SNMP trap for no space pool or not.

Parameters are available to be passed in to filter specific pools. The following are the available parameters:

- ▶ **name:** <string>  
Filter the list of pools by name. Only the pools that match the criteria are returned in the response.
- ▶ **library\_name:** <string>  
Filter the list of pools by library name. Only the pools that match the criteria are returned in the response.

Example 7-115 shows how to request the pools resource through **curl** commands.

*Example 7-115 REST pool command*

---

```
[root@l1fseesrv1]# curl -X GET 'http://localhost:7100/ibmsa/v1/pools?pretty=true'
[
  {
    "active_space": 0,
    "capacity": 11464284831744,
    "device_type": "LT0",
    "fill_policy": "Default",
    "format_class": "DEFAULT",
    "free_space": 11464273297408,
    "id": "14adb6cf-d1f5-46ef-a0bb-7b3881bdb4ec",
    "library_id": "0000013400190402",
    "library_name": "l1fs_lib",
    "low_space_warning_enable": true,
    "low_space_warning_threshold": 0,
    "mount_limit": 0,
    "name": "primary_cache",
    "no_space_warning_enable": true,
    "node_group": "l1fseesrv1",
    "num_of_tapes": 2,
    "owner": "System",
    "reclaimable_space": 11534336,
    "used_space": 11534336,
    "worm": "no"
  },
  {
    "active_space": 0,
    "capacity": 11464284831744,
    "device_type": "LT0",
    "fill_policy": "Default",
    "format_class": "DEFAULT",
    "free_space": 11464273297408,
    "id": "902b097a-7a34-4847-a346-0e6d97444a21",
    "library_id": "0000013400190402",
```

```

    "library_name": "ltfs_lib",
    "low_space_warning_enable": true,
    "low_space_warning_threshold": 0,
    "mount_limit": 0,
    "name": "copy_cache",
    "no_space_warning_enable": true,
    "node_group": "ltfseesrv1",
    "num_of_tapes": 2,
    "owner": "System",
    "reclaimable_space": 11534336,
    "used_space": 11534336,
    "worm": "no"
  }
]

```

---

Example 7-116 shows how to call pools with specifying specific fields and sorting the output in descending order.

*Example 7-116 REST API pools endpoint*

---

```

[root@ltfseesrv1 ~]# curl -X GET
'http://localhost:7100/ibmsa/v1/pools?pretty=true&fields=capacity,name,library_name,free_space,num_of_tapes,device_type&sort=-free_space'
[
  {
    "capacity": 11464284831744,
    "device_type": "LT0",
    "free_space": 11464273297408,
    "library_name": "AnacondaQ",
    "name": "copy_cache",
    "num_of_tapes": 2
  },
  {
    "capacity": 11464284831744,
    "device_type": "LT0",
    "free_space": 11454007738368,
    "library_name": "AnacondaQ",
    "name": "primary_cache",
    "num_of_tapes": 2
  }
]

```

---

## 7.29.2 Tapes endpoint

The **tapes** endpoint returns an array of JSON objects regarding tape information. The following is an example command of calling **tapes**:

```
curl -X GET 'http://localhost:7100/ibmsa/v1/tapes'
```

The following is the response data when requesting the **tapes** endpoint:

► **id:** <string>

Tape ID. Because barcode is unique within a tape library only, the id is in format of <barcode>@<library\_id>.

- ▶ **barcode:** <string>  
Barcode of the tape.
- ▶ **status:** <string>  
Tape status.
- ▶ **media\_type:** <string>  
Media type of a tape. Media type is set even if the tape is not assigned to any pool yet. Empty string if the tape is not supported by IBM Spectrum Archive.
- ▶ **media\_generation:** <string>  
Media generation of a tape. Media generation determines a possible format that the tape can be written in.
- ▶ **format\_density:** <string>  
Format of a tape. Empty string if the tape is not assigned to any pool.
- ▶ **worm:** <bool>  
Whether WORM is enabled for the tape.
- ▶ **Appendable:** <string>  
Appendable value either yes or no, determines if tape can be written to.
- ▶ **capacity:** <number>  
Capacity of the tape, in bytes.  $capacity = used\_space + free\_space$ .
- ▶ **used\_space:** <number>  
Used space of the tape, in bytes.  $used\_space = active\_space + reclaimable\_space$ .
- ▶ **free\_space:** <number>  
Free space of the tape, in bytes.
- ▶ **active\_space:** <number>  
Active space (used space consumed by active-referred files) of the tape, in bytes.
- ▶ **reclaimable\_space:** <number>  
Reclaimable space (used space consumed by unreferred files) of the tape, in bytes. This amount is the estimated size of the unreferenced space that is available for reclamation on the tape.
- ▶ **address:** <number>  
Address of this tape.
- ▶ **pool\_name:** <string>  
Pool name that this tape is assigned to. Empty string if the tape is not assigned to any pool.
- ▶ **pool\_id:** <string>  
Pool id that this tape is assigned to. Empty string if the tape is not assigned to any pool.
- ▶ **library\_name:** <string>  
Library name that the pool is belongs to.
- ▶ **library\_id:** <string>  
Library ID (serial number) that the pool is belongs to.
- ▶ **drive:** <string>  
Drive serial number that the tape is mounted on. Empty string if the tape is not mounted.

► **offline\_msg: <string>**

Offline message that can be specified when performing offline-export.

The following are available parameters to use to filter tape requests:

► **barcode: <string>**

Filter the list of tapes by barcode. Only the tapes that match the criteria are returned in the response.

► **library\_name: <string>**

Filter the list of tapes by library name. Only the tapes that match the criteria are returned in the response.

► **pool\_id: <string>**

Filter the list of tapes by pool ID. Only the tapes that match the criteria are returned in the response.

► **pool\_name: <string>**

Filter the list of tapes by pool name. Only the tapes that match the criteria are returned in the response.

Example 7-117 shows the output of requesting tapes resource.

*Example 7-117 REST API tapes endpoint*

---

```
[root@l1tfseesrv1 ~]# curl -X GET
'http://localhost:7100/ibmsa/v1/tapes?pretty=true'
[
{
  "active_space": 95869806,
  "address": "1056",
  "appendable": "yes",
  "barcode": "1FB117L6",
  "capacity": 2408088338432,
  "drive_id": "",
  "format_density": "L6",
  "free_space": 2407987675136,
  "id": "1FB117L605d5cc893-8853-4e2e-b89d-b4ee2b57bf2a",
  "library_id": "5d5cc893-8853-4e2e-b89d-b4ee2b57bf2a",
  "library_name": "liba",
  "media_generation": "LT0-Gen6",
  "media_type": "L6",
  "non_appendable_space": 0,
  "offline_msg": "",
  "pool_id": "e830e3795f364108a1907542e086cb8f",
  "pool_name": "pool2",
  "reclaimable_space": 4793490,
  "status": "Valid",
  "used_space": 100663296,
  "worm": false
}
]
```

---

### 7.29.3 Libraries endpoint

The `libraries` endpoint returns information regarding the library that the node is connected to, such as the library ID, name and model type. The following is an example of the `libraries curl` command:

```
curl -X GET 'http://localhost:7100/ibmsa/v1/libraries/'
```

The following is the response data that is returned when requesting this endpoint:

- ▶ **id:** <string>  
Serial number of the library
- ▶ **name:** <string>  
User-defined name of the library
- ▶ **model:** <string>  
Mode type of the library

The available filtering parameter for the `libraries` endpoint is the name of the library.

### 7.29.4 Nodegroups endpoint

The `nodegroups` endpoint returns information regarding the node groups that the nodes are part of, such as the nodegroup ID, name, number of nodes, library ID, and library name. The following is an example of calling the `nodegroups` endpoint:

```
curl -X GET 'http://localhost:7100/ibmsa-rest/v1/nodegroups/'
```

The following is the response data that is returned when requesting this endpoint:

- ▶ **id:** <string>  
Nodegroup ID. Because nodegroup name is unique within a tape library only, the ID is in format of <nodegroup\_name>@<library\_id>.
- ▶ **name:** <string>  
User-specified name of the node group.
- ▶ **num\_of\_nodes:** <number>  
The number of nodes assigned to the node group.
- ▶ **library\_id:** <string>  
The library ID (serial number) that the node group belongs to.
- ▶ **library\_name:** <string>  
The name of the library that the node group belongs to.

The available filtering parameters for `nodegroups` are `name` (nodegroup name), and `library_name`. These parameters filter out nodegroups that do not meet the values passed in.

### 7.29.5 Nodes endpoint

The `nodes` endpoint returns information about each EE node assigned to the cluster. The following is an example of calling the `nodes` endpoint:

```
curl -X GET 'http://localhost:7100/ibmsa/v1/nodes/'
```

The following is the response data when requesting the nodes endpoint:

- ▶ **id: <number>**  
The node ID, which is the same value as the corresponding to IBM Spectrum Scale node IDs.
- ▶ **ip: <string>**  
The IP address of the node. Specifically, this is the 'Primary network IP address' in GPFS.
- ▶ **hostname: <string>**  
The host name of the node (the 'GPFS daemon node interface name' in GPFS).
- ▶ **status: <string>**  
The LE status of the node.
- ▶ **num\_of\_drives: <number>**  
The number of drives attached to the node.
- ▶ **control\_node: <bool>**  
True if the node is configured as control node.
- ▶ **active\_control\_node: <bool>**  
True if the node is configured as a control node and is active.
- ▶ **library\_id: <string>**  
The library ID of the library that the node is attached to.
- ▶ **library\_name: <string>**  
The name of the library that the node is attached to.
- ▶ **nodegroup\_id: <string>**  
The ID of the nodegroup that the node is belongs to.
- ▶ **nodegroup\_name: <string>**  
The name of the nodegroup that the node is belongs to.

The available filtering parameters for the nodes endpoint are `library_name` and `nodegroup_name`. These parameters will filter out nodes that do not match the passed in values.

## 7.29.6 Drives endpoint

The drives endpoint returns information about each visible drive within the EE cluster. The following is an example of calling the drives endpoint:

```
curl -X GET 'http://localhost:7100/ibmsa/v1/drives/'
```

The following is the response data when requesting the drives endpoint:

- ▶ **id: <string>**  
Serial number of the drive.
- ▶ **status: <string>**  
Drive status.
- ▶ **type: <string>**  
Drive type, which can be empty if a drive is not assigned to any node group.

- ▶ **role: <string>**  
Three character string to represent the drive role. Can be empty if a drive is not assigned to any node group.
- ▶ **library\_id: <string>**  
The ID of the library that the drive belongs to.
- ▶ **library\_name: <string>**  
The name of the library that the drive belongs to.
- ▶ **library\_address: <number>**  
The address of the drive within the library.
- ▶ **node\_id: <string>**  
The ID of the node that the drive is assigned to.
- ▶ **tape\_barcode: <string>**  
The barcode of a tape that the drive is mounted to. Empty if no tape is mounted on the drive.
- ▶ **nodegroup\_name: <string>**  
The name of a nodegroup that the drive is assigned to.

The available filtering parameters for the drive endpoint are `library_name` and `nodegroup_name`. These parameters will filter out drives that do not match the passed drives in values.





## Hints, tips, and preferred practices

This chapter provides you with hints, tips, and preferred practices for the IBM Spectrum Archive Enterprise Edition (IBM Spectrum Archive EE). It covers various aspects about IBM Spectrum Scale, including reuse of tape cartridges, scheduling, and disaster recovery (DR). Some aspects might overlap with functions that are described in Chapter 7, “Operations” on page 129 and Chapter 10, “Troubleshooting” on page 273. However, it is important to list them here in the context of hints, tips, and preferred practices.

This chapter includes the following topics:

- ▶ Preventing migration of the .SPACEMAN and metadata directories
- ▶ Maximizing migration performance with redundant copies
- ▶ Changing the SSH daemon settings
- ▶ Setting mmapplypolicy options for increased performance
- ▶ Setting the inode size for the GPFS file system for increased performance
- ▶ Determining the file states for all files within the GPFS file system
- ▶ Memory considerations on the GPFS file system for increased performance
- ▶ Increasing the default maximum number of inodes in IBM Spectrum Scale
- ▶ Configuring IBM Spectrum Scale settings for performance improvement
- ▶ Real world use cases for mmapplypolicy
- ▶ Capturing a core file on RHEL with abrtcd
- ▶ Antivirus considerations
- ▶ Automatic email notification with rsyslog
- ▶ Overlapping IBM Spectrum Scale policy rules
- ▶ Storage pool assignment
- ▶ Tape cartridge removal
- ▶ Reusing LTFS formatted tape cartridges
- ▶ Reusing non-LTFS tape cartridges
- ▶ Moving tape cartridges between pools
- ▶ Offline tape cartridges
- ▶ Scheduling reconciliation and reclamation
- ▶ License Expiration Handling
- ▶ Disaster recovery
- ▶ IBM Spectrum Archive EE problem determination

- ▶ Collecting IBM Spectrum Archive EE logs for support
- ▶ Backing up file systems that are not managed by IBM Spectrum Archive EE

**Important:** All of the command examples in this chapter use the command without the full file path name because we added the IBM Spectrum Archive EE directory (/opt/ibm/1tfsee/bin) to the PATH variable of the operating system.

## 8.1 Preventing migration of the .SPACEMAN and metadata directories

This section describes an IBM Spectrum Scale policy rule that you should have in place to help ensure the correct operation of your IBM Spectrum Archive EE system.

You can prevent migration of the .SPACEMAN directory and the IBM Spectrum Archive EE metadata directory of a GPFS file system by excluding these directories by having an IBM Spectrum Scale policy rule in place. Example 8-1 shows how an `exclude` statement can look in an IBM Spectrum Scale migration policy file where the metadata directory starts with the text `/ibm/glues/.ltfsee`.

*Example 8-1 IBM Spectrum Scale sample directory exclude statement in the migration policy file*

---

```
define(  
    user_exclude_list,  
    (  
        PATH_NAME LIKE '/ibm/glues/.ltfsee/%'  
        OR PATH_NAME LIKE '%/.SpaceMan/%'  
        OR PATH_NAME LIKE '%/.snapshots/%'  
    )  
)
```

---

For more information and detailed examples, see 7.10.2, “Threshold-based migration” on page 161 and 7.10.3, “Manual migration” on page 166.

## 8.2 Maximizing migration performance with redundant copies

To minimize drive mounts/unmounts and to maximize performance with multiple copies, set the mount limit per tape cartridge pool to equal the number of tape drives in the node group divided by the number of copies. The mount limit attribute of a tape cartridge pool specifies the maximum allocated number of drives that are used for migration for the tape cartridge pool. A value of 0 means no limit and is also the default value.

For example, if there are four drives and two copies initially, set the mount limit to 2 for the primary tape cartridge pool and 2 for the copy tape cartridge pool. These settings maximize the migration performance because both the primary and copy jobs are run in parallel by using two tape drives each for each tape cartridge pool. This action also avoids unnecessary mounts/unmounts of tape cartridges.

To show the current mount limit setting for a tape cartridge pool, run the following command:

```
ltfsee pool show -p <poolname> [-l <libraryname>]
```

To set the mount limit setting for a tape cartridge pool, run the following command:

```
ltfsee pool set -p <poolname> [-l <libraryname>] -a <attribute> -v <value>
```

To set the mount limit attribute to 2, run the `lzfsee pool show` and `lzfsee pool set` commands, as shown in Example 8-2.

*Example 8-2 Set the mount limit attribute to 2*

---

```
[root@lzfsm11 ~]# lzfsee pool show -p primary_lzfsm11 -l lib_lzfsm11
Attribute  Value
poolname   primary_lzfsm11
poolid     17fba0a9-a49b-4f16-baf5-5e0395136468
devtype    LT0
format     0xFFFFFFFF
worm       no (0)
nodegroup  G0
fillpolicy Default
owner      System
mountlimit 0

[root@lzfsm11 ~]# lzfsee pool set -p primary_lzfsm11 -l lib_lzfsm11 -a mountlimit
-v 2

[root@lzfsm11 ~]# lzfsee pool show -p primary_lzfsm11 -l lib_lzfsm11
Attribute  Value
poolname   primary_lzfsm11
poolid     17fba0a9-a49b-4f16-baf5-5e0395136468
devtype    LT0
format     0xFFFFFFFF
worm       no (0)
nodegroup  G0
fillpolicy Default
owner      System
mountlimit 2
```

---

## 8.3 Changing the SSH daemon settings

The default values for `MaxSessions` and `MaxStartups` are too low and must increase to allow for successful operations with LTFSEE:

- ▶ **MaxSessions** specifies the maximum number of open sessions that is permitted per network connection. The default is 10.
- ▶ **MaxStartups** specifies the maximum number of concurrent unauthenticated connections to the SSH daemon. Additional connections are dropped until authentication succeeds or the **LoginGraceTime** expires for a connection. The default is 10:30:100.

To change **MaxSessions** and **MaxStartups** values to 1024, complete the following steps:

1. Edit the `/etc/ssh/sshd_config` file to set the `MaxSessions` and `MaxStartups` values to 1024:

```
MaxSessions = 1024
MaxStartups = 1024
```
2. Restart the `sshd` service by running the following command:

```
systemctl restart sshd.service
```

**Note:** If SSH is slow, this might indicate that several things might be wrong. Often disabling GSSAPI Authentication and reversing DNS lookup will resolve this problem and speed up SSH. Thus, set the following lines in the `sshd_config` file:

```
GSSAPIAuthentication no
UseDNS no
```

## 8.4 Setting `mmapplypolicy` options for increased performance

The default values of the `mmapplypolicy` command options must be changed when running with IBM Spectrum Archive EE. The values for these three options should be increased for enhanced performance:

▶ **-B MaxFiles**

Specifies how many files are passed for each invocation of the EXEC script. The default value is 100. If the number of files exceeds the value that is specified for `MaxFiles`, `mmapplypolicy` starts the external program multiple times.

The preferred value for IBM Spectrum Archive EE is 10000.

▶ **-m ThreadLevel**

The number of threads that are created and dispatched within each `mmapplypolicy` process during the policy execution phase. The default value is 24.

The preferred value for IBM Spectrum Archive EE is 2x the number of drives.

▶ **--single-instance**

Ensures that, for the specified file system, only one instance of `mmapplypolicy` that is started with the `--single-instance` option can run at one time. If another instance of `mmapplypolicy` is started with the `--single-instance` option, this invocation does nothing and terminates.

As a preferred practice, set the `--single_instance` option when running with IBM Spectrum Archive EE.

▶ **-s LocalWorkDirectory**

Specifies the directory to be used for temporary storage during `mmapplypolicy` command processing. The default directory is `/tmp`. The `mmapplypolicy` command stores lists of candidate and chosen files in temporary files within this directory.

When you run `mmapplypolicy`, it creates several temporary files and file lists. If the specified file system or directories contain many files, this process can require a significant amount of temporary storage. The required storage is proportional to the number of files (NF) being acted on and the average length of the path name to each file (AVPL).

To make a rough estimate of the space required, estimate NF and assume an AVPL of 80 bytes. With an AVPL of 80, the space required is roughly 300 X NF bytes of temporary space.

▶ **-N {all | mount | Node[,Node...] | NodeFile | NodeClass}**

Specifies a set of nodes to run parallel instances of policy code for better performance. The nodes must be in the same cluster as the node from which the `mmapplypolicy` command is issued. All node classes are supported.

If the `-N` option is not specified, then the command runs parallel instances of the policy code on the nodes that are specified by the `defaultHelperNodes` attribute of the `mmchconfig` command. If the `defaultHelperNodes` attribute is not set, then the list of helper nodes depends on the file system format version of the target file system. If the target file system is at file system format version 5.0.1 or later (file system format number 19.01 or later), then the helper nodes are the members of the node class `managerNodes`. Otherwise, the command runs only on the node where the `mmapplypolicy` command is issued.

**Note:** When using `-N` specify only the class defined for IBM Spectrum Archive EE nodes, except for `-I defer` or `-I prepare`.

► **-g GlobalWorkDirectory**

Specifies a global work directory in which one or more nodes can store temporary files during `mmapplypolicy` command processing. For more information about specifying more than one node to process the command, see the description of the `-N` option. For more information about temporary files, see the description of the `-s` option.

The global directory can be in the file system that `mmapplypolicy` is processing or in another file system. The file system must be a shared file system, and it must be mounted and available for reading and writing by every node that will participate in the `mmapplypolicy` command processing.

If the `-g` option is not specified, then the global work directory is the directory that is specified by the `sharedTmpDir` attribute of the `mmchconfig` command. If the `sharedTmpDir` attribute is not set to a value, then the global work directory depends on the file system format version of the target file system:

- If the target file system is at file system format version 5.0.1 or later (file system format number 19.01 or later), then the global work directory is the directory `.mmSharedTmpDir` at the root level of the target file system.
- If the target file system is at a file system format version that is earlier than 5.0.1 then the command does not use a global work directory.

If the global work directory that is specified by `-g` option or by the `sharedTmpDir` attribute begins with a forward slash (/) then it is treated as an absolute path. Otherwise it is treated as a path that is relative to the mount point of the file system or the location of the directory to be processed.

If both the `-g` option and the `-s` option are specified, then temporary files can be stored in both the specified directories. In general, the local work directory contains temporary files that are written and read by a single node. The global work directory contains temporary files that are written and read by more than one node.

If both the `-g` option and the `-N` option are specified, then `mmapplypolicy` uses high-performance, fault-tolerant protocols during execution.

**Note:** It is always a preferred practice to specify the `temp` directory to something other than `/tmp` in case the temporary files get large. This can be the case in large file systems, and the use of the IBM Spectrum Scale file system is suggested.

## 8.5 Setting the inode size for the GPFS file system for increased performance

When you create the GPFS file systems, there is an option that is called **-i InodeSize** for the **mmcrfs** command. The option specifies the byte size of inodes. The supported inode sizes are 512, 1024, and 4096 bytes. The preferred inode size is 4096 (the default for newly created IBM Spectrum Scale file systems) for all IBM Spectrum Scale file systems for IBM Spectrum Archive EE. This size should include both user data file systems and the IBM Spectrum Archive EE metadata file system.

## 8.6 Determining the file states for all files within the GPFS file system

Typically, to determine the state of a file and to which tape cartridges the file is migrated, you run the **ltfsee info files** command. However, it is not practical to run this command for every file on the GPFS file system.

In Example 8-3, the file is in the migrated state and is only on the tape cartridge JC0090JC.

*Example 8-3 Example of the ltfsee info files command*

---

```
[root@l1tfssn2 ~]# ltfsee info files /ibm/glues/LTFS_EE_FILE_v20mH1Wq_gzvnyt.bin
Name: /ibm/glues/LTFS_EE_FILE_v20mH1Wq_gzvnyt.bin
Tape id:JC0090JC   Status: migrated
```

---

Thus, use list rules in an IBM Spectrum Scale policy instead. Example 8-4 is a sample set of list rules to display files and file system objects. For those files that are in the migrated or premigrated state, the output line contains the tape cartridges on which that file is.

*Example 8-4 Sample set of list rules to display the file states*

---

```
define(
    user_exclude_list,
    (
        PATH_NAME LIKE '/ibm/glues/.ltfsee/%'
        OR PATH_NAME LIKE '%/.SpaceMan/%'
        OR PATH_NAME LIKE '%/lost+found/%'
        OR NAME = 'dsmerror.log'
    )
)

define(
    is_premigrated,
    (MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE '%V%')
)

define(
    is_migrated,
    (MISC_ATTRIBUTES LIKE '%V%')
)

define(
    is_resident,
```

```

        (NOT MISC_ATTRIBUTES LIKE '%M%')
    )

define(
    is_symlink,
    (MISC_ATTRIBUTES LIKE '%L%')
)

define(
    is_dir,
    (MISC_ATTRIBUTES LIKE '%D%')
)

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL LIST 'file_states'
EXEC '/root/file_states.sh'

RULE 'EXCLUDE_LISTS' LIST 'file_states' EXCLUDE
WHERE user_exclude_list

RULE 'MIGRATED' LIST 'file_states'
FROM POOL 'system'
SHOW('migrated ' || xattr('dmapi.IBMTPS'))
WHERE is_migrated

RULE 'PREMIGRATED' LIST 'file_states'
FROM POOL 'system'
SHOW('premigrated ' || xattr('dmapi.IBMTPS'))
WHERE is_premigrated

RULE 'RESIDENT' LIST 'file_states'
FROM POOL 'system'
SHOW('resident ')
WHERE is_resident
AND (FILE_SIZE > 0)

RULE 'SYMLINKS' LIST 'file_states'
DIRECTORIES_PLUS
FROM POOL 'system'
SHOW('symlink ')
WHERE is_symlink

RULE 'DIRS' LIST 'file_states'
DIRECTORIES_PLUS
FROM POOL 'system'
SHOW('dir ')
WHERE is_dir
AND NOT user_exclude_list

RULE 'EMPTY_FILES' LIST 'file_states'
FROM POOL 'system'
SHOW('empty_file ')
WHERE (FILE_SIZE = 0)

```

---

The policy runs a script that is named `file_states.sh`, which is shown in Example 8-5. If the policy is run daily, this script can be modified to keep several versions to be used for history purposes.

*Example 8-5 Example of file\_states.sh*

---

```
if [[ $1 == 'TEST' ]]; then
    rm -f /root/file_states.txt
elif [[ $1 == 'LIST' ]]; then
    cat $2 >> /root/file_states.txt
fi
```

---

To run the IBM Spectrum Scale policy, run the `mmappolicy` command with the `-P` option and the file states policy. This action produces a file that is called `/root/file_states.txt`, as shown in Example 8-6.

*Example 8-6 Sample output of the /root/file\_states.txt file*

---

```
58382 1061400675 0 migrated 1 JC0090JC --
/ibm/glues/LTFS_EE_FILE_yLNo6EzFmWGRV5IiwQj7Du5f48DG1h5146TY6HQ4Ua_OVfnnC.bin
58385 1928357707 0 dirs -- /ibm/glues/empty_dir_2
58388 1781732144 0 migrated 1 JC0090JC --
/ibm/glues/LTFS_EE_FILE_lnBQ1ggyQzH4QVkh54RXbX69KKdtu14HQB0u595Yx_78n50Y.bin
58389 1613275299 0 dirs -- /ibm/glues/empty_dir
58390 711115255 0 migrated 1 JC0090JC --
/ibm/glues/LTFS_EE_FILE_nJsfmm0LnmYN8kekbp4z_VJQ70j.bin
58391 386339945 0 migrated 1 JC0090JC --
/ibm/glues/LTFS_EE_FILE_0FbAueZL_9vYKRd.bin
58392 1245114627 0 dirs -- /ibm/glues/.empty_dir_3
58394 1360657693 0 migrated 1 JC0090JC --
/ibm/glues/LTFS_EE_FILE_JZg280w78VtJHG5mi25_13hZhP.bin
58395 856815733 0 migrated 1 JC0090JC --
/ibm/glues/LTFS_EE_FILE_I9UIwx4ZmVNwZsGMT03_q1RjKL.bin
```

---

In the `/root/file_states.txt` file, the files states and file system objects can be easily identified for all IBM Spectrum Scale files, including the tape cartridges where the files or file system objects are.

## 8.7 Memory considerations on the GPFS file system for increased performance

To make IBM Spectrum Scale more resistant to out of memory scenarios, adjust the `vm.min_free_kbytes` kernel tunable. This tunable controls the amount of free memory that Linux kernel keeps available (that is, not used in any kernel caches).

When `vm.min_free_kbytes` is set to its default value, some configurations might encounter memory exhaustion symptoms when free memory should in fact be available. Setting `vm.min_free_kbytes` to a higher value of 5-6% of the total amount of physical memory, up to a max of 2 GB, helps to avoid such a situation.

To modify `vm.min_free_kbytes`, complete the following steps:

1. Check the total memory of the system by running the following command:  

```
#free -k
```
2. Calculate 5-6% of the total memory in KB with a max of 2000000.
3. Add `vm.min_free_kbytes = <value from step 2>` to the `/etc/sysctl.conf` file.
4. Run `sysctl -p /etc/sysctl.conf` to permanently set the value.

## 8.8 Increasing the default maximum number of inodes in IBM Spectrum Scale

The IBM Spectrum Scale default maximum number of inodes is fine for most configurations. However, for large systems that might have millions of files or more, the maximum number of inodes to set at file system creation time might need to be changed or increased after file system creation time. The maximum number of inodes must be larger than the expected sum of files and file system objects being managed by IBM Spectrum Archive EE (including the IBM Spectrum Archive EE metadata files if there is only one GPFS file system).

Inodes are allocated when they are used. When a file is deleted, the inode is reused, but inodes are never deallocated. When setting the maximum number of inodes in a file system, there is an option to preallocate inodes. However, in most cases there is no need to preallocate inodes because by default inodes are allocated in sets as needed.

If you do decide to preallocate inodes, be careful not to preallocate more inodes than will be used. Otherwise, the allocated inodes unnecessarily consume metadata space that cannot be reclaimed.

Further considerations when managing inodes:

1. For file systems that are supporting parallel file creates, as the total number of free inodes drops below 5% of the total number of inodes, there is the potential for slowdown in the file system access. Take this situation into consideration when creating or changing your file system.
2. Excessively increasing the value for the maximum number of files might cause the allocation of too much disk space for control structures.

To view the current number of used inodes, number of free inodes, and maximum number of inodes, run the following command:

```
mmdf Device
```

To set the maximum inode limit for the file system, run the following command:

```
mmchfs Device --inode-limit MaxNumInodes[:NumInodesToPreallocate]
```

## 8.9 Configuring IBM Spectrum Scale settings for performance improvement

The performance results in 3.4.3, “Performance” on page 44 were obtained by modifying the following IBM Spectrum Scale configuration attributes to optimize IBM Spectrum Scale I/O. In most environments, only a few of the configuration attributes need to be changed. The following values were found to be optimal in our lab environment and are suitable for most environments:

- ▶ `pagepool` = 50-60% of the physical memory of the server
- ▶ `workerThreads` = 1024
- ▶ `numaMemoryInterleave` = yes
- ▶ `maxStatCache` = 0
- ▶ `maxFilesToCache` = 128k

The file system block size should be 2 MB due to a disk subsystem of eight data disks plus one parity disk with a stripe size of 256 KB.

**Note:** The `workerThread` setting is only available in IBM Spectrum Scale v4.2.1 or later. The `workerThreads` parameter controls an integrated group of variables that tune the file system performance in environments that are capable of high sequential and random read and write workloads and small file activity.

## 8.10 Real world use cases for `mmappolicy`

Typically, customers who use IBM Spectrum Archive with IBM Spectrum Scale manage one of two types of archive systems. The first is a traditional archive configuration where files are rarely accessed or updated. This configuration is intended for users who plan on keeping all their data on tape only. The second type is an active archive configuration. This configuration is more intended for users who continuously accesses the files. Each use case requires the creation of different IBM Spectrum Scale policies.

### 8.10.1 Creating a traditional archive system policy

A traditional archive system uses a single policy that scans the IBM Spectrum Archive name space for any files over 5 MB and migrates them to tape. This process saves the customer disk space immediately for new files to be generated. See “Using a cron job” on page 169 for information about how to automate the execution of this policy periodically.

**Note:** In the following policies, some optional attributes are added to provide efficient (pre)migration such as the `SIZE` attribute. This attribute specifies how many files to pass in to the `EXEC` script at a time. The preferred setting, which is listed in the following examples, is to set it to 20 GiB.

Example 8-7 shows a simple migration policy that chooses files greater than 5 MB to be candidate migration files and stubs them to tape. This is a good base policy that you can modify to your specific needs. For instance, if you need to have files on three storage pools, modify the OPTS parameter to include a third <pool>@<library>.

*Example 8-7 Simple migration file*

---

```

define(user_exclude_list,(PATH_NAME LIKE '/ibm/gpfs/.ltfsee/%' OR PATH_NAME LIKE
'/ibm/gpfs/.SpaceMan/%'))
define(is_premigrated,(MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE
'%V%'))
define(is_migrated,(MISC_ATTRIBUTES LIKE '%V%'))
define(is_resident,(NOT MISC_ATTRIBUTES LIKE '%M%'))

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'LTFSEE_FILES'
EXEC '/opt/ibm/ltfsee/bin/ltfsee'
OPTS '-p primary@lib_ltfseevm copy@lib_ltfseevm'
SIZE(20971520)

RULE 'LTFSEE_FILES_RULE' MIGRATE FROM POOL 'system'
TO POOL 'LTFSEE_FILES'
WHERE FILE_SIZE > 5242880
AND (CURRENT_TIMESTAMP - MODIFICATION_TIME > INTERVAL '5' MINUTES)
AND is_resident OR is_premigrated
AND NOT user_exclude_list

```

---

## 8.10.2 Creating active archive system policies

An active archive system requires two policies to maintain the system. The first is a premigration policy that selects all files over 5 MB to premigrate to tape, allowing users to still quickly obtain their files from disk. To see how to place this premigration policy into a cron job to run every 6 hours, see “Using a cron job” on page 169.

Example 8-8 shows a simple premigration policy for files greater than 5 MB.

*Example 8-8 Simple premigration policy for files greater than 5 MB*

---

```

define(user_exclude_list,(PATH_NAME LIKE '/ibm/gpfs/.ltfsee/%' OR PATH_NAME LIKE
'/ibm/gpfs/.SpaceMan/%'))
define(is_premigrated,(MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE
'%V%'))
define(is_migrated,(MISC_ATTRIBUTES LIKE '%V%'))
define(is_resident,(NOT MISC_ATTRIBUTES LIKE '%M%'))

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'LTFSEE_FILES'
EXEC '/opt/ibm/ltfsee/bin/ltfsee'
OPTS '-p primary@lib_ltfseevm copy@lib_ltfseevm'
SIZE(20971520)

RULE 'LTFSEE_FILES_RULE' MIGRATE FROM POOL 'system'
THRESHOLD(0,100,0)
TO POOL 'LTFSEE_FILES'

```

```
WHERE FILE_SIZE > 5242880
AND (CURRENT_TIMESTAMP - MODIFICATION_TIME > INTERVAL '5' MINUTES)
AND is_resident
AND NOT user_exclude_list
```

---

The second policy is a fail-safe policy that needs to be set so when a low disk space event is triggered, the fail-safe policy can be called. Adding the WEIGHT attribute to the policy enables the user to choose whether they want to start stubbing large files first or least recently used files. When the fail-safe policy starts running, it frees up the disk space to a set percentage.

The following are the commands for setting a fail-safe policy and calling `mmaddcallback`:

- ▶ `mmchpolicy gpfs failsafe_policy.txt`
- ▶ `mmaddcallback MIGRATION --command /usr/lpp/mmfs/bin/mmapplypolicy --event lowDiskSpace --parms "%fsName -B 20000 -m <2x the number of drives> --single-instance"`

After setting the policy with the `mmchpolicy` command, run `mmaddcallback` with the fail-safe policy. This policy runs periodically to check whether the disk space has reached the threshold where stubbing is required to free up space.

Example 8-9 shows a simple `failsafe_policy.txt`, which gets triggered when the IBM Spectrum Scale disk space reaches 80% full, and stubs least recently used files until the disk space has 50% occupancy.

*Example 8-9 failsafe\_policy.txt*

---

```
define(user_exclude_list,(PATH_NAME LIKE '/ibm/gpfs/.ltfsee/%' OR PATH_NAME LIKE
'/ibm/gpfs/.SpaceMan/%'))
define(is_premigrated,(MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE
'%V%'))
define(is_migrated,(MISC_ATTRIBUTES LIKE '%V%'))
define(is_resident,(NOT MISC_ATTRIBUTES LIKE '%M%'))

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'LTFSEE_FILES'
EXEC '/opt/ibm/ltfsee/bin/ltfsee'
OPTS '-p primary@lib_ltfsee copy@lib_ltfsee copy2@lib_ltfsee'
SIZE(20971520)

RULE 'LTFSEE_FILES_RULE' MIGRATE FROM POOL 'system'
THRESHOLD(80,50)
WEIGHT(CURRENT_TIMESTAMP - ACCESS_TIME)
TO POOL 'LTFSEE_FILES'
WHERE FILE_SIZE > 5242880
AND is_premigrated
AND NOT user_exclude_list
```

---

### 8.10.3 IBM Spectrum Archive EE migration policy with AFM

For customers using IBM Spectrum Archive EE with IBM Spectrum Scale AFM, the migration policy would need to change to accommodate the additional exclude directories wherever migrations are occurring. Example 8-10 uses the same migration policy that is shown in Example 8-7 on page 236 with the addition of extra exclude and check parameters.

*Example 8-10 Updated migration policy to include AFM*

---

```
define(user_exclude_list,(PATH_NAME LIKE '/ibm/gpfs/.ltfsee/%' OR PATH_NAME LIKE
'/ibm/gpfs/.SpaceMan/%' OR PATH_NAME LIKE '%/.snapshots/%' OR PATH_NAME LIKE
'/ibm/gpfs/fset1/.afm/%' OR PATH_NAME LIKE '/ibm/gpfs/fset1/.ptrash/%'))

define(is_premigrated,(MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE
'%V%'))

define(is_migrated,(MISC_ATTRIBUTES LIKE '%V%'))

define(is_resident,(NOT MISC_ATTRIBUTES LIKE '%M%'))

define(is_cached,(MISC_ATTRIBUTES LIKE '%u%'))

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'LTFSEE_FILES'
EXEC '/opt/ibm/ltfsee/bin/ltfsee'
OPTS '-p primary@lib_ltfseevm copy@lib_ltfseevm'
SIZE(20971520)

RULE 'LTFSEE_FILES_RULE' MIGRATE FROM POOL 'system'
THRESHOLD(0,100,0)
TO POOL 'LTFSEE_FILES'
WHERE FILE_SIZE > 5242880
AND (CURRENT_TIMESTAMP - MODIFICATION_TIME > INTERVAL '5' MINUTES)
AND is_resident
AND is_cached
AND NOT user_exclude_list
```

---

## 8.11 Capturing a core file on RHEL with abrt

The Automatic Bug Reporting Tool (ABRT) consists of the **abrt** daemon and a number of system services and utilities to process, analyze, and report detected problems. The daemon runs silently in the background most of the time, and springs into action when an application crashes or a kernel fault is detected. The daemon then collects the relevant problem data, such as a core file if there is one, the crashing application's command-line parameters, and other data of forensic utility.

For `abrt` to work with IBM Spectrum Archive EE, two configuration directives must be modified in the `/etc/abrt/abrt-action-save-package-data.conf` file:

- ▶ `OpenGPGCheck = yes/no`

Setting the `OpenGPGCheck` directive to `yes`, which is the default setting, tells ABRT to analyze and handle only crashes in applications that are provided by packages that are signed by the GPG keys, which are listed in the `/etc/abrt/gpg_keys` file. Setting `OpenGPGCheck` to `no` tells ABRT to detect crashes in all programs.

- ▶ `ProcessUnpackaged = yes/no`

This directive tells ABRT whether to process crashes in executable files that do not belong to any package. The default setting is `no`.

Here are the preferred settings:

`OpenGPGCheck = no`

`ProcessUnpackaged = yes`

## 8.12 Antivirus considerations

Although in-depth testing occurs with IBM Spectrum Archive EE and many industry-leading antivirus software programs, there are a few considerations to review periodically:

- ▶ Configure any antivirus software to exclude IBM Spectrum Archive EE and HSM work directories:
  - The library mount point (the `/l tfs` directory).
  - All IBM Spectrum Archive EE space-managed GPFS file systems (which includes the `.SPACEMAN` directory).
  - The IBM Spectrum Archive EE metadata directory (the GPFS file system that is reserved for IBM Spectrum Archive EE internal usage).
- ▶ Use antivirus software that supports sparse or offline files. Be sure that it has a setting that allows it to skip offline or sparse files to avoid unnecessary recall of migrated files.

## 8.13 Automatic email notification with rsyslog

Rsyslog and its mail output module (`ommail`) can be used to send syslog messages from IBM Spectrum Archive EE through email. Each syslog message is sent through its own email. Users should pay special attention to applying the correct amount of filtering to prevent heavy spamming. The `ommail` plug-in is primarily meant for alerting users of certain conditions and should be used in a limited number of cases. For more information, see the rsyslog `ommail` home page:

<http://www.rsyslog.com/doc/ommail.html>

Here are two examples of how rsyslog `ommail` can be used with IBM Spectrum Archive EE by modifying the `/etc/rsyslog.conf` file:

- ▶ If users want to send an email on all IBM Spectrum Archive EE registered error messages, the regular expression is `"GLES[A-Z][0-9]*E"`, as shown in Example 8-11.

*Example 8-11 Email for all IBM Spectrum Archive EE registered error messages*

---

```
$ModLoad ommail
$ActionMailSMTPServer us.ibm.com
```

```
$ActionMailFrom ltfsee@ltfsee_host1.tuc.stglabs.ibm.com
$ActionMailTo ltfsee_user1@us.ibm.com
$template mailSubject,"LTFS EE Alert on %hostname%"
$template mailBody,"%msg%"
$ActionMailSubject mailSubject
:msg, regex, "GLES[A-Z][0-9]*E" :ommail::mailBody
```

---

- ▶ If users want to send an email on any failed migration per scan list, the regular expression is “[1-9][0-9]\* failed”, as shown in Example 8-12.

*Example 8-12 Email for any failed migration per scan list*

---

```
$ModLoad ommail
$ActionMailSMTPServer us.ibm.com
$ActionMailFrom ltfsee@ltfsee_host1.tuc.stglabs.ibm.com
$ActionMailTo ltfsee_user1@us.ibm.com
$template mailSubject,"LTFS EE Alert on %hostname%"
$template mailBody,"%msg%"
$ActionMailSubject mailSubject
:msg, regex, "[1-9][0-9]* failed" :ommail::mailBody
```

---

## 8.14 Overlapping IBM Spectrum Scale policy rules

This section describes how you can avoid migration failures during your IBM Spectrum Archive EE system operations by having only non-overlapping IBM Spectrum Scale policy rules in place.

After a file is migrated to a tape cartridge pool and is in the migrated state, it cannot be migrated to other tape cartridge pools (unless it is recalled back from physical tape to file system space).

Do not use overlapping IBM Spectrum Scale policy rules within different IBM Spectrum Scale policy files that can select the same files for migration to different tape cartridge pools. If a file was migrated, a later migration fails. The migration result for any file that already is in the migrated state is fail.

In Example 8-13, an attempt is made to migrate four files to tape cartridge pool pool2. Before the migration attempt, Tape ID 2MA260L5 is already in tape cartridge pool pool1, and Tape ID 2MA262L5 has one migrated and one pre-migrated file. The state of the files on these tape cartridges before the migration attempt is shown by the **ltfsee info** command in Example 8-13.

*Example 8-13 Display the state of files by using the ltfsee info command*

---

```
# ltfsee info files *.ppt
Name: fileA.ppt
Tape id:2MA260L5@lib_lto Status: migrated
Name: fileB.ppt
Tape id:2MA260L5@lib_lto Status: migrated
Name: fileC.ppt
Tape id:2MA262L5@lib_lto Status: migrated
Name: fileD.ppt
Tape id:2MA262L5@lib_lto Status: premigrated
```

---

The mig scan list file that is used in this example contains these entries, as shown in Example 8-14.

*Example 8-14 Sample content of a scan list file*

---

```
-- /ibm/gpfs/fileA.ppt
-- /ibm/gpfs/fileB.ppt
-- /ibm/gpfs/fileC.ppt
-- /ibm/gpfs/fileD.ppt
```

---

The attempt to migrate the files produces the results that are shown in Example 8-15.

*Example 8-15 Migration of files by running the ltfsee migration command*

---

```
# ltfsee migrate mig -p pool2@lib_lto
GLES167I(00400): A list of files to be migrated has been sent to LTFS EE using
scan id 3428780289.
GLES159E(00440): Not all migration has been successful.
GLES1038I(00448): Migration result: 1 succeeded, 3 failed, 0 duplicate, 0
duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too
early for migration.
```

---

The files on Tape ID 2MA260L5 (fileB.ppt and fileD.pdf) already are in tape cartridge pool pool1. Therefore, the attempt to migrate them to tape cartridge pool pool2 produces a migration result failed.

For the files on Tape ID 2MA262L5, the attempt to migrate fileB.ppt file also produces a migration result fail because the file is already migrated. Only the attempt to migrate the pre-migrated fileA.ppt file succeeds. Therefore, one operation succeeded and three other operations failed.

## 8.15 Storage pool assignment

This section describes how you can facilitate your IBM Spectrum Archive EE system export activities by using different storage pools for logically different parts of an IBM Spectrum Scale namespace.

If you put different logical parts of an IBM Spectrum Scale namespace (such as the project directory) into different LTFS tape cartridge pools, you can Normal Export tape cartridges that contain only the files from that specific part of the IBM Spectrum Scale namespace (such as project abc). Otherwise, you must first recall all the files from the namespace of interest (such as the project directory of all projects), migrate the recalled files to an empty tape cartridge pool, and then Normal Export that tape cartridge pool.

The concept of different tape cartridge pools for different logical parts of an IBM Spectrum Scale namespace can be further isolated by using IBM Spectrum Archive node groups. A node group consists of one or more nodes that are connected to the same tape library. When tape cartridge pools are created, they can be assigned to a specific node group. For migration purposes, it allows certain tape cartridge pools to be used with only drives within the owning node group.

## 8.16 Tape cartridge removal

This section describes the information that must be reviewed before you physically remove a tape cartridge from the library of your IBM Spectrum Archive EE environment.

For more information, see 7.7.2, “Moving tape cartridges” on page 149, and “The `ltfsee premigrate` command” on page 316.

### 8.16.1 Reclaiming tape cartridges before you remove or export them

If tape cartridges are going bad, reclaim tape cartridges before you remove or Normal Export them. To avoid failed recall operations, reclaim a tape cartridge using one of these methods:

- ▶ Running the `ltfsee reclaim` command before you remove it from the LTFS file system (by running the `ltfsee tape move` command)
- ▶ Normal Export it from the LTFS library (by running the `ltfsee export` command without the `--offline` option)

### 8.16.2 Exporting tape cartridges before physically removing them from the library

A preferred practice is always to export a tape cartridge before it is physically removed from the library. If a removed tape cartridge is modified and then reinserted in the library, unpredictable behavior can occur.

## 8.17 Reusing LTFS formatted tape cartridges

In some scenarios, you might want to reuse tape cartridges for your IBM Spectrum Archive EE setup, which were used before as an LTFS formatted media in another LTFS setup.

Because these tape cartridges still might contain data from the previous usage, IBM Spectrum Archive EE recognizes the old content because LTFS is a self-describing format.

Before such tape cartridges can be reused within your IBM Spectrum Archive EE environment, they must be reformatted before they are added to an IBM Spectrum Archive EE tape cartridge pool. This task can be done by running the `ltfsee operation` commands (`reconcile` or `reclaim` and an implicit format through the `add pool` command).

### 8.17.1 Reformatting LTFS tape cartridges through `ltfsee` commands

If a tape cartridge was used as an LTFS tape, you can check its contents after it is added to the IBM Spectrum Archive EE system and loaded to a drive. You can run the `ls -la` command to display content of the tape cartridge, as shown in Example 8-16.

*Example 8-16 Display content of a used LTFS tape cartridge (non-IBM Spectrum Archive EE)*

```
[root@ltfs97 ~]# ls -la /ltfs/153AGWL5
total 41452613
drwxrwxrwx  2 root root          0 Jul 12 2012 .
drwxrwxrwx 12 root root          0 Jan  1 1970 ..
-rwxrwxrwx  1 root root    18601 Jul 12 2012 api_test.log
-rwxrwxrwx  1 root root    50963 Jul 11 2012 config.log
```

```
-rwxrwxrwx 1 root root 1048576 Jul 12 2012 dummy.000
-rwxrwxrwx 1 root root 21474836480 Jul 12 2012 perf_fcheck.000
-rwxrwxrwx 1 root root 20971520000 Jul 12 2012 perf_migrec
lrwxrwxrwx 1 root root 25 Jul 12 2012 symfile ->
/Users/piste/mnt/testfile
```

---

You can also discover if it was an IBM Spectrum Archive EE tape cartridge before or just a standard LTFS tape cartridge that is used by IBM Spectrum Archive LE or IBM Spectrum Archive SDE release. Review the hidden directory `.LTFSEE_DATA`, as shown in Example 8-17. This example indicates that this cartridge was previously used as an IBM Spectrum Archive EE tape cartridge.

*Example 8-17 Display content of a used LTFS tape cartridge (IBM Spectrum Archive EE)*

```
[root@ltfs97 ~]# ls -la /ltfs/051AGWL5
total 0
drwxrwxrwx 4 root root 0 Apr  9 18:00 .
drwxrwxrwx 11 root root 0 Jan  1 1970 ..
drwxrwxrwx 3 root root 0 Apr  9 18:00 ibm
drwxrwxrwx 2 root root 0 Apr  9 18:00 .LTFSEE_DATA
```

---

The procedure for reuse and reformatting of a previously used LTFS tape cartridge depends on whether it was used before as an IBM Spectrum Archive LE or IBM Spectrum Archive SDE tape cartridge or as an IBM Spectrum Archive EE tape cartridge.

Before you start with the reformat procedures and examples, it is important that you confirm the following starting point. You can see the tape cartridges that you want to reuse by running the `ltfsee info tapes` command in the status Unknown, as shown in Example 8-18.

*Example 8-18 Output of the ltfsee info tapes command*

```
[root@ltfs97 ~]# ltfsee info tapes
```

Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)	Pool	Library	Address	Drive	Appendable
TST200L6	Valid	L6	2242	0	2242	0	test1	TS3310	4117	-	yes
LTFS01L7	Valid	L7	5338	0	5338	0	test	TS3310	4119	-	yes
PFE001L6	Valid	L6	2242	95	2146	0	test1	TS3310	4118	-	yes

---

## Reformatting and reusing an LTFS SDE/LE tape cartridge

In this case, you run the `ltfsee pool` command to add this tape cartridge to an IBM Spectrum Archive EE tape cartridge pool and format it at the same time with the `-F` option, as shown in the following example and in Example 8-19:

```
ltfsee pool add myfirstpool <tapeID> -F
```

*Example 8-19 Reformat a used LTFS SDE/LE tape cartridge*

```
[root@ltfs97 ~]# ltfsee pool add myfirstpool 037AGWL5 -F
Tape 037AGWL5 successfully formatted.
Adding tape 037AGWL5 to storage pool myfirstpool
```

---

## Reformatting and reusing an IBM Spectrum Archive EE tape cartridge

If you want to reuse an IBM Spectrum Archive EE tape cartridge, the `ltfsee pool add -p myfirstpool <tapeID> -F` command does not work. It does not work because IBM Spectrum Archive EE detects that it was a previously used IBM Spectrum Archive EE tape cartridge.

Before you use the command for adding it to a pool and use the **-F** format option, you must resolve and clean the relationships from the old files on the tape cartridge to your IBM Spectrum Archive EE environment. You can do this task by running the **ltfsee reconcile -t** command, first against the specific tape cartridge, as shown in the following example and in Example 8-20:

```
ltfsee reconcile -t <tapeID>
ltfsee pool add -p myfirstpool <tapeID> -F
```

*Example 8-20 Reconcile and reformat a used IBM Spectrum Archive EE tape cartridge*

---

```
[root@ltfs97 ~]# ltfsee reconcile -t 058AGWL5 -p myfirstpool
GLESS016I(00109): Reconciliation requested
GLESS049I(00610): Tapes to reconcile: 058AGWL5
GLESS050I(00619): GPFS filesystems involved: /ibm/glues
GLESS053I(00647): Number of pending migrations: 0
GLESS054I(00651): Creating GPFS snapshots:
GLESS055I(00656): Creating GPFS snapshot for /ibm/glues ( /dev/gpfs )
GLESS056I(00724): Scanning GPFS snapshots:
GLESS057I(00728): Scanning GPFS snapshot of /ibm/glues ( /dev/gpfs )
GLESS058I(00738): Removing GPFS snapshots:
GLESS059I(00742): Removing GPFS snapshot of /ibm/glues ( /dev/gpfs )
GLESS060I(00760): Processing scan results:
GLESS061I(00764): Processing scan results for /ibm/glues ( /dev/gpfs )
GLESS063I(00789): Reconciling the tapes:
GLESS001I(00815): Reconciling tape 058AGWL5 has been requested
GLESS002I(00835): Reconciling tape 058AGWL5 complete
GLESL172I(02984): Synchronizing LTFS EE tapes information
[root@ltfs97 ~]# ltfsee pool add -p myfirstpool 058AGWL5 -F
Tape 058AGWL5 successfully formatted.
Adding tape 058AGWL5 to storage pool myfirstpool
```

---

## 8.18 Reusing non-LTFS tape cartridges

In some scenarios, you might want to reuse tape cartridges for your IBM Spectrum Archive EE setup, which were used before as non-LTFS formatted media in another server setup behind your tape library (such as backup tape cartridges from a Tivoli Storage Manager Server or an IBM Spectrum Protect environment).

Although these tape cartridges still might contain data from the previous usage, they can be used within IBM Spectrum Archive EE the same way as new, unused tape cartridges. For more information about how to add new tape cartridge media to an IBM Spectrum Archive EE tape cartridge pool, see 7.7.1, “Adding tape cartridges” on page 147.

## 8.19 Moving tape cartridges between pools

This section describes preferred practices to consider when you want to move a tape cartridge between tape cartridge pools. This information also relates to the function that is described in 7.7.2, “Moving tape cartridges” on page 149.

### **8.19.1 Avoiding changing assignments for tape cartridges that contain files**

If a tape cartridge contains any files, a preferred practice is to not move the tape cartridge from one tape cartridge pool to another tape cartridge pool. If you remove the tape cartridge from one tape cartridge pool and then add it to another tape cartridge pool, the tape cartridge includes files that are targeted for multiple pools.

Before you export files you want from that tape cartridge, you must recall any files that are not supposed to be exported in such a scenario.

For more information, see 7.8, “Tape storage pool management” on page 154.

### **8.19.2 Reclaiming a tape cartridge and changing its assignment**

Before you remove a tape cartridge from one tape cartridge pool and add it to another tape cartridge pool, a preferred practice is to reclaim the tape cartridge so that no files remain on the tape cartridge when it is removed. This action prevents the scenario that is described in 8.19.1, “Avoiding changing assignments for tape cartridges that contain files” on page 245.

For more information, see 7.8, “Tape storage pool management” on page 154 and 7.19, “Reclamation” on page 192.

## **8.20 Offline tape cartridges**

This section describes how you can help maintain the file integrity of offline tape cartridges by not modifying the files of offline exported tape cartridges. Also, a reference to information about solving import problems that are caused by modified offline tape cartridges is provided.

### **8.20.1 Do not modify the files of offline tape cartridges**

When a tape cartridge is offline and outside the library, do not modify its IBM Spectrum Scale offline files on disk and do not modify its files on the tape cartridge. Otherwise, some files that exist on the tape cartridge might become unavailable to IBM Spectrum Scale.

### **8.20.2 Solving problems**

For more information about solving problems that are caused by trying to import an offline exported tape cartridge that was modified while it was outside the library, see “Importing offline tape cartridges” on page 198.

## **8.21 Scheduling reconciliation and reclamation**

This section provides information about scheduling regular reconciliation and reclamation activities.

The reconciliation process resolves any inconsistencies that develop between files in the IBM Spectrum Scale and their equivalents in LTFS. The reclamation function frees up tape cartridge space that is occupied by non-referenced files and non-referenced content that is present on the tape cartridge. In other words, this is inactive content of data, but still occupying space on the physical tape.

It is preferable to schedule periodic reconciliation and reclamation, ideally during off-peak hours and at a frequency that is most effective. A schedule helps ensure consistency between files and efficient use of the tape cartridges in your IBM Spectrum Archive EE environment.

For more information, see 7.15, “Repairing files to their resident state” on page 187 and 7.19, “Reclamation” on page 192.

## 8.22 License Expiration Handling

License validation is done by the IBM Spectrum Archive EE program. If the license covers only a certain period (as in the case for the IBM Spectrum Archive EE Trial Version, which is available for three months), it expires if this time is passed. The behavior of IBM Spectrum Archive EE changes after that period in the following cases:

- ▶ The state of the nodes changes to the following defined value:

`NODE_STATUS_LICENSE_EXPIRED`

The node status can be determined through an internal node command, which is also used in other parts of IBM Spectrum Archive EE.

- ▶ When the license is expired, IBM Spectrum Archive EE can still read data, but it is impossible to write and migrate data. In such a case, not all IBM Spectrum Archive EE commands are usable.

When the license is expired and detected by the scheduler of the main IBM Spectrum Archive EE management components (MMM), it shuts down. This feature is necessary to have a proper clean-up if some jobs are still running or unscheduled. By doing so, a user is aware that IBM Spectrum Archive EE does not function because of the license expiration.

To give a user the possibility to access files that were previously migrated to tape, it is possible for IBM Spectrum Archive EE to restart, but it operates with limited functions. All functions that write to tape cartridges are not available. During the start of IBM Spectrum Archive EE (through MMM), it is detected that some nodes have the status of `NODE_STATUS_LICENSE_EXPIRED`.

IBM Spectrum Archive EE fails the following commands immediately:

- ▶ **migrate**
- ▶ **import**
- ▶ **export**
- ▶ **reconcile**
- ▶ **reclaim**

These commands are designed to write to a tape cartridge in certain cases. Therefore, they fail with an error message. The transparent access of a migrated file is not affected. The deletion of the link and the data file on a tape cartridge because of a write or truncate recall is omitted.

In summary, the following steps occur after expiration:

1. The status of the nodes changes to the state `NODE_STATUS_LICENSE_EXPIRED`.
2. IBM Spectrum Archive EE shuts down to allow a proper clean-up.
3. IBM Spectrum Archive EE can be started again with limited functions.

## 8.23 Disaster recovery

This section describes the preparation of an IBM Spectrum Archive EE DR setup and the steps that you must perform before and after a disaster to recover your IBM Spectrum Archive EE environment.

### 8.23.1 Tiers of disaster recovery

Understanding DR strategies and solutions can be complex. To help categorize the various solutions and their characteristics (for example, costs, recovery time capabilities, and recovery point capabilities), definitions of the various levels and required components can be defined. The idea behind such a classification is to help those concerned with DR to determine the following issues:

- ▶ What solution they have
- ▶ What solution they require
- ▶ What it requires to meet greater DR objectives

In 1992, the SHARE user group in the United States, along with IBM, defined a set of DR tier levels. This action was done to address the need to describe and quantify various different methodologies for successful mission-critical computer systems DR implementations. So, within the IT Business Continuance industry, the tier concept continues to be used, and is useful for describing today's DR capabilities.

The tiers' definitions are designed so that emerging DR technologies can also be applied, as shown in Table 8-1.

Table 8-1 Summary of disaster recovery tiers (SHARE)

Tier	Description
6	Zero data loss
5	Two-site two-phase commit
4	Electronic vaulting to hotsite (active secondary site)
3	Electronic vaulting
2	Offsite vaulting with a hotsite (PTAM + hot site)
1	Offsite vaulting (Pickup Truck Access Method (PTAM))
0	Offsite vaulting (PTAM)

In the context of the IBM Spectrum Archive EE product, this section focuses only on the tier 1 strategy because this is the only supported solution that you can achieve with a product that handles physical tape media (offsite vaulting).

For more information about the other DR tiers and general strategies, see *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844.

#### Tier 1: Offsite vaulting

A tier 1 installation is defined as having a disaster recovery plan (DRP) that backs up and stores its data at an offsite storage facility, and determines some recovery requirements. As shown in Figure 8-1 on page 248, backups are taken that are stored at an offsite storage facility.

This environment can also establish a backup platform, although it does not have a site at which to restore its data, nor the necessary hardware on which to restore the data, such as compatible tape devices.

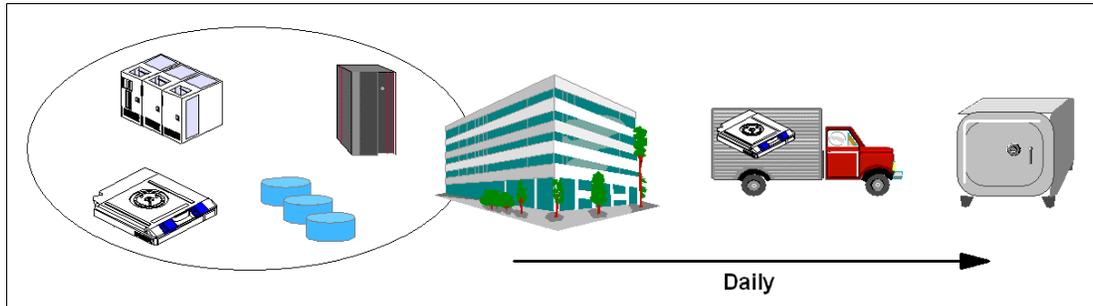


Figure 8-1 Tier 1 - offsite vaulting (PTAM)

Because vaulting and retrieval of data is typically handled by couriers, this tier is described as the PTAM. PTAM is a method that is used by many sites because this is a relatively inexpensive option. However, it can be difficult to manage because it is difficult to know exactly where the data is at any point.

There is probably only selectively saved data. Certain requirements were determined and documented in a contingency plan and there is optional backup hardware and a backup facility that is available. Recovery depends on when hardware can be supplied, or possibly when a building for the new infrastructure can be located and prepared.

Although some customers are on this tier and seemingly can recover if there is a disaster, one factor that is sometimes overlooked is the recovery time objective (RTO). For example, although it is possible to recover data eventually, it might take several days or weeks. An outage of business data for this long can affect business operations for several months or even years (if not permanently).

**Important:** With IBM Spectrum Archive EE, the recovery time can be improved because after the import of the vaulting tape cartridges into a recovered production environment, the user data is immediately accessible without the need to copy back content from the tape cartridges into a disk or file system.

### 8.23.2 Preparing IBM Spectrum Archive EE for a tier 1 disaster recovery strategy (offsite vaulting)

IBM Spectrum Archive EE has all the tools and functions that you need to prepare a tier 1 DR strategy for offsite vaulting of tape media.

The fundamental concept is based on the IBM Spectrum Archive EE function to create replicas and redundant copies of your file system data to tape media during migration (see 7.10.4, “Replicas and redundant copies” on page 170). IBM Spectrum Archive EE enables the creation of a replica plus two more redundant replicas (copies) of each IBM Spectrum Scale file during the migration process.

The first replica is the primary copy, and other replicas are called *redundant copies*. Redundant copies must be created in tape cartridge pools that are different from the tape cartridge pool of the primary replica and different from the tape cartridge pools of other redundant copies.

Up to two redundant copies can be created, which means that a specific file from the GPFS file system can be stored on three different physical tape cartridges in three different IBM Spectrum Archive EE tape cartridge pools.

The tape cartridge where the primary replica is stored and the tapes that contain the redundant copies are referenced in the IBM Spectrum Scale inode with an IBM Spectrum Archive EE DMAPI attribute. The primary replica is always listed first.

Redundant copies are written to their corresponding tape cartridges in the IBM Spectrum Archive EE format. These tape cartridges can be reconciled, exported, reclaimed, or imported by using the same commands and procedures that are used for standard migration without replica creation.

Redundant copies must be created in tape cartridge pools that are different from the pool of the primary replica and different from the pools of other redundant copies. Therefore, create a DR pool named DRPool that exclusively contains the media you plan to Offline Export for offline vaulting. You must also plan for the following issues:

- ▶ Which file system data is migrated (as another replica) to the DR pool?
- ▶ How often do you plan to export and remove of physical tapes for offline vaulting?
- ▶ How do you handle media lifecycle management with the tape cartridges for offline vaulting?
- ▶ What are the DR steps and procedure?

If the primary replica of the IBM Spectrum Archive EE server and IBM Spectrum Scale do not exist due to a disaster, the redundant copy that is created and stored in an external site (offline vaulting) will be used for the disaster recovery.

Section 8.23.3, “IBM Spectrum Archive EE tier 1 DR procedure” on page 250 describes the steps that are used to perform these actions:

- ▶ Recover (import) the offline vaulting tape cartridges to a newly installed IBM Spectrum Archive EE environment
- ▶ Re-create the GPFS file system information
- ▶ Regain access to your IBM Spectrum Archive EE data

**Important:** The migration of a pre-migrated file does not create new replicas.

Example 8-21 shows you a sample migration policy to migrate all files to three pools. To have this policy run periodically, see “Using a cron job” on page 169.

*Example 8-21 Sample of a migration policy*

```
define(user_exclude_list,(PATH_NAME LIKE '/ibm/gpfs/.ltfsee/%' OR PATH_NAME LIKE
'/ibm/gpfs/.SpaceMan/%'))
define(is_premigrated,(MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE
'%V%'))
define(is_migrated,(MISC_ATTRIBUTES LIKE '%V%'))
define(is_resident,(NOT MISC_ATTRIBUTES LIKE '%M%'))

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'LTFSEE_FILES'
EXEC '/opt/ibm/ltfsee/bin/ltfsee'
OPTS '-p primary@lib_ltfseevm copy@lib_ltfseevm DR@lib_ltfseevm'
```

```
SIZE(20971520)

RULE 'LTFSEE_FILES_RULE' MIGRATE FROM POOL 'system'

TO POOL 'LTFSEE_FILES'

AND is_premigrated
AND NOT user_exclude_list
```

---

After you create redundant replicas of your file system data on different IBM Spectrum Archive EE tape cartridge pools for offline vaulting, you can Normal Export the tape cartridges by running the IBM Spectrum Archive EE export command. For more information, see 7.21.2, “Exporting” on page 199.

**Important:** The IBM Spectrum Archive EE export command does not eject the tape cartridge to the physical I/O station of the attached tape library. To eject the DR tape cartridges from the library to take them out for offline vaulting, you can run the **ltfsee tape move** command with the option **ieslot**. For more information, see 7.7, “Tape library management” on page 147 and 11.1, “Command-line reference” on page 300.

### 8.23.3 IBM Spectrum Archive EE tier 1 DR procedure

To perform a DR to restore an IBM Spectrum Archive EE server and IBM Spectrum Scale with tape cartridges from offline vaulting, complete the following steps:

1. Before you start a DR, a set of IBM Spectrum Archive EE exported tape cartridges for an Offline vault must be created. In addition, a “new” Linux server and an IBM Spectrum Archive EE cluster environment, including IBM Spectrum Scale, must be set up.
2. Confirm that the new installed IBM Spectrum Archive EE cluster is running and ready for the import operation by running the following commands:
  - # **ltfsee status**
  - # **ltfsee info tapes**
  - # **ltfsee info pools**
3. Insert the tape cartridges for DR into the tape library I/O station.
4. Use your tape library management GUI to assign the DR tape cartridges to the IBM Spectrum Archive EE logical tape library partition of your new IBM Spectrum Archive EE server.
5. From the IBM Spectrum Archive EE program, retrieve the updated inventory information from the logical tape library by running the following command:

```
# ltfsee retrieve
```
6. Import the DR tape cartridges into the IBM Spectrum Archive EE environment by running the **ltfsee rebuild** command. The **ltfsee rebuild** command features various options that you can specify. Therefore, it is important to become familiar with these options, especially when you are performing DR. For more information, see Chapter 11, “Reference” on page 299.

When you rebuild from one or more tape cartridges, the **ltfsee rebuild** command adds the specified tape cartridge to the IBM Spectrum Archive EE library and imports the files on that tape cartridge into the IBM Spectrum Scale namespace.

This process puts the stub file back in to the IBM Spectrum Scale namespace, but the imported files stay in a migrated state, which means that the data remains on tape. The data portion of the file is not copied to disk during the import.

Rebuild from tape cartridges not included in any tape cartridge pool cannot be targeted for migration. If you want a rebuild from a tape to be available for migration, you must assign it to a tape cartridge pool.

## Restoring file system objects and files from tape

If a GPFS file system fails, the migrated files and the saved file system objects (empty regular files, symbolic links, and empty directories) can be restored from tapes by running the **ltfsee rebuild** command.

The **ltfsee rebuild** command operates similarly to the **ltfsee import** command. The **ltfsee rebuild** command reinstantiates the stub files in IBM Spectrum Scale for migrated files. The state of those files changes to the migrated state. Additionally the **ltfsee rebuild** command re-creates the file system objects in IBM Spectrum Scale for saved file system objects.

**Note:** When a symbolic link is saved to tape and then restored by the **ltfsee rebuild** command, the target of the symbolic link is kept. However, this process might cause the link to break. Therefore, after a symbolic link is restored, it might need to be moved manually to its original location on IBM Spectrum Scale.

### *Typical recovery procedure by using the ltfsee rebuild command*

Here is a typical user scenario for recovering migrated files and saved file system objects from tape by running the **ltfsee rebuild** command:

1. Re-create the GPFS file system or create a GPFS file system.
2. Restore the migrated files and saved file system objects from tape by running the **ltfsee rebuild** command:

```
ltfsee rebuild -P /gpfs/ltfsee/rebuild -p PrimPool -t LTFS01L6 LTFS02L6  
LTFS03L6
```

`/gpfs/ltfsee/rebuild` is a directory in IBM Spectrum Scale to be restored to, `PrimPool` is the storage pool to import the tapes into, and `LTFS01L6`, `LTFS02L6`, and `LTFS02L6` are tapes that contain migrated files or saved file system objects.

### *Rebuild processing for unexported tapes that are not reconciled*

The **ltfsee rebuild** command might encounter tapes that are not reconciled when the command is applied to tapes that are not exported from IBM Spectrum Archive EE. In this case, the following situations can occur with the processing to restore files and file system objects, and should be handled as described for each case:

- ▶ The tapes might have multiple generations of a file or a file system object. If so, the **ltfsee rebuild** command restores an object from the latest one that is on the tapes that are specified by the command.
- ▶ The tapes might not reflect the latest file information from IBM Spectrum Scale. If so, the **ltfsee rebuild** command restores files or file system objects that were removed from IBM Spectrum Scale.

### **Rebuild and restore considerations**

Observe these additional considerations when you run the `ltfsee rebuild` command:

- ▶ While the `ltfsee rebuild` command is running, do not modify or access the files or file system objects to be restored. During the rebuild process, an old generation of the file can appear on IBM Spectrum Scale.
- ▶ Avoid running the `ltfsee rebuild` command with any tape that is being used by IBM Spectrum Archive EE. Otherwise, significant errors are likely to occur because two stub files are created in IBM Spectrum Scale that are migrated to one tape.

**Note:** Observe the following advice if you run the `ltfsee rebuild` command with a tape in use by IBM Spectrum Archive EE.

You must resolve the situation (two stub files for one migrated file) before any other operations are made against the files or file system objects that correspond to the tape.

For example, consider the case where a file is migrated to a tape, and then the file is accidentally removed from IBM Spectrum Scale. To restore the file from the tape involves these steps:

1. Run the `ltfsee rebuild` command with the tape, which rebuilds all files and file system objects in the tape to a specified directory.
2. Choose the file to be restored and move it to the original location in IBM Spectrum Scale.
3. Remove all other files and file system objects in the specified directory before any actions (migration, recall, save, reconcile, reclaim, import, export, or rebuild) are made against the tape.

**Important:** With the first release of IBM Spectrum Archive EE, access control list (ACL) file system information is *not* supported and not recovered if you rebuild the IBM Spectrum Scale data after a disaster through the replica tape cartridges. All recovered files under IBM Spectrum Scale have a generic permission setting, as shown in the following example:

```
-rw----- 1 root root 104857600 Apr 16 14:51 file1.img
```

The IBM Spectrum Scale Extended Attributes (EAs) are supported by IBM Spectrum Archive EE and are restored. These EAs are imported because they are used by various applications to store more file information. For more information about limitations, see IBM Spectrum Archive Library Edition *Extended Attributes*:

[https://www.ibm.com/support/knowledgecenter/en/STZMZN\\_2.4.0/ltfs\\_hints\\_extended\\_attribute\\_linux.html#ltfs\\_hints\\_extended\\_attribute\\_linux](https://www.ibm.com/support/knowledgecenter/en/STZMZN_2.4.0/ltfs_hints_extended_attribute_linux.html#ltfs_hints_extended_attribute_linux)

## **8.24 IBM Spectrum Archive EE problem determination**

If you discover an error message or a problem while you are running and operating the IBM Spectrum Archive EE program, you can check the IBM Spectrum Archive EE log file as a starting point for problem determination.

The IBM Spectrum Archive EE log file can be found in the following directory:

```
/var/log/ltfsee.log
```

In Example 8-22, we attempted to migrate two files (document10.txt and document20.txt) to a pool (myfirstpool) that contained two new formatted and added physical tapes (055AGWL5 and 055AGWL5). We encountered an error that only one file was migrated successfully. We checked the ltfsee.log to determine why the other file was not migrated.

*Example 8-22 Check the ltfsee.log file*

---

```
GLESL167I(00400): A list of files to be migrated has been sent to LTFS EE using scan id
1887703553.
GLESL159E(00440): Not all migration has been successful.
GLESL038I(00448): Migration result: 1 succeeded, 1 failed, 0 duplicate, 0 duplicate wrong
pool, 0 not found, 0 too small to qualify for migration, 0 too early for migration.
[root@ltfs97 glues]#
[root@ltfs97 glues]# vi /var/log/ltfsee.log
2016-12-14T09:17:08.538724-07:00 ltfs97 mmm[7889]: GLESM148E(00538): File
/ibm/gpfs/document20.txt is already migrated and will be skipped.
2016-12-14T09:17:18.540054-07:00 ltfs97 ltfsee[1948]: GLESL159E(00440): Not all migration
has been successful.
```

---

In Example 8-22, you can see from the message in the IBM Spectrum Archive EE log file that one file we tried to migrate was already in a migrated state and therefore skipped to migrate again, as shown in the following example:

```
GLESM148E(00883): File /ibm/glues/document20.txt is already migrated and will be skipped.
```

For more information about problem determination, see Chapter 10, “Troubleshooting” on page 273.

## 8.25 Collecting IBM Spectrum Archive EE logs for support

If you discover a problem with your IBM Spectrum Archive EE program and open a ticket at the IBM Support Center, you might be asked to provide a package of IBM Spectrum Archive EE log files.

A Linux script is available with IBM Spectrum Archive EE that collects all of the needed files and logs for your convenience to provide them to IBM Support. This task also compresses the files into a single package.

To generate the compressed .tar file and provide it on request to IBM Support, run the following command:

```
ltfsee_log_collection
```

Example 8-23 shows the output of the `ltfsee_log_collection` command. During the log collection run, you are asked what information you want to collect. If you are unsure, select Y to select all the information. At the end of the output, you can find the file name and where the log package was stored.

*Example 8-23 The ltfsee\_log\_collection command*

---

```
[root@ltfs97 ~]# ltfsee_log_collection
```

LTFS Enterprise Edition - log collection program

This program collects the following information from your GPFS cluster.

- (1) Log files that are generated by GPFS, LTFS Enterprise Edition
- (2) Configuration information that are configured to use GPFS and LTFS

Enterprise Edition

- (3) System information including OS distribution and kernel, and hardware information (CPU and memory)

If you want to collect all the above information, input 'y'.  
If you want to collect only (1) and (2), input 'p' (partial).  
If you don't want to collect any information, input 'n'.

The collected data will be zipped in the `ltfsee_log_files_<date>_<time>.tar.gz`.  
You can check the contents of it before submitting to IBM.

```
Input > y
Create a temporary directory '/root/ltfsee_log_files'
copy all log files to the temporary directory
/var/adm/ras/mmfs.log.latest
/var/adm/ras/mmfs.log.previous
/opt/tivoli/tsm/client/hsm/bin/dsmerror.log
/var/log/ltfs.log
/var/log/ltfsee.log
/etc/ltfs.conf
/etc/ltfs.conf.local
/etc/rc.gpfshsm
/opt/tivoli/tsm/client/ba/bin/dsm.opt
/opt/tivoli/tsm/client/ba/bin/dsm.sys
/var/log/messages
get gpfs/hsm configuration
get running process information
get system configuration
list all nodes in the cluster
zip temporary directory
ltfsee_log_files/
ltfsee_log_files/reconcile/
ltfsee_log_files/reconcile/gpfs/
ltfsee_log_files/reconcile/gpfs/gpfs_lists_per_tape/
ltfsee_log_files/reconcile/gpfs/gpfs_lists_per_tape/058AGWL5
ltfsee_log_files/reconcile/gpfs/gpfs_lists_per_tape/YAM049L5
ltfsee_log_files/reconcile/gpfs/gpfs_policies/
ltfsee_log_files/reconcile/gpfs/gpfs_policies/17750464391302654144
ltfsee_log_files/reconcile/gpfs/ltfs_lists/
ltfsee_log_files/reconcile/gpfs/ltfs_lists/YAM049L5
ltfsee_log_files/reconcile/gpfs/gpfs_fss/
ltfsee_log_files/reconcile/gpfs/gpfs_fss/gpfs_fss_invlvd
ltfsee_log_files/reconcile/gpfs/gpfs_fss/gpfs_fss_all
ltfsee_log_files/reconcile/gpfs/gpfs_fss/gpfs_fss_all_dmapi_mntd
ltfsee_log_files/reconcile/gpfs/gpfs_fss/gpfs_fss_all_dmapi
ltfsee_log_files/reconcile/gpfs/gpfs_lists/
ltfsee_log_files/reconcile/gpfs/gpfs_lists/17750464391302654144
ltfsee_log_files/htohru9.ltd.sd1/
ltfsee_log_files/htohru9.ltd.sd1/process.log
ltfsee_log_files/htohru9.ltd.sd1/mmfs.log.latest
ltfsee_log_files/htohru9.ltd.sd1/ltfs.conf
ltfsee_log_files/htohru9.ltd.sd1/ltfsee.log
ltfsee_log_files/htohru9.ltd.sd1/messages
ltfsee_log_files/htohru9.ltd.sd1/dsm.sys
ltfsee_log_files/htohru9.ltd.sd1/dsmerror.log
```

```
ltfsee_log_files/htohru9.ltd.sdl/ltfsee.log
ltfsee_log_files/htohru9.ltd.sdl/ltfsee_info.log
ltfsee_log_files/htohru9.ltd.sdl/information
ltfsee_log_files/htohru9.ltd.sdl/ltfsee.conf.local
ltfsee_log_files/htohru9.ltd.sdl/rc.gpfshsm
ltfsee_log_files/htohru9.ltd.sdl/mmfs.log.previous
ltfsee_log_files/htohru9.ltd.sdl/dsm.opt
ltfsee_log_files/mm1sfs.log
ltfsee_log_files/mm1sconfig.log
ltfsee_log_files/dsmmigfs_query.log
remove temporary directory
The log files collection process is completed
zipped file name: ltfsee_log_files_20130410_145325.tar.gz
```

---

## 8.26 Backing up file systems that are not managed by IBM Spectrum Archive EE

The Tivoli Storage Manager Backup/Archive client and the Tivoli Storage Manager hierarchical storage management (HSM) client from the IBM Spectrum Protect family are components of IBM Spectrum Archive EE and are installed as part of the IBM Spectrum Archive EE installation process.

Therefore, it is possible to use them to back up local server file systems that are not part of the GPFS or IBM Spectrum Scale cluster for IBM Spectrum Archive EE. However, the backup of a GPFS or IBM Spectrum Scale file system that is used by IBM Spectrum Archive EE is *not* supported.

### 8.26.1 Considerations

Consider the following points when you are using the Tivoli Storage Manager Backup/Archive client in the IBM Spectrum Archive EE environment:

- ▶ Licensing

You must purchase a separate Tivoli Storage Manager client license from the IBM Spectrum Protect family to back up your server by using the Tivoli Storage Manager Backup/Archive client that is supplied with IBM Spectrum Archive EE.

- ▶ Compatibility

The Tivoli Storage Manager Backup/Archive client Version 7.1.1.3 from the IBM Spectrum Protect family is installed with IBM Spectrum Archive EE. This version is required by IBM Spectrum Archive EE and was modified to run with IBM Spectrum Archive EE. Therefore, it cannot be upgraded or downgraded independently of the IBM Spectrum Archive EE installation.

This version might cause compatibility issues if, for example, you are running a Tivoli Storage Manager V5.5 server. For more information about the current versions of Tivoli Storage Manager server that are compatible with this Tivoli Storage Manager client, see this website:

<http://www.ibm.com/support/docview.wss?uid=swg21053218>

## 8.26.2 Backing up a GPFS or IBM Spectrum Scale environment

This section describes the standard method of backing up a GPFS or IBM Spectrum Scale environment (GPFS file system) to a Tivoli Storage Manager server of the IBM Spectrum Protect family with and without HSM. For more information, see *General Parallel File System Version 4 Release 2.1 Advanced Administration Guide*, SC23-7032, or see the *IBM Spectrum Scale: Advanced Administration Guide*, which is available at this website:

<https://www.ibm.com/support/knowledgecenter/STXKQY>

The backup of a GPFS file system that is managed by the IBM Spectrum Archive EE environment is *not* supported in the Version 1 release of IBM Spectrum Archive EE. The primary reason is that attempting to back up the stub of a file that was migrated to LTFS causes it to be automatically recalled from LTFS (tape) to the IBM Spectrum Scale. This is not an efficient way to perform backups, especially when you are dealing with large numbers of files.

However, a GPFS file system that is independent of IBM Spectrum Archive EE and not managed by IBM Spectrum Archive EE can be backed up. Normally, the **mmbackup** command is used to back up the files of a GPFS file system to Tivoli Storage Manager servers by using the Tivoli Storage Manager Backup/Archive Client of the IBM Spectrum Protect family. In addition, the **mmbackup** command can operate with regular Tivoli Storage Manager backup commands for backup. After a file system is backed up, you can restore files by using the interfaces that are provided by Tivoli Storage Manager of the IBM Spectrum Protect family.

If HSM is also installed, Tivoli Storage Manager and HSM coordinate to back up and migrate data. If you back up and migrate files to the same Tivoli Storage Manager server, the HSM client can verify that current backup versions of your files exist before you migrate them. If they are not backed up, they must be backed up before migration (if the **MIGREQUIRESBKUP=Yes** option is set for management classes).

## 8.26.3 Backing up non-IBM Spectrum Scale file systems

The backup of non-GPFS file systems space can be accomplished with minor modifications to the embedded Tivoli Storage Manager configuration files that are set up during a standard IBM Spectrum Archive EE installation.

Example 8-24 shows the default `dsm.sys` that is created after the installation of IBM Spectrum Archive EE.

*Example 8-24 Default dsm.sys file*

---

```
SErvername server_a
  COMMethod      TCPIP
  TCPPort        1500
  TCPServeraddress node.domain.company.COM

HSMBACKENDMODE TSMFREE
ERRORLOGNAME /opt/tivoli/tsm/client/hsm/bin/dsmerror.log
```

---

To permit the backup of non GPFS file system to a Tivoli Storage Manager server of the IBM Spectrum Protect family, another server stanza must be added to the file, as shown in Example 8-25.

*Example 8-25 Modified dsm.sys file*

---

```
MIGRATESERVER server_a

SErvername server_a
  COMMMethod      TCPip
  TCPPort         1500
  TCPServeraddress node.domain.company.COM

HSMBACKENDMODE TSMFREE
ERRORLOGNAME /opt/tivoli/tsm/client/hsm/bin/dsmerror.log

SErvername TSM
  COMMMethod      TCPip
  TCPPort         1500
  TCPServeraddress 192.168.10.20
  ERRORLOGNAME /opt/tivoli/tsm/client/ba/bin/dsmerror.log
  PASSWORDACCESS GENERATE
  EXCLUDE.DIR /ibm/glues
```

---

These are the minimum configuration options that are needed to perform a backup. However, if other parameters are required in your environment, you can add them to the SErvername TSM28 stanza.

The **migrate server** option specifies the name of the server to which you want to migrate files from your client node. If you do not specify a server with the **migrate server** option, your files migrate to the server that you specify with the **defaultserver** option. If you do not specify a server with either of these options, your files migrate to the server that you identify in the first stanza of your dsm.sys file.

You also must exclude the GPFS to avoid it being backed up as part of an incremental backup. In our case, the file system is named /ibm/glues/.

Example 8-26 shows the default dsm.opt file that was created after the installation of IBM Spectrum Archive EE.

*Example 8-26 Default dsm.opt file*

---

```
* SErvername      A server name defined in the dsm.sys file

HSMDISABLEAUTOMIGDAEMONS YES
```

---

Although it is possible to add another server stanza to the default dsm.opt file, it is safer to create a separate dsm.opt file specifically for Tivoli Storage Manager backups. In Example 8-27, we created a file that is called dsm\_tsm.opt with the same SErvername as specified in the dsm.sys file. These are the minimum configuration options that are needed to perform a backup. However, if there are other parameters that are required in your environment, you can add them to the “SErvername TSM” stanza.

*Example 8-27 Other dsm\_tsm.opt file*

---

```
SErvername TSM
```

---

To perform a backup of a non-GPFS file system to a Tivoli Storage Manager server of the IBM Spectrum Protect family, the Tivoli Storage Manager Backup/Archive client must be started by using the **-optfile** option, as shown in Example 8-28.

*Example 8-28 Start the Tivoli Storage Manager Backup/Archive Client*

---

```
[root@ltfs97 bin]# dsmc -optfile=dsm_tsm.opt
IBM Tivoli Storage Manager
  Command Line Space Management Client Interface
    Client Version 7, Release 1, Level 0.3
    Client date/time: 09/23/2014 13:45:36
    (c) Copyright by IBM Corporation and other(s) 1990, 2014. All Rights Reserved.

Node Name: LTFS97
Session established with server TSM: Windows
  Server Version 7, Release 1, Level 0
  Server date/time: 09/23/2014 13:45:42  Last access: 09/23/2014 13:43:32
tsm>
tsm> q incl excl
*** FILE INCLUDE/EXCLUDE ***
Mode Function  Pattern (match from top down)  Source File
-----
No exclude filespace statements defined.
Excl Directory /.../.TsmCacheDir              TSM
Excl Directory /.../.SpaceMan                 Operating System
Excl Directory /ibm/glues
/opt/tivoli/tsm/client/ba/bin/dsm.sys
Exclude HSM      /etc/adsm/SpaceMan/config/.../* Operating System
Exclude HSM      /.../.SpaceMan/.../*      Operating System
Exclude Restore  /.../.SpaceMan/.../*      Operating System
Exclude Archive  /.../.SpaceMan/.../*      Operating System
No DFS include/exclude statements defined.
tsm>
```

---

For more information about how to back up a server by using Tivoli Storage Manager of the IBM Spectrum Protect family, see the client manual found at this website:

[http://www.ibm.com/support/knowledgecenter/api/redirect/tsminfo/v6r4/topic/com.ibm.itism.client.doc/b\\_ba\\_guide\\_unx\\_lnx.pdf](http://www.ibm.com/support/knowledgecenter/api/redirect/tsminfo/v6r4/topic/com.ibm.itism.client.doc/b_ba_guide_unx_lnx.pdf)

## 8.26.4 IBM TS4500 Automated Media Verification with IBM Spectrum Archive EE

In some use cases where IBM Spectrum Archive EE is deployed, you might have the requirement to periodically ensure that the files and data that is migrated from the IBM Spectrum Scale file system to physical tape is still readable and can be recalled back from tape to the file system without any error. Especially in a more long-term archival environment, a function that checks the physical media based on a schedule that the user can implement is highly appreciated.

Starting with the release of the IBM TS4500 Tape Library R2, a new, fully transparent function is introduced within the TS4500 operations that is named *policy-based automatic media verification*. This new function is hidden from any ISV software, similar to the automatic cleaning.

No ISV certification is required. It can be enabled/disabled through the logical library with additional settings to define the verify period (for example, every 6 months) and the first verification date.

One or more designated media verification drives (MVDs) must be assigned to a logical library in order for the verification to take place. A preferred practice is to have two MVDs assigned at a time to ensure that no false positives occur because of a faulty tape drive. Figure 8-2 shows an example of such a setup.

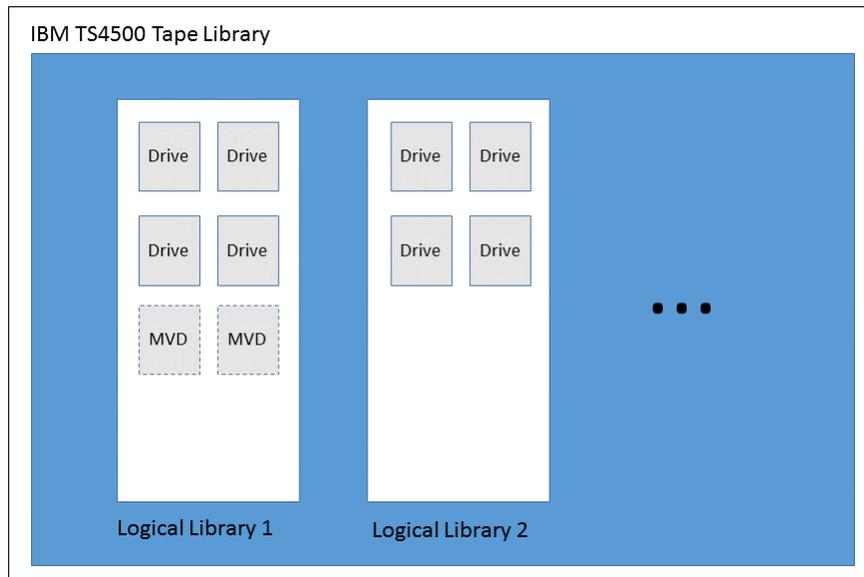


Figure 8-2 TS4500 with one logical library showing two MVDs configured

**Note:** MVDs defined within a logical library are not accessible from the host or application by using the drives of this particular logical library for production.

Verify that results are simple pass/fail, but verify that failures are retried on a second physical drive, if available, before being reported. A failure is reported through all normal notification options (email, syslog, and SNMP). MVDs are not reported as mount points (SCSI DTEs) to the ISV application, so MVDs *do not* need to be connected to the SAN.

During this process, whenever access from the application or host is required to the physical media under media verification, the tape library stops the current verification process. It then dismounts the needed tape from the MVD, and mounts it to a regular tape drive within the same logical library for access by the host application to satisfy the requested mount. At a later point, the media verification process continues.

The library GUI Cartridges page adds columns for last verification date/time, verification result, and next verification date/time (if automatic media verification is enabled). If a cartridge being verified is requested for ejection or mounting by the ISV software (which thinks the cartridge is in a storage slot), the verify task is automatically canceled, a checkpoint occurs, and the task resumes later (if/when the cartridge is available).

The ISV eject or mount occurs with a delay comparable to a mount to a drive being cleaned (well within the preferred practice SCSI Move Medium timeout values). The GUI also supports a manual stop of the verify task.

The last verification date/time is written in the cartridge memory (CM) and read upon first mount after being newly inserted into a TS4500, providing persistence and portability (similar to a cleaning cartridge usage count).

All verify mounts are recorded in the mount history CSV file, allowing for more granular health analysis (for example, outlier recovered error counts) by using Tape System Reporter (TSR) or Rocket Server graph.

The whole media verification process is transparent to IBM Spectrum Archive EE as the host. No definitions and configurations need to be done within IBM Spectrum Archive EE. All setup activities are done only through the TS4500 management interface.

Figure 8-3 to Figure 8-6 on page 262 show screen captures from the TS4500 tape library web interface that show you how to assign an MVD to a logical library. It is a two-step process because you must define a drive to be an MVD and then assign this drive before the logical library (if it was not assigned before).

1. You can select the menu option **Drives by Logical Library** to assign an unassigned drive to a logical library by right-clicking the unassigned drive icon. A menu opens where you select **Assign**, as shown in Figure 8-3.

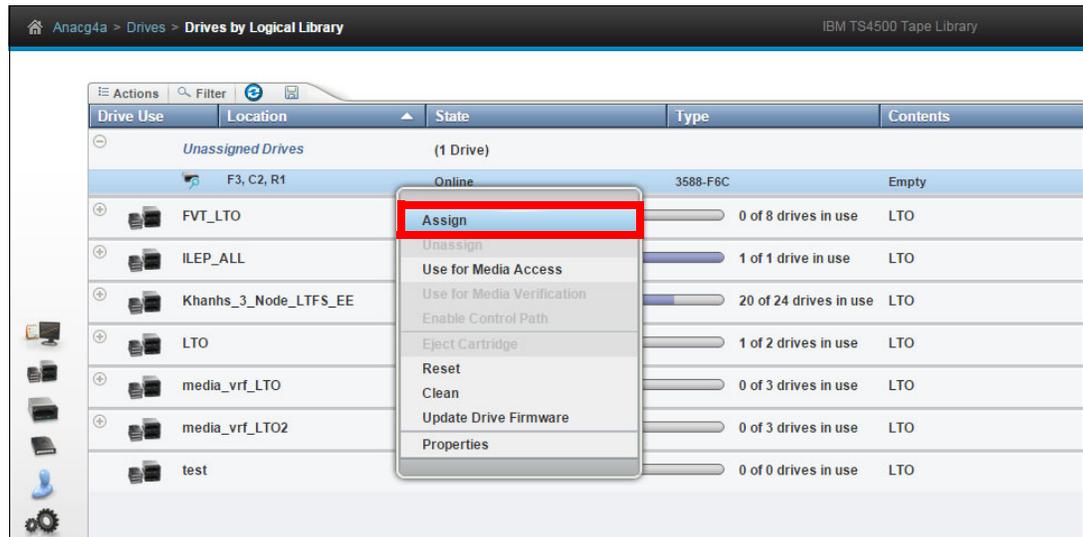


Figure 8-3 Assign a tape drive to a logical library through the TS4500 web interface (step 1)

- Another window opens where you must select the specific logical library to which the unassigned drive is supposed to be added, as shown in Figure 8-4.

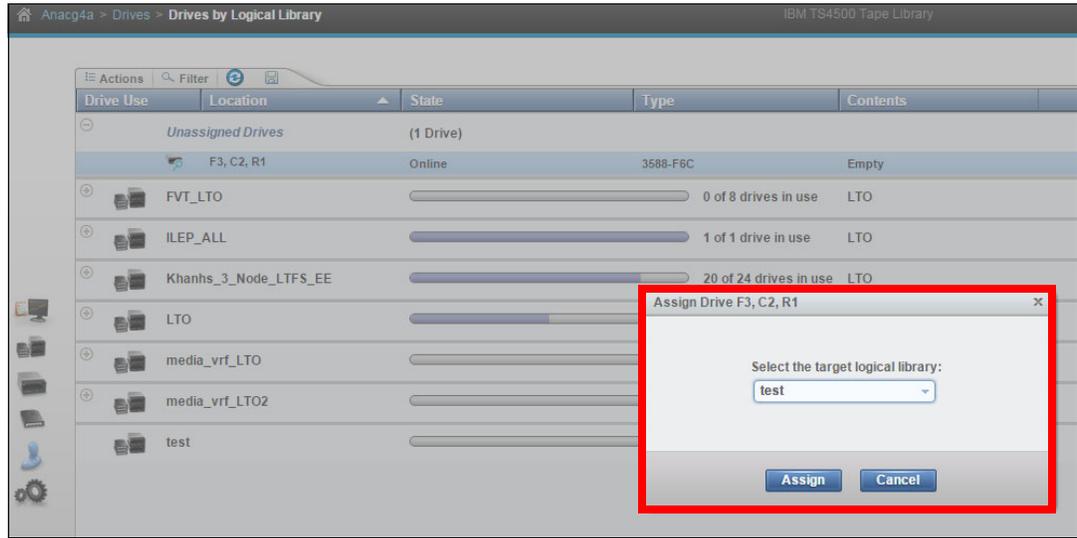


Figure 8-4 Assign a tape drive to a logical library through the TS4500 web interface (step 2)

- If the drive to be used as an MVD is configured within the logical library, change its role, as shown in Figure 8-5 and Figure 8-6 on page 262.

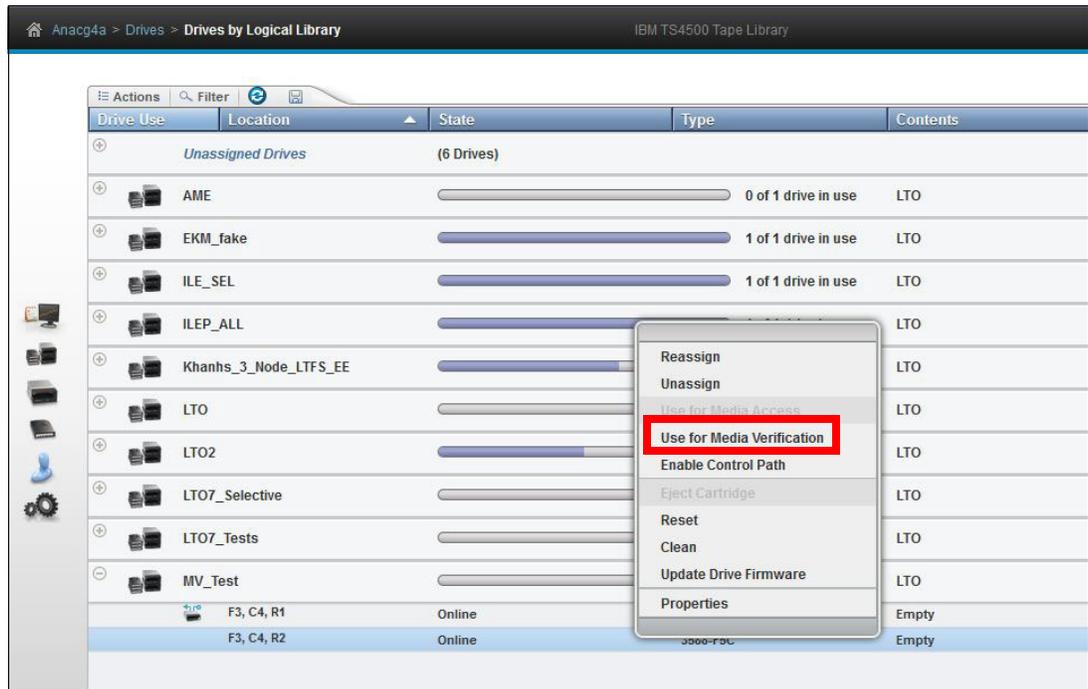


Figure 8-5 Reserve a tape drive as the media verification drive through the TS4500 web interface

- You must right-click the assigned drive within the logical library. A menu opens and you select **Use for Media Verification** from the list of the provided options. A confirmation dialog box opens. Click **Yes** to proceed.

- After making that configuration change to the drive, you see a new icon in front of it to show you the new role (Figure 8-6).

⊕	 LT07_Selective	<div style="width: 100%; height: 10px; background: linear-gradient(to right, #ccc, #ccc);"></div>	0 of 1 drive in use	LTO
⊕	 LT07_Tests	<div style="width: 100%; height: 10px; background: linear-gradient(to right, #ccc, #ccc);"></div>	0 of 1 drive in use	LTO
⊖	 MV_Test	<div style="width: 100%; height: 10px; background: linear-gradient(to right, #ccc, #ccc);"></div>	0 of 2 drives in use	LTO
	 F3, C4, R1	Online	3588-F5C	Empty
	 F3, C4, R2	Online	3588-F5C	Empty

Figure 8-6 Display a tape drive as the media verification drive through the TS4500 web interface

**Note:** The MVD flag for a tape drive is a global setting, which means that after it is assigned, the drive keeps its role as an MVD even it is unassigned and then assigned to a new logical library. Unassigning does not disable this role.

To unassign a drive from being an MVD, follow the same procedure again, and select (after the right-click) **Use for Media Access**. This action changes the drive role back to normal operation for the attached host application to this logical library.

Figure 8-7 shows you the TS4500 web interface dialog box for enabling automatic media verification on an existing logical library. You must go to the Cartridges by Logical Library page. Then, select **Modify Media Verification** for the selected logical library. The Automatic Media Verification dialog box opens where you can enter the media verification schedule.

**Modify Media Verification on Logical Library MV\_Test** ✕

Automatic media verification:  **Enable**

Verify media every:   ▾

First verification date:  ▾

Media verification drives:   ▾

Figure 8-7 Modify Media Verification dialog box to set up a schedule

By using this dialog box, you can enable/disable an automatic media verification schedule. Then, you can configure how often the media should be verified and the first verification date. Finally, you can select the MVDs, which are selected by the library to perform the scheduled media verification test.

If you go to the Cartridges by Logical Library page and select **Properties** for the selected logical library, a dialog box opens where you can see the current media verification configuration for that logical library, as shown by Figure 8-8.

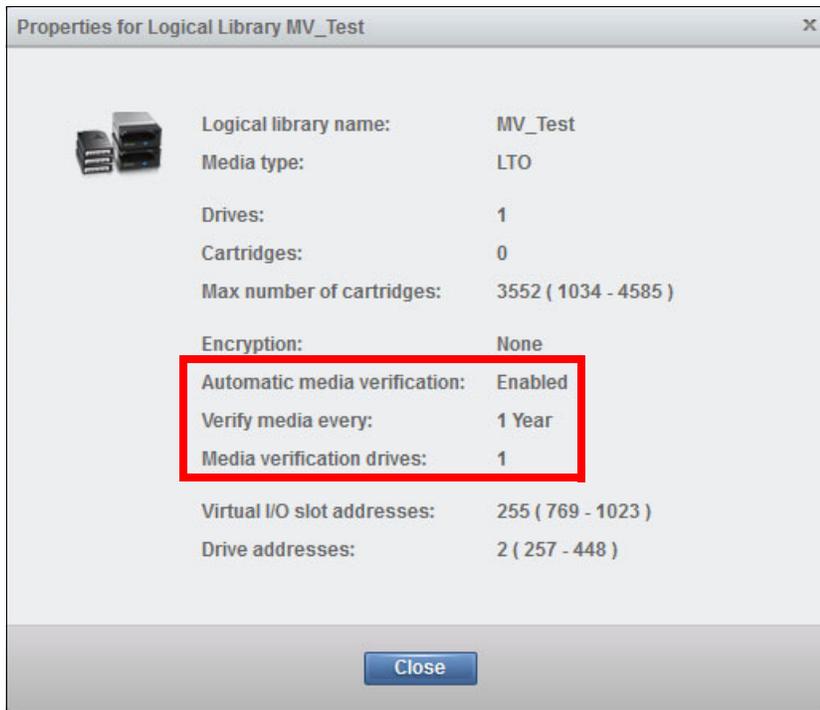


Figure 8-8 TS4500 properties for a logical library

For more information and the usage of the TS4500 R2 media verification functions, see *IBM IBM TS4500 R4 Tape Library Guide*, SG24-8235, and IBM TS4500 documentation at IBM Knowledge Center:

[http://www.ibm.com/support/knowledgecenter/STQRQ9/com.ibm.storage.ts4500.doc/ts4500\\_ichome.html](http://www.ibm.com/support/knowledgecenter/STQRQ9/com.ibm.storage.ts4500.doc/ts4500_ichome.html)





# Use cases

This chapter describes various use case examples for IBM Spectrum Archive EE.

This chapter includes the following topics:

- ▶ Use cases overview
- ▶ AFM use cases

## 9.1 Use cases overview

The typical use cases for IBM Spectrum Archive EE can be broken into three categories: Archive, tiered storage, and data exchange, as shown in Figure 9-1.

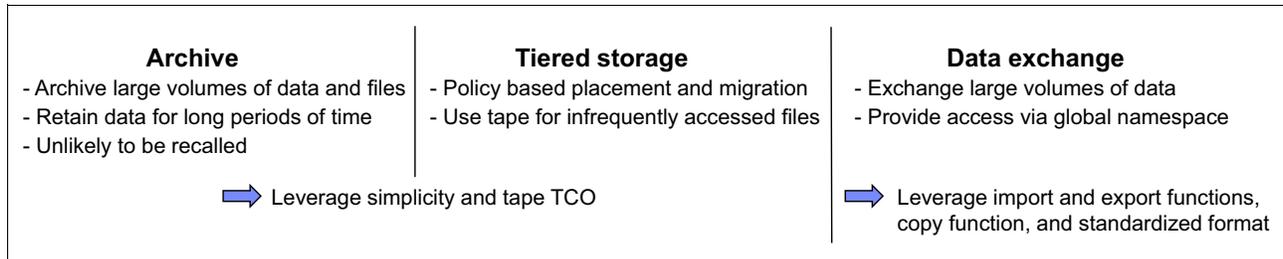


Figure 9-1 Typical use cases for IBM Spectrum Archive EE

For more information about each use case, see Figure 9-2, Figure 9-3 on page 267, and Figure 9-4 on page 268.

### 9.1.1 Use case for archive

Figure 9-2 summarizes the requirements, solution, and benefits of an IBM Spectrum Archive EE use case for archiving data.

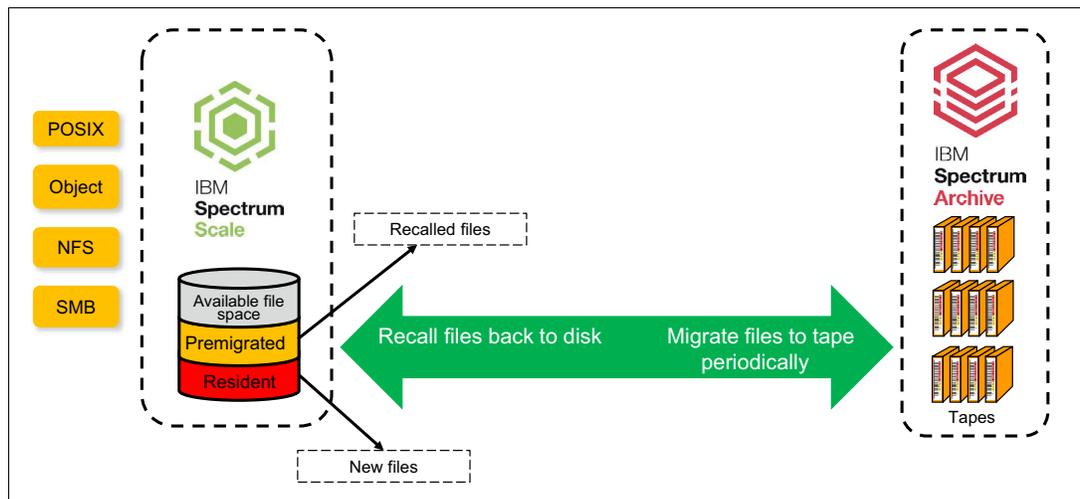


Figure 9-2 Use case for archive

Some of the requirements for the archive use case are:

- ▶ Large amount of data, larger files
- ▶ Infrequently accessed
- ▶ Longer retention periods
- ▶ Easy data access

The solution is based on the following:

- ▶ Archive storage based on IBM Spectrum Scale, IBM Spectrum Archive EE, and standard file system interfaces

Some of the archive use case benefits are:

- ▶ Simplicity with file system interface
- ▶ Scalable with IBM Spectrum Scale and IBM Spectrum Archive
- ▶ EE Low TCO with IBM tape

## 9.1.2 Use case for tiered and scalable storage

Figure 9-3 summarizes the requirements, solution, and benefits of an IBM Spectrum Archive EE use case for tiered and scalable storage.

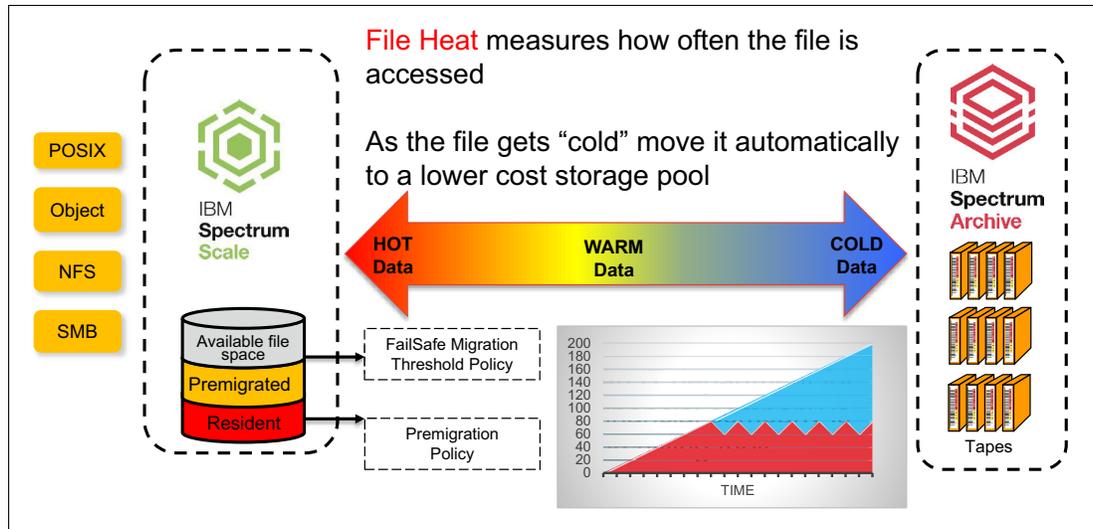


Figure 9-3 Use case for tiered and scalable storage

Some of the requirements for the tiered and scalable use case are:

- ▶ Archive to file systems
- ▶ Simple backup solution
- ▶ Easy data access for restore

The solution is based on the following:

- ▶ Archive storage based on IBM Spectrum Scale, IBM Spectrum Archive EE, and standard file system interfaces

Some of the tiered and scalable use case benefits are:

- ▶ Easy to use with standard copy tools
- ▶ Scalable with IBM Spectrum Scale and IBM Spectrum Archive EE
- ▶ Low TCO with IBM tape

### 9.1.3 Use case data exchange

Figure 9-4 summarizes the requirements, solution, and benefits of an IBM Spectrum Archive EE use case for data exchange.

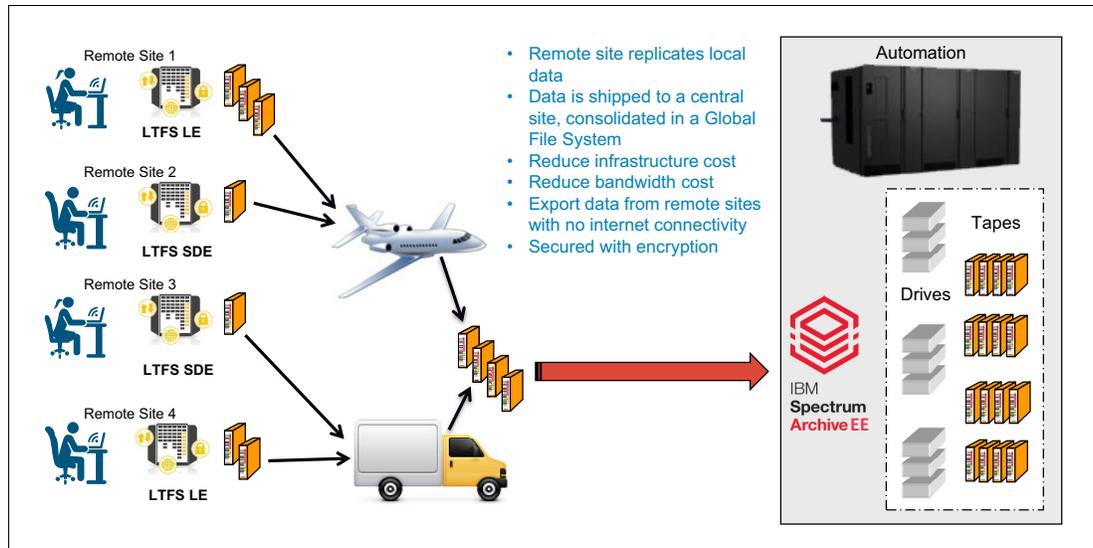


Figure 9-4 Use case for data exchange

Some of the requirements for the data exchange use case are:

- ▶ Export entire directories to tape
- ▶ Import files and directories with seamless data access
- ▶ Leverage global name space

The solution is based on the following:

- ▶ IBM Spectrum Scale, IBM Spectrum Archive EE, and standard file system interfaces using export and import functions

Some of the data exchange use case benefits are:

- ▶ Export of tape copies
- ▶ Efficient import without reading data
- ▶ Import and export within global namespace

## 9.2 AFM use cases

This section covers the use of a home and cache sites, and delves into two typical use cases of IBM Spectrum Archive EE with IBM Spectrum Scale AFM explaining the Centralized Archive Repository scenario and the Asynchronous Archive Replication scenario.

Active file management (AFM) uses a home-and-cache model in which a single home provides the primary storage of data, and exported data is cached in a local GPFS file system:

- Home** A home site is an NFS export of a remote cluster. This export can be a local file system in the remote cluster, a GPFS file system, or a GPFS fileset in the remote cluster. AFM is supported when a remote file system is mounted on the cache cluster using GPFS protocols. This configuration requires that a multicluster setup exists between the home and cache before AFM can use the home cluster's file system mount for AFM operations.
- Cache** A cache site is a remote cluster with a GPFS fileset that has a mount point to the exported NFS file system of the home cluster's file system. A cache site uses a proprietary protocol over NFS. Each AFM-enabled fileset has a single home cluster associated with it (represented by the host name of the home server).

### 9.2.1 Centralized Archive Repository

In an environment where data needs to be allocated together to create a bigger picture, archiving, or for disaster recovery planning, users can make a Centralized Archive Repository. A Centralized Archive Repository uses a single home cluster that can have multiple NFS exports to many cache sites.

In this setup, IBM Spectrum Archive EE is configured on the home cluster to archive all the data generated from each cache cluster. The idea behind this solution is to have a single home repository that has a large disk space, and multiple cache sites that cannot afford large disk space.

**Note:** AFM supports multiple cache modes, and this solution can be used with single writer or independent writer. However, with the release of IBM Spectrum Archive EE V1.2.3.0, only the independent writer is currently supported.

When files are generated on the cache clusters, they are asynchronously replicated to the home site. When these files are no longer being accessed on the cache clusters, the files can be evicted, freeing up disk space at the cache clusters. They can then be migrated onto tape at the home cluster. If evicted files need to be accessed again at the cache clusters, they can simply be recovered by opening the file for access or by using AFM's prefetch operation to retrieve multiple files back to disk from the home site.

Figure 9-5 shows a configuration of a single home cluster with multiple cache clusters to form a Centralized Archive Repository.

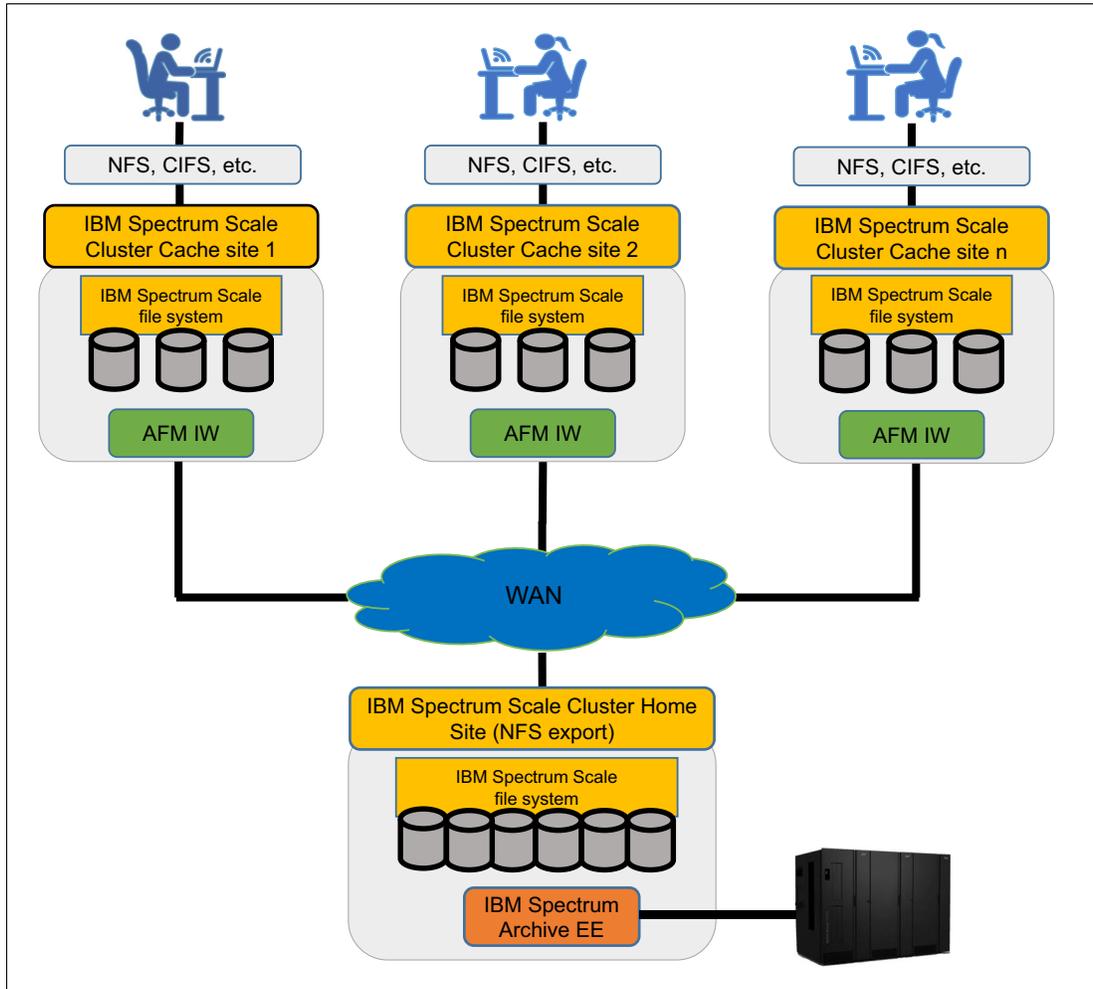


Figure 9-5 Centralized Archive Repository

Some examples of customers who can benefit from this solution are research groups that are spread out geographically and rely on each group's data, such as universities. Medical groups and media companies can also benefit.

## 9.2.2 Asynchronous Archive Replication

Asynchronous Archive Replication is an extension to the stretched cluster configuration. In it, users require the data created is replicated to a secondary site and can be migrated to tape at both sites. By incorporating IBM Spectrum Scale AFM to the stretched cluster idea, there are no limits on how far away the secondary site is located. In addition to geolocation capabilities, data created on home or cache is asynchronously replicated to the other site.

Asynchronous Archive Replication requires two remote clusters configured, one being the home cluster and the other being a cache cluster with the independent writer mode. By using the independent writer mode in this configuration, users can create files at either site and the data/metadata is asynchronously replicated to the other site.

**Note:** With independent writer, the cache site always wins during file modifications. If files are created at home, only metadata is transferred to the cache at the next update or refresh. To obtain the file's data from the home site at the cache site, use AFM's prefetch operation to get the data or open specific files. The data is then propagated to the cache nodes.

Figure 9-6 shows a configuration of an Asynchronous Archive Replication solution between a home and cache site.

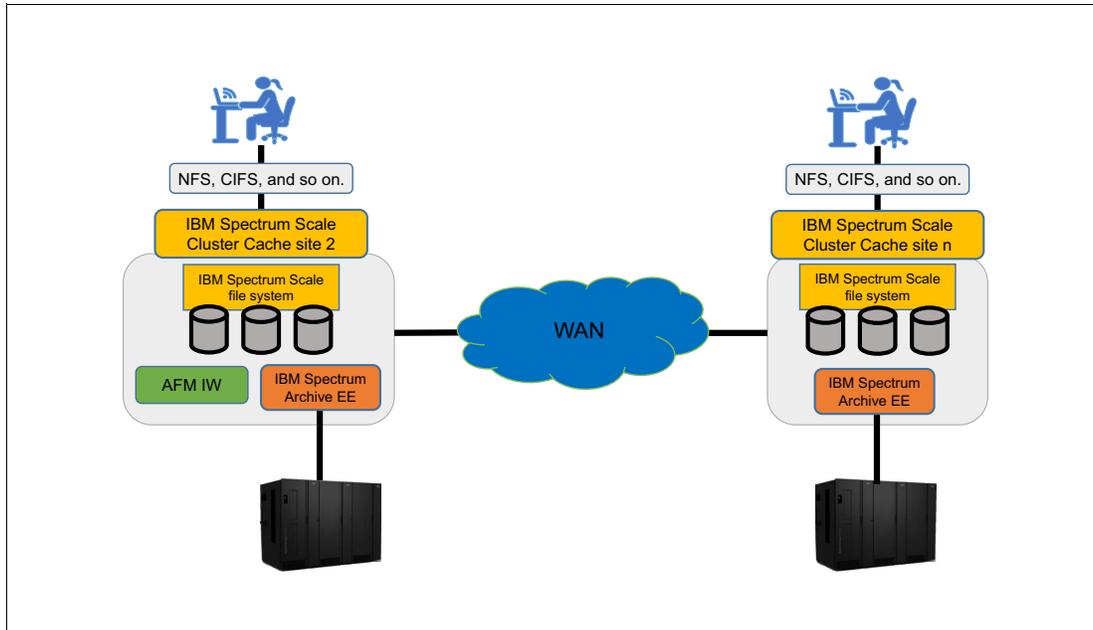


Figure 9-6 Asynchronous Archive Replication





# Troubleshooting

This chapter describes the process that you can use to troubleshoot issues with IBM Spectrum Archive EE.

This chapter includes the following topics:

- ▶ Overview
- ▶ Hardware
- ▶ Recovering data from a write failure tape
- ▶ Recovering data from a read failure tape
- ▶ Handling export errors
- ▶ Software
- ▶ Recovering from system failures

## 10.1 Overview

This section provides a simple health check procedure for IBM Spectrum Archive EE.

### 10.1.1 Quick health check

If you are having issues with an existing IBM Spectrum Archive EE environment, Figure 10-1 shows a simple flowchart that you can follow as the first step to troubleshooting problems with the IBM Spectrum Archive EE components.

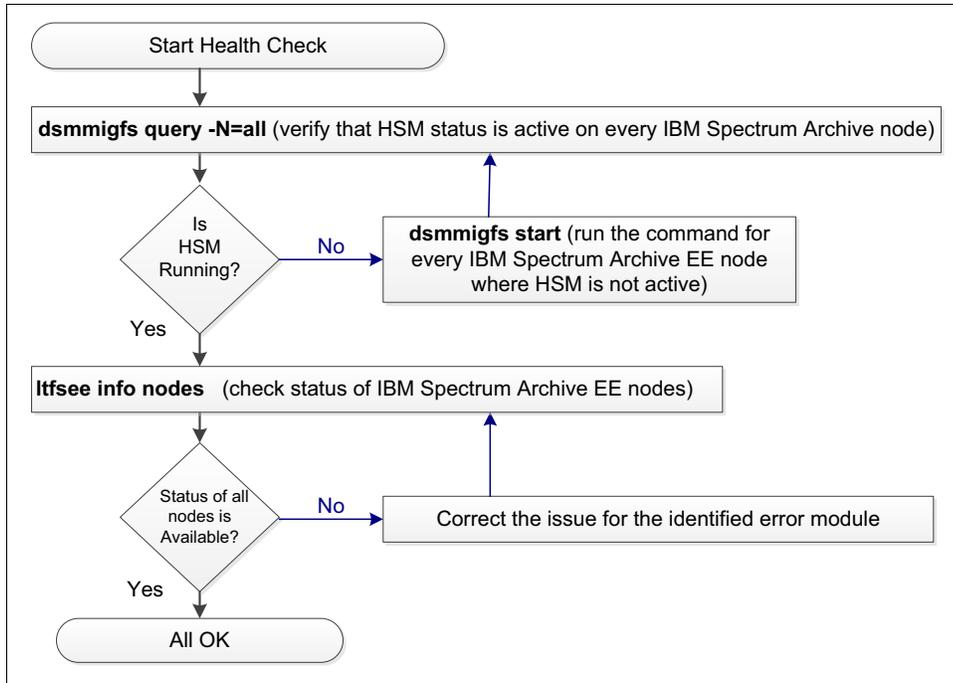


Figure 10-1 Quick health check procedure

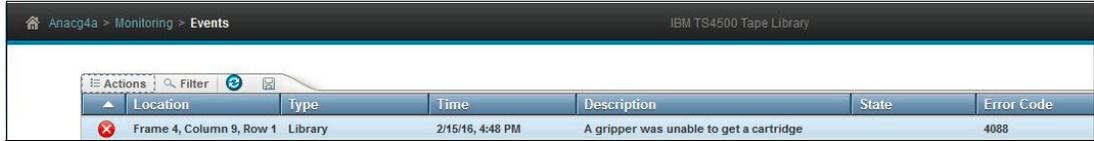
If your issue remains after you perform these simple checks, follow the procedures that are described in the remainder of this chapter to perform more detailed troubleshooting. If the problem cannot be resolved, contact IBM Spectrum Archive Support.

## 10.2 Hardware

The topics in this section provide information that can help you to identify and resolve problems with the hardware that is used by IBM Spectrum Archive EE.

## 10.2.1 Tape library

If the TS4500 tape library has a problem, it reports an error in the events page on the TS4500 Management GUI. When an error occurs, IBM Spectrum Archive might not work. Figure 10-2 shows an example of a library error.



The screenshot shows the 'Events' page for the 'IBM TS4500 Tape Library'. A table displays one error entry with a red 'X' icon in the 'Location' column.

Location	Type	Time	Description	State	Error Code
Frame 4, Column 9, Row 1	Library	2/15/16, 4:48 PM	A gripper was unable to get a cartridge		4088

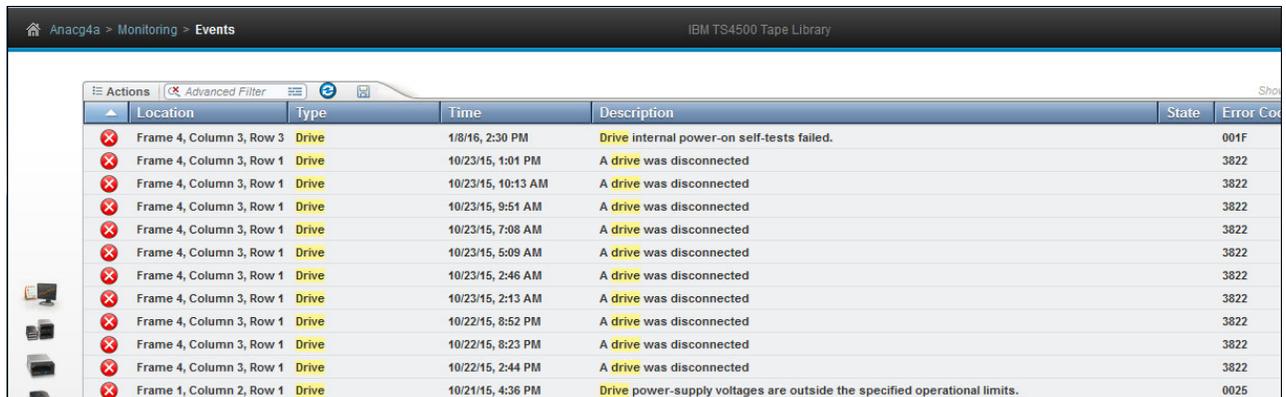
Figure 10-2 Tape library error log

For more information about how to solve tape library errors, see the *IBM TS4500 R4 Tape Library Guide*, SG24-8235.

## 10.2.2 Tape drives

If an LTO tape drive has a problem, it reports the error on a single-character display (SCD). If a TS1140 (or later) tape drive has a problem, it reports the error on an 8-character message display. When this error occurs, IBM Spectrum Archive might not work. To obtain information about a drive error, determine which drive is reporting the error and then access the events page to see the error by using the TS4500 Management GUI.

Figure 10-3 shows an example from the web interface of a tape drive that has an error and is no longer responding.



The screenshot shows the 'Events' page for the 'IBM TS4500 Tape Library' with a list of 14 error entries. Each entry has a red 'X' icon in the 'Location' column and a yellow 'Drive' label in the 'Type' column.

Location	Type	Time	Description	State	Error Code
Frame 4, Column 3, Row 3	Drive	1/8/16, 2:30 PM	Drive internal power-on self-tests failed.		001F
Frame 4, Column 3, Row 1	Drive	10/23/15, 1:01 PM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/23/15, 10:13 AM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/23/15, 9:51 AM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/23/15, 7:08 AM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/23/15, 5:09 AM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/23/15, 2:46 AM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/23/15, 2:13 AM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/22/15, 8:52 PM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/22/15, 8:23 PM	A drive was disconnected		3822
Frame 4, Column 3, Row 1	Drive	10/22/15, 2:44 PM	A drive was disconnected		3822
Frame 1, Column 2, Row 1	Drive	10/21/15, 4:36 PM	Drive power-supply voltages are outside the specified operational limits.		0025

Figure 10-3 Tape drive error

If you right-click the event and select **Display fix procedure**, another window opens and shows suggestions about how to fix the problem. If a drive display reports a specific drive error code, see the tape drive maintenance manual for a solution. For more information about analyzing the operating system error logs, see 10.6.1, “Linux” on page 289.

If a problem is identified in the tape drive and the tape drive must be repaired, the drive must first be removed from the IBM Spectrum Archive EE system. For more information, see “Taking a tape drive offline” on page 277.

## Managing tape drive dump files

This section describes how to manage the automatic erasure of drive dump files. IBM Spectrum Archive automatically generates two tape drive dump files in the /tmp directory when it receives unexpected sense data from a tape drive. Example 10-1 shows the format of the dump files.

### Example 10-1 Dump files

---

```
[root@ltfs97 tmp]# ls -la *.dmp
-rw-r--r-- 1 root root 3681832 Apr  4 14:26 ltfs_1068000073_2013_0404_142634.dmp
-rw-r--r-- 1 root root 3681832 Apr  4 14:26 ltfs_1068000073_2013_0404_142634_f.dmp
-rw-r--r-- 1 root root 3681832 Apr  4 14:42 ltfs_1068000073_2013_0404_144212.dmp
-rw-r--r-- 1 root root 3697944 Apr  4 14:42 ltfs_1068000073_2013_0404_144212_f.dmp
-rw-r--r-- 1 root root 3697944 Apr  4 15:45 ltfs_1068000073_2013_0404_154524.dmp
-rw-r--r-- 1 root root 3683424 Apr  4 15:45 ltfs_1068000073_2013_0404_154524_f.dmp
-rw-r--r-- 1 root root 3683424 Apr  4 17:21 ltfs_1068000073_2013_0404_172124.dmp
-rw-r--r-- 1 root root 3721684 Apr  4 17:21 ltfs_1068000073_2013_0404_172124_f.dmp
-rw-r--r-- 1 root root 3721684 Apr  4 17:21 ltfs_1068000073_2013_0404_172140.dmp
-rw-r--r-- 1 root root 3792168 Apr  4 17:21 ltfs_1068000073_2013_0404_172140_f.dmp
```

---

The size of each drive dump file is approximately 2 MB. By managing the drive dump files that are generated, you can save disk space and enhance IBM Spectrum Archive performance.

It is not necessary to keep dump files after they are used for problem analysis. Likewise, the files are not necessary if the problems are minor and can be ignored. A script program that is provided with IBM Spectrum Archive EE periodically checks the number of drive dump files and their date and time. If some of the dump files are older than two weeks or if the number of dump files exceeds 1000 files, the script program erases them.

The script file is started by using Linux **crontab** features. A `cron_ltfs_limit_dumps.sh` file is in the `/etc/cron.daily` directory. This script file is started daily by the Linux operating system. The interval to run the script can be changed by moving the `cron_ltfs_limit_dumps.sh` file to other cron folders, such as `cron.weekly`. For more information about how to change the crontab setting, see the manual for your version of Linux.

In the `cron_ltfs_limit_dumps.sh` file, the automatic drive dump erase policy is specified by the option of the `ltfs_limit_dump.sh` script file, as shown in the following example:

```
/opt/IBM/ltfs/bin/ltfs_limit_dumps.sh -t 14 -n 1000
```

You can modify the policy by editing the options in the `cron_ltfs_limit_dumps.sh` file. The expiration date is set as a number of days by the `-t` option. In the example, a drive dump file is erased when it is more than 14 days old. The number of files to keep is set by the `-n` option. In our example, if the number of files exceeds 1,000, older files are erased so that the 1,000-file maximum is not exceeded. If either of the options are deleted, the dump files are deleted by the remaining policy.

By editing these options in the `cron_ltfs_limit_dumps.sh` file, the number of days that files are kept and the number of files that are stored can be modified.

Although not recommended, you can disable the automatic erasure of drive dump files by removing the `cron_ltfs_limit_dumps.sh` file from the cron folder.

## Taking a tape drive offline

This section describes how to take a drive offline from the IBM Spectrum Archive EE system to perform diagnostic operations while the IBM Spectrum Archive EE system stays operational. To accomplish this task, use software such as the IBM Tape Diagnostic Tool (ITDT) or the IBM LTFS Format Verifier, which are described in 11.3, “System calls and IBM tools” on page 328.

**Important:** If the diagnostic operation you intend to perform requires that a tape cartridge be loaded into the drive, ensure that you have an empty non-pool tape cartridge available in the logical library of IBM Spectrum Archive EE. If a tape cartridge is in the tape drive when the drive is removed, the tape cartridge is automatically moved to the home slot.

To perform diagnostic tests, complete the following steps:

1. Identify the node ID number of the drive to be taken offline by running the **ltfsee info drives** command. Example 10-2 shows the tape drives in use by IBM Spectrum Archive EE.

*Example 10-2 Identify the tape drive to remove*

```
[root@ltfssn1 ~]# ltfsee info drives
Drive S/N      Status      Type      Role  Library      Address  Node ID  Tape
Node Group
000000014A00  Not mounted TS1140    mrg   TS4500_3592  257     1        -
SINGLE_NODE
0000078D8320  Mounted     TS1150    mrg   TS4500_3592  258     1
JCC539JC SINGLE_NODE
```

In this example, we take the tape drive with serial number 000000014A00 on cluster node 1 offline.

2. Remove the tape drive from the IBM Spectrum Archive EE inventory by specifying the **ltfsee drive remove <drive serial number>** command. Example 10-3 shows the removal of a single tape drive from IBM Spectrum Archive EE.

*Example 10-3 Remove the tape drive*

```
[root@ltfssn1 ~]# ltfsee drive remove -d 000000014A00
GLES121I(00282): Drive serial 000000014A00 is removed from LTFS EE drive list.
```

3. Check the success of the removal. Run the **ltfsee info drives** command and verify that the output shows that the MMM attribute for the drive is in the stock state. Example 10-4 shows the status of the drives after it is removed from IBM Spectrum Archive EE.

*Example 10-4 Check the tape drive status*

```
[root@ltfssn1 ~]# ltfsee info drives
Drive S/N      Status      Type      Role  Library      Address  Node ID  Tape      Node
Group
0000078D8320  Mounted     TS1150    mrg   TS4500_3592  258     1        JCC539JC
SINGLE_NODE
000000014A00  Stock       UNKNOWN   ---   TS4500_3592  257     -        -        -
```

4. Identify the primary device number of the drive for subsequent operations by running the **cat /proc/scsi/IBMtape** command. The command output lists the device number in the Number field. Example 10-5 shows the output of this command with the offline tape drive 13 in bold.

*Example 10-5 List the tape drives in Linux*

---

```
[root@1tfs97 /]# cat /proc/scsi/IBMtape
lin_tape version: 1.76.0
lin_tape major number: 248
Attached Tape Devices:
```

Number	model	S/N	HBA	SCSI	FO Path
0	ULT3580-TD5	00078A218E	qla2xxx	2:0:0:0	NA
1	ULT3580-TD5	1168001144	qla2xxx	2:0:1:0	NA
2	ULTRIUM-TD5	9A700L0077	qla2xxx	2:0:2:0	NA
3	ULT3580-TD6	1013000068	qla2xxx	2:0:3:0	NA
4	03592E07	0000013D0485	qla2xxx	2:0:4:0	NA
5	ULT3580-TD5	00078A1D8F	qla2xxx	2:0:5:0	NA
6	ULT3580-TD6	00013B0037	qla2xxx	2:0:6:0	NA
7	03592E07	001013000652	qla2xxx	2:0:7:0	NA
8	03592E07	0000078DDAC3	qla2xxx	2:0:8:0	NA
9	03592E07	001013000255	qla2xxx	2:0:9:0	NA
10	ULT3580-TD5	1068000073	qla2xxx	2:0:10:0	NA
11	03592E07	0000013D0734	qla2xxx	2:0:11:0	NA
12	ULT3580-TD6	00013B0084	qla2xxx	2:0:12:0	NA
<b>13</b>	<b>3592E07</b>	<b>00000014A00</b>	<b>qla2xxx</b>	<b>2:0:13:0</b>	<b>NA</b>
14	ULT3580-TD5	1068000070	qla2xxx	2:0:14:0	NA
15	ULT3580-TD5	1068000016	qla2xxx	2:0:15:0	NA
16	03592E07	0000013D0733	qla2xxx	2:0:19:0	NA
17	03592E07	0000078DDBF1	qla2xxx	2:0:20:0	NA
18	ULT3580-TD5	00078AC0A5	qla2xxx	2:0:21:0	NA
19	ULT3580-TD5	00078AC08B	qla2xxx	2:0:22:0	NA
20	03592E07	0000013D0483	qla2xxx	2:0:23:0	NA
21	03592E07	0000013D0485	qla2xxx	2:0:24:0	NA
22	03592E07	0000078D13C1	qla2xxx	2:0:25:0	NA

---

5. If your diagnostic operations require a tape cartridge to be loaded into the drive, complete the following steps. Otherwise, you are ready to perform diagnostic operations on the drive, which has the drive address `/dev/IBMtapenumber`, where number is the device number that is obtained in step 4:
  - a. Move the tape cartridge to the drive from the I/O station or home slot. You can move the tape cartridge by using ITDT (in which case the drive must have the control path), or the TS4500 Management GUI.
  - b. Perform the diagnostic operations on the drive, which has the drive address `/dev/IBMtapenumber`, where number is the device number that is obtained in step 4.
  - c. When you are finished, return the tape cartridge to its original location.
6. Add the drive to the IBM Spectrum Archive EE inventory again by running the **1tfssee drive add drive\_serial node\_id** command, where **node\_id** is the same node that the drive was assigned to originally in step 1 on page 277.

Example 10-6 shows the tape drive that is readed to IBM Spectrum Archive EE.

*Example 10-6 Add again the tape drive*

```
[root@l1fssn1 ~]# ltfsee drive add -d 000000014A00 -n 1
GLESL119I(00174): Drive 000000014A00 added successfully.
```

Running the **ltfsee info drives** command again shows that the tape drive is no longer in a “stock” state. Example 10-7 shows the output of this command.

*Example 10-7 Check the tape drive status*

```
[root@l1fssn1 ~]# ltfsee info drives
Drive S/N      Status      Type      Role  Library      Address  Node ID  Tape      Node Group
000000014A00  Not mounted TS1140    mrg   TS4500_3592  257     1        -         SINGLE_NODE
0000078D8320  Mounted     TS1150    mrg   TS4500_3592  258     1        JCC539JC SINGLE_NODE
```

## 10.2.3 Tape cartridge

Table 10-1 shows all possible status conditions for an IBM Spectrum Archive tape cartridge as displayed by the **ltfsee info tapes** command.

*Table 10-1 Tape cartridge status*

Tape cartridge status	File system access	Description	How to recover tape cartridges	Valid IBM Spectrum Archive EE commands
Valid	Yes	The Valid status indicates that the cartridge is valid.	N/A	All but imports, recover, and pool add
Exported	No	The Exported status indicates that the cartridge is valid and is exported.	N/A	ltfsee import ltfsee tape move
Offline	No	The Offline status indicates that the cartridge is valid and is exported offline.	N/A	ltfsee import with offline option ltfsee tape move
Unknown	Unknown	The tape cartridge contents are unknown.	The status change after it is used. Run <b>ltfsee tape validate</b> command to check the tape into a drive.	ltfsee pool remove ltfsee reclaim (source/target) ltfsee export ltfsee rebuild ltfsee recall
Write Protected	Read-only	The tape cartridge is physically (or logically) in a write-protected state.	Remove the physical write-protection.	ltfsee pool remove ltfsee recall

Tape cartridge status	File system access	Description	How to recover tape cartridges	Valid IBM Spectrum Archive EE commands
Critical	Read-only	Indicates that the tape had a write failure. The index in memory is dirty.	Run the <b>recover</b> command using the <b>-c</b> option to generate a scan list and bring the files back into resident state, then run the <b>recover</b> command using the <b>-r</b> option to double check any missed files. The <b>-r</b> option removes the tape from the tape cartridge pool if no new files were detected remaining on tape. Save the tape just in case there are issues recovering files and contact IBM Spectrum Archive support to determine the root cause of the critical tape.	ltfsee recover ltfsee recall
Warning	Read-only	Indicates that the tape had a read failure.	Files on Warning tapes can be recovered by using the <code>relocate_replica.sh</code> script. The <code>relocate_replica.sh</code> script preserves the pool order the files are migrated on and will recall then remigrate the files to a new tape within the pool of the warning tape. In the end, the warning tape will be moved out of the pool. Contact IBM Spectrum Archive support to determine the root cause of the warning state.	ltfsee pool remove ltfsee reconcile ltfsee reclaim(source) ltfsee rebuild ltfsee recall
Unavailable	No	Indicates that the cartridge is not available in the IBM Spectrum Archive EE system. A tape that is newly inserted into the tape library is in this state.	If this is a brand new cartridge, run the <b>ltfsee pool add</b> command with the <b>-f</b> option to format the tape and added to the resource. Otherwise, add the tape cartridge to LTFS by running the <b>import</b> command if it contains data.	ltfsee pool remove ltfsee pool add ltfsee import
Invalid	No	The tape cartridge is inconsistent with the LTFS format and must be checked by using the <b>-c</b> option.	Perform a pool add with <b>-c</b> option to attempt to recover the tape.	ltfsee pool add ltfsee pool remove
Unformatted	No	The tape cartridge is not formatted and must be formatted by using the <b>-f</b> option.	Format the tape cartridge.	ltfsee pool add ltfsee pool remove
Inaccessible	No	The tape cartridge is not allowed to move in the library or might be stuck in the drive.	Remove the stuck tape cartridge and fix the cause.	ltfsee pool remove

Tape cartridge status	File system access	Description	How to recover tape cartridges	Valid IBM Spectrum Archive EE commands
Error	No	Indicates that the tape cartridge reported a medium error.	The tape cartridge status returns to Valid by physically removing the medium from the library, then adding it to the library again. If this state occurs again, contact IBM Spectrum Archive support to determine the root cause of the error state.	ltfsee pool remove
Not supported	No	The tape cartridge is an older generation or an LTO Write-Once, Read-Many (WORM) tape cartridge.	LTFS supports the following tape cartridges: <ul style="list-style-type: none"> <li>▶ LTO-8</li> <li>▶ LTO-8 (M8)</li> <li>▶ LTO-7</li> <li>▶ LTO-6</li> <li>▶ LTO-5</li> <li>▶ 3592 Extended data (JD)</li> <li>▶ 3592 Advanced data (JC)</li> <li>▶ 3592 Extended data (JB)</li> <li>▶ 3592 Economy data (JK)</li> <li>▶ 3592 Economy data (JL)</li> <li>▶ 3592 Extended WORM data (JZ)</li> <li>▶ 3592 Advanced WORM data (JY)</li> </ul>	None
Duplicate	No	The tape cartridge has the same bar code as another tape cartridge.	Remove one of the duplicate tape cartridges from the library.	None
Disconnected	No	The Disconnected status indicates that the EE and LE components that are used by this tape cannot communicate. The admin channel connection might be disconnected.	Check the EE and LE components to see whether they are running.	ltfsee pool remove
Unusable	No	The tape has become unusable.	Perform a pool add with <code>-c</code> option to attempt to recover the tape.	ltfsee recall ltfsee export ltfsee rebuild ltfsee reclaim (source/target) ltfsee reconcile ltfsee pool remove

Tape cartridge status	File system access	Description	How to recover tape cartridges	Valid IBM Spectrum Archive EE commands
Write Fenced	Read-only	Indicates that the tape had a write failure but the index was successfully written.	Run the <b>recover</b> command with the <b>-c</b> option to generate a scan list and bring the files back into resident state, then run the <b>recover</b> command using the <b>-r</b> option to double check any missed files. The <b>-r</b> option removes the tape from the <code>ltfsee</code> pool if no new files were detected remaining on tape. Save the tape just in case there are issues recovering files and contact IBM Spectrum Archive support to determine the root cause of the write fenced tape.	<code>ltfsee pool remove</code> <code>ltfsee recover</code> <code>ltfsee tape move</code> <code>ltfsee recall</code>

### Unknown status

This status is only a temporary condition that can be caused when a new tape cartridge is added to the tape library, but was not yet mounted in a tape drive to load the index.

### Write-protected status

This status is caused by setting the write-protection tag on the tape cartridge. If you want to use this tape cartridge in IBM Spectrum Archive EE, you must remove the write protection because a write-protected tape cartridge cannot be added to a tape cartridge pool. After the write-protection is removed, you must run the `ltfsee retrieve` command to update the status of the tape to `Valid LTFS`.

### Critical or Warning status

This status can be caused by actual physical errors on the tape. Automatic recovery has been added to IBM Spectrum Archive V1.2.2. See 10.3, “Recovering data from a write failure tape” on page 284 and 10.4, “Recovering data from a read failure tape” on page 286 for recovery procedures.

### Unavailable status

This status is caused by a tape cartridge being removed from LTFS. The process of adding it to LTFS (see 7.7.1, “Adding tape cartridges” on page 147) changes the status back to `Valid LTFS`. Therefore, this message requires no other corrective action.

### Invalid LTFS status

If an error occurs while writing to a tape cartridge, it might be displayed with an `Invalid LTFS` status that indicates an inconsistent LTFS format. Example 10-8 shows an `Invalid LTFS` status.

*Example 10-8 Check the tape cartridge status*

```
[root@l1fssn1 ~]# ltfsee info tapes
```

Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)	Pool	Library	Address	Drive	Appendable
TST200L6	Valid	L6	2242	0	2242	0	test1	TS3310	4117	-	yes
LTFS01L7	Valid	L7	5338	0	5338	0	test	TS3310	4119	-	yes
PFE001L6	Invalid	L6	2242	95	0	0	test1	TS3310	4118	-	no

This status can be changed back to Valid LTFS by checking the tape cartridge. To do so, run the command that is shown in Example 10-9.

*Example 10-9 Add tape cartridge to pool with check option*

```
[root@l1fssn1 ~]# ltfsee pool add -p test1 -t PFE001L6 -c
GLESLO42I(00640): Adding tape PFE001L6 to storage pool test1.
Tape PFE001L6 successfully checked.
Added tape PFE001L6 to pool test1 successfully.
```

Example 10-10 shows the updated tape cartridge status for JCC541JC.

*Example 10-10 Check the tape cartridge status*

```
[root@l1fssn1 ~]# ltfsee info tapes
```

Tape ID	Status	Type	Capacity(GiB)	Used(GiB)	Free(GiB)	Reclaimable(GiB)	Pool	Library	Address	Drive	Appendable
TST200L6	Valid	L6	2242	0	2242	0	test1	TS3310	4117	-	yes
LTFS01L7	Valid	L7	5338	0	5338	0	test	TS3310	4119	-	yes
PFE001L6	Valid	L6	2242	95	2146	0	test1	TS3310	4118	-	yes

### Unformatted status

This status usually is observed when a scratch tape is added to LTFS without formatting it. It can be fixed by removing and readding it with the **-format** option, as described in 7.7.3, “Formatting tape cartridges” on page 150.

If the tape cartridge was imported from another system, the IBM LTFS Format Verifier can be useful for checking the tape format. For more information about performing diagnostic tests with the IBM LTFS Format Verifier, see 11.3.2, “Using the IBM LTFS Format Verifier” on page 329.

### Inaccessible status

This status is most often the result of a stuck tape cartridge. Removing the stuck tape cartridge and then moving it back to its homeslot, as shown in 7.7.2, “Moving tape cartridges” on page 149, should correct the Inaccessible status.

### Error status

Tape cartridges with an error status can often be the result of errors on the tape. This cartridge cannot be used until the condition is cleared. Stop IBM Spectrum Archive EE and clear the dcache for the files on tape and then restart IBM Spectrum Archive EE, as described in 7.4, “Starting and stopping IBM Spectrum Archive EE” on page 143. If the tape continues to go into an error state contact IBM Spectrum Archive support to determine the root cause.

### Not-supported status

Only LTO-8, M8, 7, 6, 5, and 3592-JB, JC, JD, JK, JL, JY, and JZ tape cartridges are supported by IBM Spectrum Archive EE. This message indicates that the tape cartridge is not one of these types and should be removed from the tape library.

## Write Fenced status

This status is caused by actual physical errors on the tape. However, the index was successfully written on one of the tape's partitions. The process of recovering such a tape removes the tape from the LTFS EE pool and will still be marked as Write Fenced. Save the tape for future reference in case the recover was unsuccessful and contact IBM Spectrum Archive support. See 10.3, "Recovering data from a write failure tape" on page 284 for steps to recover a Write Fenced tape.

## 10.3 Recovering data from a write failure tape

The following are steps for recovering data from a write failure tape, including both Critical and Write Fenced tapes:

1. Verify that the tape is in either Critical or Write Fenced by running **ltfsee info tapes**.
2. Run **ltfsee recover -c** on the Critical/Write Fenced tape to recall all the files on the tape and make them resident again.
3. Run **ltfsee recover -r** on the Critical/Write Fenced tape to perform a final file check on the tape and finally remove the tape from the pool.

If the tape was Critical after step 3, the drive is now unlocked for further use.

Example 10-11 shows the commands and output to recover a Write Fenced tape.

*Example 10-11 Recovering data from a write fenced tape*

---

```
[root@ltfssrv18 ~]# ltfsee info tapes | grep IM1196L6
IM1196L6 Write Fenced LT06 0 0 0 0 primary lib_lib0
4112 - no

[root@ltfssrv18 ~]# ltfsee recover -p primary -t IM1196L6 -c
Scanning GPFS file systems for finding migrated/saved objects in tape IM1196L6.
Tape IM1196L6 has 101 files to be recovered. The list is saved to
/tmp/ltfssrv18.17771.ibm.gpfs.recoverlist.
Bulk recalling files in tape IM1196L6
GLESL268I(00151): 101 file name(s) have been provided to recall.
GLESL263I(00207): Recall result: 101 succeeded, 0 failed, 0 duplicate, 0 not
migrated, 0 not found.
Making 101 files resident in tape IM1196L6.
Changed to resident: 10/101.
Changed to resident: 20/101.
Changed to resident: 30/101.
Changed to resident: 40/101.
Changed to resident: 50/101.
Changed to resident: 60/101.
Changed to resident: 70/101.
Changed to resident: 80/101.
Changed to resident: 90/101.
Changed to resident: 100/101.
Scanning remaining objects migrated/saved in tape IM1196L6.
Scanning non EE objects in tape IM1196L6.
Recovery of tape IM1196L6 is successfully done. 101 files are recovered. The list
is saved to /tmp/ltfssrv18.17771.ibm.gpfs.recoverlist.

[root@ltfssrv18 ~]# ltfsee recover -p primary -t IM1196L6 -r
```

```

Scanning GPFS file systems for finding migrated/saved objects in tape IM1196L6.
Tape IM1196L6 has no file to be recovered.
Removed the tape IM1196L6 from the pool write_fenced.
[root@l1tfssrv18 ~]# l1tfsee info tapes | grep IM1196L6
IM1196L6 Write Fenced LT06 0          0          0          0 -
lib_l1tfssrv18 4112    -   no

```

---

Example 10-12 shows the commands and output to recover data from a Critical tape.

*Example 10-12 Recovering data from a critical tape*

---

```

[root@l1tfssrv18 ~]# l1tfsee info tapes | grep Critical
JD0335JD Critical    TS1150(J5)          9022    0    9022    0
PrimPool lib_l1tfssrv18 260    0000078D8322    no

```

```

[root@l1tfssrv18 ~]# l1tfsee recover -p PrimPool -t JD0335JD -c
Scanning GPFS file systems for finding migrated/saved objects in tape JD0335JD.
Tape JD0335JD has 100 files to be recovered. The list is saved to
/tmp/l1tfssrv18.tuc.stglabs.ibm.com.18168.ibm.glues.recoverlist.
Bulk recalling files in tape JD0335JD
GLES268I(00151): 100 file name(s) have been provided to recall.
GLES263I(00207): Recall result: 100 succeeded, 0 failed, 0 duplicate, 0 not
migrated, 0 not found.
Making 100 files resident in tape JD0335JD.
Changed to resident: 10/100.
Changed to resident: 20/100.
Changed to resident: 30/100.
Changed to resident: 40/100.
Changed to resident: 50/100.
Changed to resident: 60/100.
Changed to resident: 70/100.
Changed to resident: 80/100.
Changed to resident: 90/100.
Changed to resident: 100/100.
Scanning remaining objects migrated/saved in tape JD0335JD.
Scanning non EE objects in tape JD0335JD.
Recovery of tape JD0335JD is successfully done. 100 files are recovered. The list
is saved to /tmp/l1tfssrv18.tuc.stglabs.ibm.com.18168.ibm.glues.recoverlist.

[root@l1tfssrv18 ~]# l1tfsee recover -p PrimPool-t JD0335JD -r
Scanning GPFS file systems for finding migrated/saved objects in tape JD0335JD.
Tape JD0335JD has no file to be recovered.
Removed the tape JD0335JD from the pool PrimPool.

[root@l1tfssrv18 ~]# l1tfsee info tapes | grep JD0335JD
JD0335JD Error      TS1150          0          0          0          0
-          lib_l1tfssrv18 1063    -   no

```

---

## 10.4 Recovering data from a read failure tape

Copy migrated jobs from a Warning tape to a valid tape within the same pool with the following steps:

1. Identify a tape with a read failure by running `ltfsee info tapes` to locate the Warning tape.
2. After the Warning tape has been identified, run the `relocate_replica.sh` script to copy the files from the Warning tape to a new tape within the same pool.
3. After a successful copy, remove the Warning tape from the library and discard it.

The syntax for the script is as follows:

```
relocate_replica.sh -t <tape id> -p <pool name>@<library name>:<pool name>@<library name>[:<pool name>@<library name>] -P <path name>
```

► -t <tape id>

Specifies tape ID to relocate replica from

► -p <pool name>@<library name>[:<pool name>@<library name>]

Specifies pool names and library names to store replicas after running this script

► -P <path name>

Specifies path to GPFS file system to scan

Example 10-13 shows system output of the steps to recover data from a read failure tape.

*Example 10-13 Recovering data from a read failure*

---

```
[root@ltfssrv18 ~]# ltfsee info tapes | grep primary
JD2065JD Warning TS1150(J5) 9022 1 9021 0
primary lib0 1036 - yes
JD2067JD Valid TS1150(J5) 9022 1 9021 0
primary lib0 1038 - yes
JD2066JD Valid TS1150(J5) 9022 0 9022 0
primary lib0 1037 - yes
```

```
[root@ltfssrv18 ~]# ltfsee info files `ls | head -1`
Name: LTFS_EE_FILE_0qIc9F01LsaBmyTvA_b1ljgl.bin
Tape id:JD2065JD@lib0:JD2070JD@lib0:JD2064JD@lib0 Status: premigrated
```

```
[root@ltfssrv18 ~]# ./relocate_replica.sh -t JD2065JD -p
primary@lib0:copy@lib0:copy2@lib0 -P /ibm/gpfs/
1. Getting pool name and library name to which tape JD2065JD belongs
2. Removing specified tape JD2065JD from pool primary@lib0
3. Creating policy file...
4. Performing policy scan...
5. Recalling migrated files to premigrated state
6. Removing replica in target tape JD2065JD
7. Creating replica in alternative tape in pool primary@lib0
8. Creating policy file...
9. Performing policy scan...
All of replicas in tape JD2065JD have been relocated successfully.
```

```
[root@ltfssrv18 ~]# ltfsee info tapes | grep JD2065JD
```

```
JD2065JD Unavailable TS1150          0          0          0          0
-      lib0      1036      -      no
```

```
[root@l1tfssrv18 ~]# l1tfsee info files `ls | head -1`
Name: L1TFS_EE_FILE_OqIc9F01LsaBmyTvA_b1ljgl.bin
Tape id:JD2066JD@lib0:JD2070JD@lib0:JD2064JD@lib0 Status: premigrated
```

**Note:** To obtain the `relocate_replica.sh` script, see Appendix A, “Additional material” on page 333.

## 10.5 Handling export errors

The following are the steps to clean up files referencing exported tapes on the IBM Spectrum Archive file system when there are export errors:

1. Stop the L1TFS EE service by running `l1tfsee stop`.
2. After the process has stopped and verified by running `pidof mmm`, gather all the IBMTPS attributes from the failed export message.
3. Run `l1tfsee_export_fix -T <IBMTPS_attribute> [<IBMTPS_attribute>]`.
4. Start the L1TFS EE service by running `l1tfsee start`.

Example 10-14 shows a typical export error and then follows the steps above in resolving the problem.

*Example 10-14 Fix IBMTPS file pointers on the GPFS file system*

```
[root@l1tfssrv18 ~]# l1tfsee export -p PrimPool -t JD3592JD
GLESS016I(00184): Reconciliation requested.
GLESM401I(00194): Loaded the global configuration.
GLESS049I(00636): Tapes to reconcile: JD3592JD .
GLESS050I(00643): GPFS file systems involved: /ibm/glues .
GLESS134I(00665): Reserving tapes for reconciliation.
GLESS135I(00698): Reserved tapes: JD3592JD .
GLESS054I(00736): Creating GPFS snapshots:
GLESS055I(00741): Deleting previous reconcile snapshot and creating a new one for
/ibm/glues ( gpfs ).
GLESS056I(00762): Scanning GPFS snapshots:
GLESS057I(00767): Scanning GPFS snapshot of /ibm/glues ( gpfs ).
GLESS060I(00843): Processing scan results:
GLESS061I(00848): Processing scan results for /ibm/glues ( gpfs ).
GLESS141I(00861): Removing stale DMAPI attributes:
GLESS142I(00866): Removing stale DMAPI attributes for /ibm/glues ( gpfs ).
GLESS063I(00899): Reconciling the tapes:
GLESS086I(00921): Reconcile is skipped for tape JD3592JD because it is already
reconciled.
GLESS137I(01133): Removing tape reservations.
GLESS058I(02319): Removing GPFS snapshots:
GLESS059I(02326): Removing GPFS snapshot of /ibm/glues ( gpfs ).
Export of tape JD3592JD has been requested...
GLES075E(00660): Export of tape JD3592JD completed with errors. Some GPFS files
still refer files in the exported tape.
GLES073I(00765): Moving tape JD3592JD.
Tape JD3592JD is unmounted because it is inserted into the drive.
```

```
GLESL043I(00151): Removing tape JD3592JD from storage pool replicate.
GLESL631E(00193): Failed to export some tapes.
Tapes (<no tape>) were successfully exported.
Tapes (<no tape>) are still in the pool and needs a retry to export them.
Tapes (JD3592JD) are in Exported state but some GPFS files may still refer files
in these tapes. TPS list to fix are:
JD3592JD065fb41ec-b42d-4fc5-8957-d57e1567aac1@0000013FA0020404
```

```
[root@l1tfssrv18 ~]# l1tfsee stop
Library name: lib_l1tfssrv18, library id: 0000013FA0020404, control node (MMM) IP
address: 9.11.121.249.
Stopped L1TFS EE service (MMM) for library lib_l1tfssrv18.
```

```
[root@l1tfssrv18 ~]# pidof mmm
[root@l1tfssrv18 ~]# l1tfsee_export_fix -T
JD3592JD065fb41ec-b42d-4fc5-8957-d57e1567aac1@0000013FA0020404
Please make sure that IBM Spectrum Archive EE is not running on any node in this
cluster.
And please do not remove/rename any file in all the DMAPI-enabled GPFS if
possible.
Type "yes" to continue.
?
```

```
yes
GLESY016I(00545): Start finding strayed stub files and fix them for GPFS gpfs
(id=14593091431837534985)
GLESY020I(00558): Listing up files that needs to be fixed for GPFS gpfs.
GLESY015I(00531): Fix of exported files completes. Total=985, Succeeded=985,
Failed=0
GLESY018I(00603): Successfully fixed files in GPFS gpfs.
GLESY025I(00615): l1tfsee_export_fix exits with RC=0
[root@l1tfssrv18 ~]# l1tfsee start
Library name: lib_l1tfssrv18, library id: 0000013FA0020404, control node (MMM) IP
address: 9.11.121.249.
GLESM397I(00221): Configuration option: DISABLE_AFM_CHECK yes.
GLESM401I(00264): Loaded the global configuration.
GLESM402I(00301): Created the Global Resource Manager.
GLESM403I(00316): Fetched the node groups from the Global Resource Manager.
GLESM404I(00324): Detected the IP address of the MMM (9.11.121.249).
GLESM405I(00335): Configured the node group (G0).
GLESM406I(00344): Created the unassigned list of the library resources.
GLESL536I(00080): Started the Spectrum Archive EE service (MMM) for library
lib_l1tfssrv18.
```

---

## 10.6 Software

IBM Spectrum Archive EE is composed of four major components, each with its own set of log files. Therefore, problem analysis is slightly more involved than other products. This section describes troubleshooting issues with each component in turn and the Linux operating system and Simple Network Management Protocol (SNMP) alerts.

## 10.6.1 Linux

The log file `/var/log/messages` contains global LINUX system messages, including the messages that are logged during system start and messages that are related to LTFS and IBM Spectrum Archive EE functions. However, three specific log files are also created:

- ▶ `ltfs.log`
- ▶ `ltfsee.log`
- ▶ `ltfsee_trc.log`

Unlike with previous LTFS/IBM Spectrum Archive products, there is no need to enable the system logging on Linux because it is automatically performed during the installation process. Example 10-15 shows the changes to the `rsyslog.conf` file and the location of the two log files.

*Example 10-15 The rsyslog.conf file*

---

```
[root@ltfssn1 ~]# cat /etc/rsyslog.conf | grep ltfs
:msg, startswith, "GLES," /var/log/ltfsee_trc.log;gles_trc_template
:msg, startswith, "GLES" /var/log/ltfsee.log;RSYSLOG_FileFormat
:msg, regex, "LTFS[ID0-9][0-9]*[EWID]" /var/log/ltfs.log;RSYSLOG_FileFormat
```

---

By default, after the `ltfs.log`, `ltfsee.log`, and `ltfsee_trc.log` files reach 1 MB, they are rotated and four copies are kept. Example 10-16 shows the log file rotation settings. These settings can be adjusted as needed within the `/etc/logrotate.d/syslog` control file.

*Example 10-16 Syslog rotation*

---

```
[root@ltfssn1 ~]# cat /etc/logrotate.d/syslog
/var/log/cron
/var/log/maillog
/var/log/messages
/var/log/secure
/var/log/spooler
{
    sharedscripts
    postrotate
    /bin/kill -HUP `cat /var/run/syslogd.pid 2> /dev/null` 2> /dev/null || true
    endscrip
}
/var/log/ltfs.log {
    size 1M
    rotate 4
    missingok
    compress
    sharedscripts
    postrotate
    /bin/kill -HUP `cat /var/run/syslogd.pid 2> /dev/null` 2> /dev/null || true
    /bin/kill -HUP `cat /var/run/rsyslogd.pid 2> /dev/null` 2> /dev/null || true
    endscrip
}
/var/log/ltfsee.log {
    size 1M
    rotate 4
    missingok
    compress
    sharedscripts
```

```

    postrotate
    /bin/kill -HUP `cat /var/run/syslogd.pid 2> /dev/null` 2> /dev/null || true
    /bin/kill -HUP `cat /var/run/rsyslogd.pid 2> /dev/null` 2> /dev/null || true
    endscript
}
/var/log/ltfsee_trc.log {
    size 1M
    rotate 4
    missingok
    compress
    sharedscripts
    postrotate
    /bin/kill -HUP `cat /var/run/syslogd.pid 2> /dev/null` 2> /dev/null || true
    /bin/kill -HUP `cat /var/run/rsyslogd.pid 2> /dev/null` 2> /dev/null || true
    endscript
}

```

---

These log files (`ltfs.log`, `ltfsee.log`, `ltfsee_trc.log`, and `/var/log/messages`) are invaluable in troubleshooting LTFs messages. The `ltfsee.log` file contains only warning and error messages. Therefore, it is easy to start looking here for the reasons of failure. For example, a typical file migration might return the information message that is shown in Example 10-17.

*Example 10-17 Simple migration with informational messages*

---

```

# ltfsee migrate mig -p PrimPool@lib_lto
GLES167I(00400): A list of files to be migrated has been sent to LTFs EE using
scan id 1842482689.
GLES159E(00440): Not all migration has been successful.
GLES1038I(00448): Migration result: 0 succeeded, 1 failed, 0 duplicate, 0
duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too
early for migration.

```

---

From the `GLES159E` message, you know that the migration was unsuccessful, but you do not know *why* it was unsuccessful. To understand why, you must examine the `ltfsee.log` file. Example 10-18 shows the end of the `ltfsee.log` file immediately after the **failed migrate** command is run.

*Example 10-18 The ltfsee.log file*

---

```

# tail /var/log/ltfsee.log
2016-12-14T09:05:38.494320-07:00 ltfs97 mmm[7889]: GLES600E(01691): Failed to
migrate/premigrate file /ibm/gpfs/file1.mpeg. The specified pool name does not
match the existing replica copy.
2016-12-14T09:05:48.500848-07:00 ltfs97 ltfsee[29470]: GLES159E(00440): Not all
migration has been successful.

```

---

In this case, the migration of the file was unsuccessful because it was previously migrated/premigrated to a different tape cartridge.

With IBM Spectrum Archive EE, there are two logging facilities. One is in a human-readable format that is monitored by users and the other is in machine-readable format that is used for further problem analysis. The former facility is logged in to `/var/log/ltfsee.log` through the “user” syslog facility and contains only warnings and errors. The latter facility is logged in to `/var/log/ltfsee_trc.log` through the “local2” Linux facility.

The messages in machine-readable format can be converted into human-readable format by the newly created tool `ltfsee_catcsvlog`, which is run by the following command:

```
/opt/ibm/ltfsee/bin/ltfsee_catcsvlog /var/log/ltfsee_trc.log
```

The `ltfsee_catcsvlog` command accepts multiple log files as command-line arguments. If no argument is specified, `ltfsee_catcsvlog` reads from `stdin`.

## Persistent problems

This section describes ways to solve persistent IBM Spectrum Archive EE problems.

If an unexpected and persistent condition occurs in the IBM Spectrum Archive EE environment, contact your IBM service representative. Provide the following information to help IBM re-create and solve the problem:

- ▶ Machine type and model of your IBM tape library in use for IBM Spectrum Archive EE
- ▶ Machine type and model of the tape drives that are embedded in the tape library
- ▶ Specific IBM Spectrum Archive EE version
- ▶ Description of the problem
- ▶ System configuration
- ▶ Operation that was performed at the time the problem was encountered

The operating system automatically generates system log files after initial configuration of the IBM Spectrum Archive EE. Provide the results of the `ltfsee_log_collection` command to your IBM service representative.

## 10.6.2 IBM Spectrum Scale

IBM Spectrum Scale writes operational messages and error data to the IBM Spectrum Scale log file. The IBM Spectrum Scale log can be found in the `/var/adm/ras` directory on each node. The IBM Spectrum Scale log file is named `mmfs.log.date.nodeName`, where `date` is the time stamp when the instance of IBM Spectrum Scale started on the node and `nodeName` is the name of the node. The latest IBM Spectrum Scale log file can be found by using the symbolic file name `/var/adm/ras/mmfs.log.latest`.

The IBM Spectrum Scale log from the prior start of IBM Spectrum Scale can be found by using the symbolic file name `/var/adm/ras/mmfs.log.previous`. All other files have a time stamp and node name that is appended to the file name.

At IBM Spectrum Scale start, files that were not accessed during the last 10 days are deleted. If you want to save old files, copy them elsewhere.

Example 10-19 shows normal operational messages that appear in the IBM Spectrum Scale log file.

*Example 10-19 Normal operational messages in an IBM Spectrum Scale log file*

---

```
[root@ltfs97 ]# cat /var/adm/ras/mmfs.log.latest
Wed Apr  3 13:25:04 JST 2013: runmmfs starting
Removing old /var/adm/ras/mmfs.log.* files:
Unloading modules from /lib/modules/2.6.32-220.el6.x86_64/extra
Loading modules from /lib/modules/2.6.32-220.el6.x86_64/extra
Module           Size  Used by
mmfs26            1749012  0
mmfslinux         311300  1 mmfs26
tracedev          29552   2 mmfs26,mmfslinux
```

```
Wed Apr 3 13:25:06.026 2013: mmfsd initializing. {Version: 3.5.0.7   Built: Dec
12 2012 19:00:50} ...
Wed Apr 3 13:25:06.731 2013: Pagepool has size 3013632K bytes instead of the
requested 29360128K bytes.
Wed Apr 3 13:25:07.409 2013: Node 192.168.208.97 (htohru9) is now the Group
Leader.
Wed Apr 3 13:25:07.411 2013: This node (192.168.208.97 (htohru9)) is now Cluster
Manager for htohru9.ltd.sdl.
```

```
Starting ADSM Space Management daemons
Wed Apr 3 13:25:17.907 2013: mmfsd ready
Wed Apr 3 13:25:18 JST 2013: mmcommon mmfsup invoked. Parameters: 192.168.208.97
192.168.208.97 all
Wed Apr 3 13:25:18 JST 2013: mounting /dev/gpfs
Wed Apr 3 13:25:18.179 2013: Command: mount gpfs
Wed Apr 3 13:25:18.353 2013: Node 192.168.208.97 (htohru9) appointed as manager
for gpfs.
Wed Apr 3 13:25:18.798 2013: Node 192.168.208.97 (htohru9) completed take over
for gpfs.
Wed Apr 3 13:25:19.023 2013: Command: err 0: mount gpfs
Wed Apr 3 13:25:19 JST 2013: finished mounting /dev/gpfs
```

---

Depending on the size and complexity of your system configuration, the amount of time to start IBM Spectrum Scale varies. Taking your system configuration into consideration, if you cannot access a file system that is mounted (automatically or by running a **mount** command) after a reasonable amount of time, examine the log file for error messages.

The IBM Spectrum Scale log is a repository of error conditions that were detected on each node, and operational events, such as file system mounts. The IBM Spectrum Scale log is the first place to look when you are attempting to debug abnormal events. Because IBM Spectrum Scale is a cluster file system, events that occur on one node might affect system behavior on other nodes, and all IBM Spectrum Scale logs can have relevant data.

A common error that might appear when trying to mount GPFS is that it cannot read superblock. Example 10-20 shows the output of the error when trying to mount GPFS.

*Example 10-20 Superblock error from mounting GPFS*

---

```
[root@l1tfsml1 ~]# mmmount gpfs
Wed May 24 12:53:59 MST 2017: mmmount: Mounting file systems ...
mount: gpfs: can't read superblock
mmmount: Command failed. Examine previous error messages to determine cause.
```

---

The cause of this error and failure to mount GPFS is that the GPFS file system had dmapi enabled, but the HSM process has not been started. To get around this error and successfully mount GPFS, issue the **systemctl start hsm** command, and make sure it is running by issuing **systemctl status hsm**. After HSM is running, wait for the recall processes to initiate. This process can be viewed by issuing **ps -afe | grep dsm**. Example 10-21 shows output of starting HSM, checking the status, and mounting GPFS.

*Example 10-21 Starting HSM and mounting GPFS*

---

```
[root@l1tfsml1 ~]# systemctl start hsm
[root@l1tfsml1 ~]# systemctl status hsm
? hsm.service - HSM Service
```

```

Loaded: loaded (/usr/lib/systemd/system/hsm.service; enabled; vendor preset:
disabled)
Active: active (running) since Wed 2017-05-24 13:04:59 MST; 4s ago
Main PID: 16938 (dsmwatchd)
CGroup: /system.slice/hsm.service
        +-16938 /opt/tivoli/tsm/client/hsm/bin/dsmwatchd nodetach

May 24 13:04:59 ltfsml1.tuc.stglabs.ibm.com systemd[1]: Started HSM Service.
May 24 13:04:59 ltfsml1.tuc.stglabs.ibm.com systemd[1]: Starting HSM Service...
May 24 13:04:59 ltfsml1.tuc.stglabs.ibm.com dsmwatchd[16938]: HSM(pid:16938):
start
[root@ltfsml1 ~]# ps -afe | grep dsm
root      7906      1  0 12:56 ?          00:00:00
/opt/tivoli/tsm/client/hsm/bin/dsmwatchd nodetach
root      9748      1  0 12:57 ?          00:00:00 dsmreca1ld
root      9773    9748  0 12:57 ?          00:00:00 dsmreca1ld
root      9774    9748  0 12:57 ?          00:00:00 dsmreca1ld
root      9900  26012  0 12:57 pts/0    00:00:00 grep --color=auto dsm
[root@ltfsml1 ~]# mmmount gpfs
Wed May 24 12:57:22 MST 2017: mmmount: Mounting file systems ...
[root@ltfsml1 ~]# df -h | grep gpfs
gpfs                280G  154G  127G  55%
/ibm/glues

```

---

If HSM is already running, double check if the dsmreca1ld daemons are running by issuing **ps -afe | grep dsm**. If no dsmreca1ld daemons are running, start them by issuing **dsmmigfs start**. After they have been started, GPFS can be successfully mounted.

### 10.6.3 IBM Spectrum Archive LE+ component

This section describes the options that are available to analyze problems that are identified by the LTFS logs. It also provides links to messages and actions that can be used to troubleshoot the source of an error.

The messages that are referenced in this section provide possible actions only for solvable error codes. The error codes that are reported by LTFS program can be retrieved from the terminal console or log files. For more information about retrieving error messages, see 10.6.1, “Linux” on page 289.

When multiple errors are reported, LTFS attempts to find a message ID and an action for each error code. If you cannot locate a message ID or an action for a reported error code, LTFS encountered a critical problem. If you try an initial action again and continue to fail, LTFS also encountered a critical problem. In these cases, contact your IBM service representative for more support.

Message ID strings start with the keyword LTFS and are followed by a four- or five-digit value. However, some message IDs include the uppercase letter I or D after LTFS, but before the four- or five-digit value. When an IBM Spectrum Archive EE command is run and returns an error, check the message ID to ensure that you do not mistake the letter I for the numeral 1.

A complete list of all LTFS messages can be found in the IBM Spectrum Archive EE section of IBM Knowledge Center, which is available on the following website:

[https://www.ibm.com/support/knowledgecenter/ST9MBR\\_1.2.6/ltfs\\_ee\\_messages.html](https://www.ibm.com/support/knowledgecenter/ST9MBR_1.2.6/ltfs_ee_messages.html)

At the end of the message ID, the following single capital letters indicate the importance of the problem:

- ▶ E: Error
- ▶ W: Warning
- ▶ I: Information
- ▶ D: Debugging

When you troubleshoot, check messages for errors only.

Example 10-22 shows a problem analysis procedure for LTFS.

*Example 10-22 LTFS messages*

---

```
cat /var/log/lufs.log
LTFS11571I State of tape '' in slot 0 is changed from 'Not Initialized' to
'Non-supported'
LTFS14546W Cartridge '' does not have an associated LTFS volume
LTFS11545I Rebuilding the cartridge inventory
LTFSI1092E This operation is not allowed on a cartridge without a bar code
LTFS14528D [localhost at [127.0.0.1]:34888]: --> CMD_ERROR
```

---

The set of 10 characters represents the message ID, and the text that follows describes the operational state of LTFS. The fourth message ID (LTFSI1092E) in this list indicates that an error was generated because the last character is the letter E. The character immediately following LTFS is the letter I. The complete message, including an explanation and appropriate course of action for LTFSI1092E, is shown in the following example:

Example 10-23.

*Example 10-23 Example of message*

---

```
LTFSI1092E This operation is not allowed on a cartridge without a bar code
Explanation
A bar code must be attached to the medium to perform this operation.
Action
Attach a bar code to the medium and try again.
```

---

Based on the description that is provided here, the tape cartridge in the library does not have a bar code. Therefore, the operation is rejected by LTFS. The required user action to solve the problem is to attach a bar code to the medium and try again.

## 10.6.4 Hierarchical storage management

During installation, hierarchical storage management (HSM) is configured to write log entries to a log file in /opt/tivoli/tsm/client/hsm/bin/dsmerror.log. Example 10-24 shows an example of this file.

*Example 10-24 The dsmerror.log file*

---

```
[root@lufs97 /]# cat dsmerror.log
03/29/2013 15:24:28 ANS9101E Migrated files matching '/ibm/glues/file1.img' could not be found.
03/29/2013 15:24:28 ANS9101E Migrated files matching '/ibm/glues/file2.img' could not be found.
03/29/2013 15:24:28 ANS9101E Migrated files matching '/ibm/glues/file3.img' could not be found.
03/29/2013 15:24:28 ANS9101E Migrated files matching '/ibm/glues/file4.img' could not be found.
03/29/2013 15:24:28 ANS9101E Migrated files matching '/ibm/glues/file5.img' could not be found.
03/29/2013 15:24:28 ANS9101E Migrated files matching '/ibm/glues/file6.img' could not be found.
04/02/2013 16:24:06 ANS9510E dsmrecalld: cannot get event messages from session 515A6F7E00000000, expected
max message-length = 1024, returned message-length = 144. Reason : Stale NFS file handle
```

```
04/02/2013 16:24:06 ANS9474E dsmrecalld: Lost my session with errno: 1 . Trying to recover.
04/02/13 16:24:10 ANS9433E dsmwatchd: dm_send_msg failed with errno 1.
04/02/2013 16:24:11 ANS9433E dsmrecalld: dm_send_msg failed with errno 1.
04/02/2013 16:24:11 ANS9433E dsmrecalld: dm_send_msg failed with errno 1.
04/02/2013 16:24:11 ANS9433E dsmrecalld: dm_send_msg failed with errno 1.
04/03/13 13:25:06 ANS9505E dsmwatchd: cannot initialize the DMAPI interface. Reason: Stale NFS file
handle
04/03/2013 13:38:14 ANS1079E No file specification entered
04/03/2013 13:38:20 ANS9085E dsmrecall: file system / is not managed by space management.
```

---

The HSM log contains information about file migration and recall, threshold migration, reconciliation, and starting and stopping the HSM daemon. You can analyze the HSM log to determine the current state of the system. For example, the logs can indicate when a recall has started but not finished within the last hour. The administrator can analyze a particular recall and react accordingly.

In addition, an HSM log might be analyzed by an administrator to optimize HSM usage. For example, if the HSM log indicates that 1,000 files are recalled at the same time, the administrator might suggest that the files can be first compressed into one .tar file and then migrated.

## 10.6.5 IBM Spectrum Archive EE logs

This section describes IBM Spectrum Archive EE logs and message IDs and provide some tips for dealing with failed recalls and lost or strayed files.

### IBM Spectrum Archive EE log collection tool

IBM Spectrum Archive EE writes its logs to the files `/var/log/ltfsee.log` and `/var/log/ltfsee_trc.log`. These files can be viewed in a text editor for troubleshooting purposes. Use the IBM Spectrum Archive EE log collection tool to collect data that you can send to IBM Support.

The `ltfsee_log_collection` tool is in the `/opt/ibm/ltfsee/bin` folder. To use the tool, complete the following steps:

1. Log on to the operating system as the root user and open a console.
2. Start the tool by running the following command:  

```
# /opt/ibm/ltfsee/bin/ltfsee_log_collection
```
3. When the following message displays, read the instructions, then enter `y` or `p` to continue:

```
LTFS Enterprise Edition - log collection program
This program collects the following information from your GPFS cluster.
a. Log files that are generated by GPFS, LTFS Enterprise Edition
b. Configuration information that is configured to use GPFS and LTFS Enterprise
   Edition
c. System information including OS distribution and kernel, and hardware
   information (CPU and memory)
```

If you want to collect all the information, enter `y`.

If you want to collect only a and b, enter `p` (partial).

If you do not want to collect any information, enter `n`.

The collected data is compressed in the `ltfsee_log_files_<date>_<time>.tar.gz` file. You can check the contents of the file before submitting it to IBM.

4. Make sure that a packed file with the name `ltfsee_log_files_[date]_[time].tar.gz` is created in the current directory. This file contains the collected log files.
5. Send the `tar.gz` file to your IBM service representative.

## Messages reference

For IBM Spectrum Archive EE, message ID strings start with the keyword GLES and are followed by a single letter and then by a three-digit value. The single letter indicates which component generated the message. For example, GLESL is used to indicate all messages that are related to the IBM Spectrum Archive EE command. At the end of the message ID, the following single uppercase letter indicates the importance of the problem:

- ▶ E: Error
- ▶ W: Warning
- ▶ I: Information
- ▶ D: Debugging

When you troubleshoot, check messages for errors only. For a list of available messages, see this website:

[https://www.ibm.com/support/knowledgecenter/ST9MBR\\_1.2.6/ltfs\\_ee\\_messages.html](https://www.ibm.com/support/knowledgecenter/ST9MBR_1.2.6/ltfs_ee_messages.html)

## Failed reconciliations

Failed reconciliations usually are indicated by the GLESS003E error message with the following description:

Reconciling tape %s failed due to a generic error.

## Lost or strayed files

Table 10-2 describes the seven possible status codes for files in IBM Spectrum Archive EE. They can be viewed for individual files by running the `ltfsee info files` command.

*Table 10-2 Status codes for files in IBM Spectrum Archive EE*

Status code	Description
Resident	The Resident status indicates that the file is resident in the GPFS namespace and is not saved, migrated, or premigrated to a tape.
Migrated	The file was migrated. The file was copied from a GPFS file system to a tape, and exists only as a stub file in the IBM Spectrum Scale namespace.
Premigrated	The file was partially migrated. An identical copy exists on your local file system and in tape.
Saved	The Saved status indicates that the file system object that has no data (a symbolic link, an empty directory, or an empty regular file) was saved. The file system object was copied from GPFS file system to a tape.
Offline	The file was migrated to a tape cartridge and then the tape cartridge was exported offline.
Lost	The file was in the migrated status, but the file is not accessible from IBM Spectrum Archive EE because the tape cartridge that the file is supposed to be on is not accessible. The file might be lost because of tape corruption or if the tape cartridge was removed from the system without first exporting it.
Strayed	The file was in the premigrated status, but the file is not accessible from IBM Spectrum Archive EE because the tape that the file is supposed to be on is not accessible. The file might be lost because of tape corruption or if the tape cartridge was removed from the system without first exporting.

The only two status codes that indicate an error are `lost` and `strayed`. Files in these states should be fixed where possible by returning the missing tape cartridge to the tape library or by attempting to repair the tape corruption. For more information, see 7.20, “Checking and repairing” on page 195. If this is not possible, they should be restored from a redundant copy. For more information, see 7.14.2, “Selective recall” on page 184. If a redundant copy is not available, the stub files must be deleted from the GPFS file system.

## 10.7 Recovering from system failures

The system failures that are described in this section are the result of hardware failures or temporary outages that result in IBM Spectrum Archive EE errors.

### 10.7.1 Power failure

When a library power failure occurs, the data on the tape cartridge that is actively being written is probably left in an inconsistent state.

To recover a tape cartridge from a power failure, complete the following steps:

1. Create a mount point for the tape library. For more information, see the procedure described in 7.2.2, “IBM Spectrum Archive Library Edition Plus component” on page 136.
2. If you do not know which tape cartridges are in use, try to access all tape cartridges in the library. If you do know which tape cartridges are in use, try to access the tape cartridge that was in use when the power failure occurred.
3. If a tape cartridge is damaged, it is identified as inconsistent and the corresponding subdirectories disappear from the file system. You can confirm which tape cartridges are damaged or inconsistent by running the `ltfsee info tapes` command. The list of tape cartridges that displays indicates the volume name, which is helpful in identifying the inconsistent tape cartridge. For more information, see 7.20, “Checking and repairing” on page 195.
4. Recover the inconsistent tape cartridge by running the `ltfsee pool add` command with the `-c` option. For more information, see 7.20, “Checking and repairing” on page 195.

### 10.7.2 Mechanical failure

When a library receives an error message from one of its mechanical parts, the process to move a tape cartridge cannot be performed.

**Important:** A drive in the library normally performs well despite a failure so that ongoing access to an opened file on the loaded tape cartridge is not interrupted or damaged.

To recover a library from a mechanical failure, complete the following steps:

1. Run the `umount` or `fusermount -u` command to start a tape cartridge unload operation for each tape cartridge in each drive.

**Important:** The `umount` command might encounter an error when it tries to move a tape cartridge from a drive to a storage slot.

To confirm that the status of all tapes, access the library operator panel or the web user interface to view the status of the tape cartridges.

2. Turn off the power to the library and remove the source of the error.
3. Follow the procedure that is described in 10.7.1, “Power failure” on page 297.

**Important:** One or more inconsistent tape cartridges might be found in the storage slots and might need to be made consistent by following the procedure that is described in “Unavailable status” on page 282.

### 10.7.3 Inventory failure

When a library cannot read the tape cartridge bar code for any reason, an inventory operation for the tape cartridge fails. The corresponding media folder does not display, but a specially designated folder that is named UNKN0000 is listed instead. This designation indicates that a tape cartridge is not recognized by the library.

If the user attempts to access the tape cartridge contents, the media folder is removed from the file system. The status of any library tape cartridge can be determined by running the **ltfsee info tapes** command. For more information, see 7.25, “Obtaining inventory, job, and scan status” on page 210.

To recover from an inventory failure, complete the following steps:

1. Remove any unknown tape cartridges from the library by using the operator panel or Tape Library Specialist web interface, or by opening the door or magazine of the library.
2. Check all tape cartridge bar code labels.

**Important:** If the bar code is removed or about to peel off, the library cannot read it. Replace the label or firmly attach the bar code to fix the problem.

3. Insert the tape cartridge into the I/O station.
4. Check to determine whether the tape cartridge is recognized by running the **ltfsee info tapes** command.
5. Add the tape cartridge to the LTFS inventory by running the **ltfsee pool add** command.

### 10.7.4 Abnormal termination

If LTFS terminates because of an abnormal condition, such as a system hang-up or after the user initiates a **kill** command, the tape cartridges in the library might remain in the tape drives. If this occurs, LTFS locks the tape cartridges in the drives and the following command is required to release them:

```
# ltf s release_device -o changer_devname=[device name]
```



# Reference

This chapter describes the commands that are used to operate IBM Spectrum Archive EE, data, and metadata formats for IBM Spectrum Scale to LTFS migrations, system calls, tools, and limitations.

This chapter includes the following topics:

- ▶ Command-line reference
- ▶ Formats for IBM Spectrum Scale to IBM Spectrum Archive EE migration
- ▶ System calls and IBM tools
- ▶ IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products

## 11.1 Command-line reference

This section describes the LTFSE Enterprise Edition (IBM Spectrum Archive EE) commands, IBM Spectrum Scale commands, and Tivoli Storage Manager space management commands.

### 11.1.1 IBM Spectrum Archive EE commands

This section describes the syntax, parameters, and function of the IBM Spectrum Archive EE commands.

You must have root user authority to run the IBM Spectrum Archive EE commands, except for the `ltfsee info` command. Some options have a short format with a “-” followed by a letter and a long format that is “--” followed by the option name, such as `-P` and `--path`, which is shown in the `ltfsee import` command in Example 11-1.

*Example 11-1 Short and long format for command options*

---

```
ltfsee import -p pool1 -t D00279L5 1FA682L5 -P /importUser
ltfsee import -p pool1 -t D00279L5 1FA682L5 --path /importUser
```

---

#### The `ltfsee cleanup` command

Run the `ltfsee cleanup` command to clean up scan results that are not removed after a scan or session finishes. For example, you can run this command to clean up any scan numbers that remain if a user runs a `ltfsee migrate`, `ltfsee premigrate`, or `ltfsee save` command, and then closes the command window before the command completes, as shown in the following example:

```
ltfsee cleanup -s <scan number> [-l library_name]
```

Option `-l` must be provided on multi-library systems and can be omitted on single-library systems. This command cleans up scan numbers that remain after a scan or session finishes, but cannot be removed by its calling process. The only scans that can be cleaned up are scans that show 0 items within the “#items Queued” column of the IBM Spectrum Archive EE information scans command output, as shown in the following example:

```
ltfsee cleanup -s 3603955717
```

#### The `ltfsee datamigrate` command

Run the `ltfsee datamigrate` command to perform pool-to-pool migration within the same node group. For example, you can run this command to move all of your data from an older generation media pool to a newer generation media pool. The following example shows the syntax and options for the `ltfsee datamigrate` command:

```
datamigrate -p <poolname> [-l <libraryname>] -d <destination_pool> [-n <number> |
-t <tape_id_1 ... tape_id_N> [-L[file_number_th]] [-q]
```

Required parameters

**-p <poolname>**

Specifies the name of the source tape storage pool from which you want to migrate data. This parameter must be specified as the first parameter.

If no other parameters are specified, all the tapes in the source tape storage pool are selected to have their data migrate to the destination pool.

Only tapes that belong to the specified tape storage pool can have their data migrate to the destination pool. Any tapes that are specified with the `-t` option that do not belong to the specified tape storage pool are ignored.

**-l <libraryname>**

The name of the tape library to which the source and destination tape pools belong. This option is required in multiple-library systems. In single-library systems, it can be omitted.

**-d <destination\_pool>**

The pool in which the data that was selected for migration is rewritten to new tapes. The tapes are written in the format type, and on the cartridge type, that is specified for the destination pool.

### Optional parameters

**-n <number>**

The specified number of tapes from which to migrate data. If `-n` is specified, up to `<number>` tapes from the source tape storage pool `<poolname>` have their data migrate to the destination pool. After the data is migrated, the source tapes are removed from their storage pool. If less than `<number>` tapes successfully migrate their data, only the successfully migrated tapes are removed. The `-n` option cannot be used with the `-t` option.

**-t <tape\_id\_1 ... tape\_id\_N>**

Specifies the tapes from which to migrate data, then remove from the source tape storage pool. If `-t` is specified, the listed tapes `<tape_id_1 ... tape_id_N>` from source tape storage pool `<poolname>` have their data migrate to tapes in the destination pool, then the listed tapes are removed from the source pool.

If not all of the listed tapes successfully migrate their data, only the successfully migrated tapes are removed from the pool.

All listed tapes must belong to the specified tape storage pool.

The `-t` option cannot be used with the `-n` option.

**-L[<file\_number\_th>]**

Defines the number of files to migrate to the destination pool.

Use the `-L` option to limit the amount of processing time that the data migration operation takes. The `-L` option works with the `-t` option with only one tape ID. When this option is used, the tape might not be removed from the storage pool, the storage pool might not gain capacity, while the unreferenced capacity of the tape increases. The tape might require more data migration operations to completely move all of the data from the tape to the destination tape pool. If this option is specified without a threshold value, the default value is 100,000.

**Note:** This option must not be separated from the argument value by a space. Example usage: `-L5000`.

**-q**

The **-q** option runs a quick-reconcile function, before the data migration function, which resolves simple inconsistency problems between the IBM Spectrum Scale file system and the IBM Spectrum Archive EE tapes. Files with inconsistencies that quick reconcile cannot resolve are left on the source tape. The **reconcile** command needs to be run on those files that are left on the source tape. All DMAPI-enabled IBM Spectrum Scale file systems must be mounted to use this option.

## The **ltfsee drive** command

Run the **ltfsee drive** command to add a tape drive to or remove a tape drive from the IBM Spectrum Archive EE system, as shown in the following example:

```
ltfsee drive <add | remove> -d <drive_serial[:attr]> [-n <node_id>] [-l  
<library_name>]
```

**Important:** A drive that is already added to a node cannot be added to another node. To change the node on which that drive is located, first remove the drive by running the **ltfsee drive remove** command and then add the drive to the wanted node by running the **ltfsee drive add** command.

This command adds a specified tape drive to the IBM Spectrum Archive EE system. When you add a tape drive, you must specify values for the following variables:

► **drive\_serial[:attr]**

The drive serial number of the drive that you want to add. The **drive\_serial** value has a maximum length of 32 characters.

**attr** is a decimal numeric parameter that is optional for this command and can be specified after “:”. This optional parameter is only valid for IBM Spectrum Archive EE versions V1R1.1.3 or later. A logical **OR** applies to set the three attributes: Migrate (4), Recall (2), and Generic (1).

For example, a number of 6 for **attr** allows migration and recall jobs, whereas generic jobs are disallowed. All of the attributes are set by default. If the tape drive to update has attributes set by IBM Spectrum Archive EE, you must remove them before adding new ones by using the **ltfsee drive remove** command. For more information about the usage of the drive role, see 7.22, “Drive Role settings for job assignment control” on page 203.

► **node\_id**

The node ID of the cluster to which you want to add the tape drive.

The **drive remove -d <drive\_serial>** command removes a specified tape drive from the IBM Spectrum Archive EE system. When you remove a tape cartridge, you must specify a value for the **drive\_serial** variable. If a tape cartridge is in the specified tape drive and a job is in progress, the tape cartridge is unloaded automatically when the job completes. The **drive\_serial** is the drive serial number of the drive that you want to remove, as shown in the following examples:

```
ltfsee drive add -d 1068002111 -n 1  
ltfsee drive remove -d 1068002111
```

## The ltfsee export command

Run the **ltfsee export** command to remove one or more tape cartridges from the IBM Spectrum Archive EE system by removing files on them from the IBM Spectrum Scale namespace, as shown in the following example:

```
ltfsee export -p <pool_name> [ -l <library_name> ] [ -t <tape_id_1 tape_id_2 ... tape_id_N> ] [-o <message>]
```

The tape cartridge is reserved so that it is no longer a target for file migrations. It is then reconciled to remove any inconsistencies between it and IBM Spectrum Scale. The process then removes all files from the Spectrum Scale file system that exist on the exported tape cartridge and remove the tape from the cartridge pool. The files on tape cartridges are unchanged by the export and are accessible by other LTFSS systems.

If the **--offline** option is specified, the inodes that refer to the files on the tapes will not be removed, but just set to an offline status. This feature means that the files remain visible in the GPFS namespace, but are no longer accessible. Tapes that have been exported with the **--offline** option can be imported by using the **--offline** option of the **import** command. The **export** command includes the following components:

- ▶ **-t, --tape <tape\_ids>**  
List of tapes to export.
- ▶ **-p, --pool <poolName>**  
Storage pool defining the list of tapes to export.
- ▶ **-l, --library <library\_name>**  
The library name, which can be omitted on single library systems.
- ▶ **-o, --offline <message>**  
Indicates that the tapes are exported into offline mode, meaning that the files remain visible in the GPFS namespace. The message, which must not exceed 500 characters, is stored in a GPFS extended attribute.

The following are some examples:

```
ltfsee export -p pool1 -t D00279L5 1FA682L5
ltfsee export -p pool1 -l lib1
ltfsee export -p pool1 -o "Moved to storage room B"
```

In the case where an error occurs during the normal export command, the state of the tape becomes Exported. However, a cleanup is still necessary because some of the files on the file system can still reference the tape that was exported. See 10.5, "Handling export errors" on page 287 on how to clean up the IBMTPS file pointers.

## The ltfsee failover command

Run the **ltfsee failover** command to manually initiate the failover process from the active control node to the non-active control node:

```
ltfsee failover [-l <libraryname>]
```

The ltfsee failover command takes this variable:

- ▶ **-l <libraryname>**  
The name of the library name in which you want the active control node failover to start, which can be omitted on single library systems.

## The ltfsee fsopt command

To query or update the file system level settings for stub size, read starts recalls, and preview size, run the **ltfsee fsopt** command:

```
ltfsee fsopt { query [-G <gpfs_filesystem(s)>] } | { update [ -S <stub_size> ]  
[ -R <read_starts_recall> ] [-P <preview_size>] [-F] -G <gpfs_filesystem(s)> }
```

To query the current file system level settings, run the **ltfsee fsopt query** command:

```
ltfsee fsopt query -G /ibm/gpfs
```

The **-G <gpfs\_filesystem(s)>** option is not required, but can be used to specify one or more GPFS file systems to query. Multiple file system names must be separated by a space.

To update the file system level settings for stub size, read start recalls, or preview size, run the following command:

```
ltfsee fsopt update [-S <stub_size>] [-R <read_starts_recalls>]  
[-P <preview_size>] [-F] -G <gpfs_filesystem(s)>
```

You must select one or more of the following parameters when you run the **ltfsee fsopt update** command:

► **-S <stub\_size>**

Defines the size of the initial file part that is kept resident on disk for migrated files. The value must be a multiple of the file system block size, and larger than or equal to the preview size.

Possible values: 0 - 1073741824

► **-R <read\_starts\_recall>**

When this feature is set to yes, reading from the resident file part starts a background recall of the file. During the background recall, data from the resident part can be read. The rest of the file can be read upon recall completion.

Possible values: yes | no | undef

► **-P <preview\_size>**

Defines the initial file part size for which reads from the resident file part do not trigger recalls. The value must be smaller than or equal to the stub size.

Possible values: 0 - 1073741824

In addition, you can specify the **-F** option, which forces the settings to update even if a stub size change is requested while there are ongoing migrations.

**Important:** Using the **-F** option might result in some ongoing migrations that use the old stub size while others use the new stub size.

The **-G <gpfs\_filesystem(s)>** option is not required, but can be used to specify one or more GPFS file systems to update. Multiple file system names must be separated by a space, as shown in the following example:

```
ltfsee fsopt update -S 10485760 -P 81920 -R yes -G /ibm/gpfs
```

## The ltfsee import command

To add one or more tape cartridges to the IBM Spectrum Archive EE system and reconstitute the files in the GPFS namespace, run the **ltfsee import** command, as shown in the following example. For more information about managing imports with the **ltfsee import** command and its options, see 7.21.1, “Importing” on page 196.

```
ltfsee import -p <pool_name> [ -l <library_name> ] -t <tape_id_1 tape_id_2 ... tape_id_N>
{ [[ -P <pathName> ] [ -o | -i | -r ] [ -R ]] | [ --offline ] }
```

Imported files will be in a migrated state, meaning that the data remains on tape. The data portion of the file is not copied to disk during the import.

Furthermore, you can specify one of the following optional parameters with the **ltfsee import** command:

▶ **-p, --pool <pool\_name>**

The storage pool to which the imported tapes are assigned.

▶ **-l, --library <library\_name>**

The library name, which can be omitted on single-library systems.

▶ **-t, --tape <tape\_ids>**

The names (IDs) of the tapes to import.

▶ **-P, --path {pathName}**

Specifies the path name where the files are imported. If omitted, the import uses the default path (`/<GPFS_FileSystem>/IMPORT`). If the path does not exist, it is created. If multiple GPFS file systems are mounted, this option is mandatory. For more information, see “Restoring file system objects and files from tape” on page 251.

▶ **-o, --overwrite**

Indicates that files are overwritten if they exist already. This option is mutually exclusive with **-i**, **-r**, and **--offline**. For more information, see “Restoring file system objects and files from tape” on page 251.

▶ **-i, --ignore**

Indicates that files are ignored if they exist already. This option is mutually exclusive with **-o**, **-r**, and **--offline**. For more information, see “Restoring file system objects and files from tape” on page 251.

▶ **-r, --rename**

Indicates that files are renamed if they exist already. This option is mutually exclusive with **-i**, **-o**, and **--offline**. For more information, see “Restoring file system objects and files from tape” on page 251.

▶ **-R, --recreate**

Indicates that files are imported in the specified path, without creating a tape ID directory. For more information, see “Restoring file system objects and files from tape” on page 251.

▶ **--offline**

Indicates that the tapes were exported with the **--offline** option. All files are restored to the original exported path. If any file was modified in the GPFS namespace or on tape, it will not be imported. This option is mutually exclusive with **-p**, **-o**, **-i**, **-r**, and **-R**, as shown in the following examples:

– **ltfsee import -p pool1 -l lib1 -t D00279L5** (single or multiple libraries system)

– **ltfsee import -p pool1 -t D00279L5** (in a single library system, the library name can be omitted)

- ltfsee import -p pool1 -l lib1 -t D00279L5 1FA682L5 -P /importUser
- ltfsee import -p pool1 -l lib1 -t D00279L5 1FA682L5 -r
- ltfsee import -p pool1 -l lib1 -t D00279L5 1FA682L5 -R -o -P /importUser
- ltfsee import -p pool1 -l lib1 -t D00279L5 1FA682L5 --offline

## The ltfsee info command

To list current information about IBM Spectrum Archive EE jobs and scans (groups of jobs), and resource inventory for tapes, drives, nodes, pools, and files of the whole IBM Spectrum Archive EE system or its part related to one tape library, run the **ltfsee info** command. You do not need to have root user authority to run the **ltfsee info** command, as shown in the following examples:

```
ltfsee info <resource> [-l <library_name>]
ltfsee info files -f <filepath_regular_expression> [-q]
ltfsee info jobs [-l <library_name>]
ltfsee info scans [-l <library_name>]
```

Where **<resource>** must be one of the following parameters:

- ▶ **libraries** (without the -l option)
- ▶ **tapes**
- ▶ **pools**
- ▶ **nodes**
- ▶ **nodegroups**
- ▶ **drives**

The following are additional options:

- ▶ **-l**  
Show the information for the specified library. If -l is omitted, show the information for the whole system.
- ▶ **-q**  
Do not check to see whether the file exists on the tape.

This option lists current IBM Spectrum Archive EE resource inventory information or information about ongoing migration and recall jobs or scans. A job refers to the migration or recall of a single file. A scan refers to a list of such jobs (usually originating from an IBM Spectrum Scale policy scan) that was submitted by using the **migrate** option, as shown in the following examples:

```
ltfsee info jobs
ltfsee info scans
ltfsee info tapes
ltfsee info drives
ltfsee info nodes
ltfsee info pools
ltfsee info files -f <filepath_regular_expression> [-q]
```

For more information, see “Sample ltfsee info command output” on page 307 to see sample command output for the following examples.

## Sample ltfsee info command output

This section presents sample output for each option of the **ltfsee info** command.

### Sample ltfsee info jobs command output

Example 11-2 shows sample output of the **ltfsee info jobs** command.

Example 11-2 Sample output of the ltfsee info jobs command

```
[root@ltfseesrv ~]# ltfsee info jobs
Job Type Status Idle(sec) Scan ID Tape Pool Library Node File Name or inode
Migration In-progress 0 852297473 IM1178L6 warning lto_ts4500 1 /ibm/gpfs/user1/LTFS_EE_FILE_Fa43W zQKcAeP2_j2XT9t.bin
Migration In-progress 0 852297473 IM1178L6 warning lto_ts4500 1 /ibm/gpfs/user1/LTFS_EE_FILE_JmkqS
oqPX9FhRoD1gTdCnDttA3Ee0A24pfQfbeva2ruRvTp_mXYbKW.bin
Migration In-progress 0 852297473 IM1178L6 warning lto_ts4500 1 /ibm/gpfs/user1/LTFS_EE_FILE_1Y08u
c068bAzXBgPOTSKWHhU21Sv_LFFDAd.bin
Migration In-progress 0 852297473 IM1178L6 warning lto_ts4500 1 /ibm/gpfs/user1/LTFS_EE_FILE_MXJ4c h_BQdX25.bin
Migration In-progress 0 852297473 IM1178L6 warning lto_ts4500 1 /ibm/gpfs/user1/LTFS_EE_FILE_KDVMc
Yoz3iQk2IUnvPcY2G7pqydvFBLWTxEQ7G31Dz_f3pKGG.bin
Migration In-progress 0 852297473 IM1178L6 warning lto_ts4500 1 /ibm/gpfs/user1/LTFS_EE_FILE_9xy6w
Oj3YQvd6Vu_055v3t.bin
```

Table 11-1 describes the status codes that can be displayed by the IBM Spectrum Archive EE **ltfsee info jobs** command.

Table 11-1 Status codes for ltfsee info jobs

Status code	Description
Unscheduled	The Unscheduled status indicates that the job request is received but is not yet scheduled for work. For example, when a migration request to a single tape storage pool is received, the migration job is queued to be scheduled for migration and is assigned a job status of Unscheduled.
Scheduled	The Scheduled status indicates that the job request is scheduled for work. For example, when an unscheduled migration job gets scheduled for migration, the job status changes to Scheduled.
In progress	The In progress status indicates that initial processing for the job is in progress. For example, when data copying from GPFS to tape begins for a job that is scheduled for migration, the job status changes to In progress.
Copied	The data copying from GPFS to tape for the migration job is completed.
Data transferred	The Data transferred status indicates that data transfer from tape to disk was successful.

### Sample ltfsee info tapes command output

Example 11-3 shows sample output of the **ltfsee info tapes** command.

Example 11-3 Sample output of the ltfsee info tapes command

```
[[root@ltfs17 ~]# ltfsee info tapes
Tape ID Status Type Capacity(GiB) Used(GiB) Free(GiB) Reclaimable(GiB) Pool Library Address Drive Appendable
TST200L6 Valid L6 2242 0 2242 0 test1 TS3310 4117 - yes
LTFS01L7 Valid L7 5338 0 5338 0 test TS3310 4119 - yes
PFE001L6 Valid L6 2242 95 2146 0 test1 TS3310 4118 - yes
```

Table 11-2 describes the status codes for the **ltfsee info tapes** command.

Table 11-2 Status codes for the ltfsee info tapes command

Status code	Description
Valid	The Valid status indicates that the cartridge is valid.
Exported	The Exported status indicates that the cartridge is valid and is exported.
Offline	Offline status indicates that the cartridge is valid and is exported offline.

Status code	Description
Unknown	The Unknown status indicates that the cartridge contents are unknown.
Unformatted	The Unformatted status indicates that the cartridge is not formatted. To use this tape with IBM Spectrum Archive EE, format the tape by using the <b>ltfsee pool</b> command with the <b>format</b> option.
Unavailable	The Unavailable status indicates that the cartridge is not available in the IBM Spectrum Archive EE system. Tapes that are newly inserted into the tape library have an Unavailable status. To recall files from a tape with this status, first import the tape by using the <b>ltfsee import</b> command. To add a tape with this status to the IBM Spectrum Archive EE system, complete the following steps: 1. Move the tape to a home slot by using the <b>ltfsee tape move</b> command. 2. Add the tape to a tape storage pool by using the <b>ltfsee pool</b> command.
Critical	The Critical status indicates that an attempt to write to this cartridge failed. To avoid data loss, recover the data on this cartridge by using the <b>ltfsee recover</b> command, keep the cartridge, and contact IBM Spectrum Archive support to determine the root cause of the critical state.
Write Fenced	The Write Fenced status indicates that an attempt to write to this cartridge failed, but was able to sync the files that were successfully migrated up to the write failure. To avoid data loss, recover the data on this cartridge by using the <b>ltfsee recover</b> command, keep the cartridge, and contact IBM Spectrum Archive support to determine the root cause of the write fenced state.
Warning	The Warning status indicates that a read failure occurred on this cartridge. To avoid data loss, move the data from this cartridge onto another cartridge by using the <code>relocate_replica.sh</code> script, keep the cartridge and contact IBM Spectrum Archive support to determine the root cause of the warning state.
Error	Indicates that the tape cartridge reported a medium error. The tape cartridge status returns to Valid by physically removing the medium from the library, then adding it to the library again. If this state occurs again, the tape cartridge should not be used. Contact IBM Spectrum Archive support to determine the root cause of the error state.
Write Protected	The Write Protected status indicates that the cartridge is write protected.
Duplicated	The Duplicated status indicates that there is more than one cartridge with this bar code. Remove the cartridge from the tape library.
Not supported	The Not supported status indicates that the cartridge is not supported by IBM Spectrum Archive EE. Only cartridges that are supported by IBM Spectrum Archive EE can be used.
Invalid	The Invalid status indicates that the cartridge is inconsistent with the LTFs format. To check and repair this tape before you add it to a tape storage pool, use the <b>ltfsee pool add</b> command with the <b>check</b> option.
Inaccessible	The Inaccessible status indicates that the tape is not allowed to move in the library. The tape might be stuck in the library or in one of the tape drives. For example, in a multiple node cluster, the tape is mounted by a tape drive that is assigned to a node that is not available.
Disconnected	The Disconnected status indicates that the EE and LE components that are used by this tape cannot communicate. The admin channel connection might be disconnected.
Unusable	The Unusable status indicates that the tape cannot be used.

### Sample `ltfsee info drives` command output

Example 11-4 shows sample output of the `ltfsee info drives` command.

Example 11-4 Sample output of the `ltfsee info drives` command

```
[root@ltfs17 ~]# ltfsee info drives
Drive S/N   Status      Type      Role  Library  Address  Node ID  Tape      Node Group
00000XXXXX Not mounted LT06      mrg   lib2     257     3        -         G1
F09CA15000 Not mounted LT05      mrg   lib2     258     3        -         G1
1068000560 Stock       UNKNOWN   ---   lib2     256     -        -         -
00078B00F3 Mounted     LT06      mrg   lib1     259     1        TY0305L6  G0
00013B0084 Mounted     LT06      mrg   lib1     260     1        TY0288L6  G0
00013B0037 Mounted     LT06      mrg   lib1     261     2        TY0303L6  G0
00078D00C2 Mounted     LT07      mrg   lib1     262     2        TY0287L6  G0
00078D00BE Mounted     LT07      mrg   lib1     263     2        TY0201L6  G0
1013000125 Mounted     LT06      mrg   lib1     264     1        TY0125L5  G0
```

Table 11-3 describes the status codes for the `ltfsee info drives` command.

Table 11-3 Status codes for `ltfsee info drives`

Status code	Description
Not mounted	The Not mounted status indicates that the tape drive is empty and there is no mount or unmount in progress for that tape drive.
Mounted	The Mounted status indicates that the tape drive mounted a cartridge. No operation is being done on that cartridge.
Mounting	The Mounting status indicates that a mount is in progress for that tape drive.
Unmounting	The Unmounting status indicates that an unmount action is in progress for that tape drive.
In use	The In use status indicates that an operation is being done on that tape drive. The <code>ltfsee info jobs</code> command should show a job in a corresponding state.
Stock	The Stock status indicates that the tape drive is removed from the library inventory or the drive information is not accessible.
Locked	The Locked status indicates that the tape drive is preserved for maintenance because the cartridge on this tape drive is in a critical condition.
Disconnected	The Disconnected status indicates that the EE and LE components that are used by this drive cannot communicate. The admin channel connection might be disconnected.

### Sample `ltfsee info nodes` command output

Example 11-5 shows a sample output of the `ltfsee info nodes` command.

Example 11-5 Sample output of the `ltfsee info nodes` command

```
[root@ltfs17 ~]# ltfsee info nodes
Node ID     Status      Node IP           Drives  Ctrl Node  Library  Node Group  Host Name
3           Available  192.168.209.161   2        Yes  lib2    G1         ltfsee05.ltd.sdl
2           Available  192.168.208.18    3        No   lib1    G0         ltfs18
1           Available  192.168.208.17    3        Yes  lib1    G0         ltfs17
```

Table 11-4 describes the status codes that can be displayed by the `ltfsee info nodes` command.

Table 11-4 Status codes for the `ltfsee info nodes` command

Status code	Description
Available	The Available status indicates that the IBM Spectrum Archive LE+ component on this node is available for operation.
License Expired	The License Expired status indicates that the IBM Spectrum Archive LE+ component on this node has an expired license.
Unknown	The Unknown status indicates that the IBM Spectrum Archive LE+ component on this node is inoperable, or the node is down.
Disconnected	The Disconnected status indicates that the EE and LE component on this node cannot communicate. The admin channel connection might be disconnected.
Not configured	The Not configured status indicates that the work directory of LE is not correct. Reconfigure the node by stopping IBM Spectrum Archive EE and running the <code>ltfsee_config -m ADD_NODE</code> command on the node. Then, restart IBM Spectrum Archive EE.
Error	The Error status indicates that a critical component of EE is not functioning or has lost communication.

### Sample `ltfsee info pools` command output

Example 11-6 shows a sample output of the `ltfsee info pools` command.

Example 11-6 Sample output of the `ltfsee info pools` command

```
[root@ltfs17 ~]# ltfsee info pools
```

Pool Name	Total(TiB)	Free(TiB)	Unref(TiB)	Tapes	Type	Library	Node	Group
pool05	0.0	0.0	0.0	0	-	lib2	G1	
pool18	12.2	12.2	0.0	6	LTO	lib1	G0	
pool17	13.1	13.1	0.0	6	LTO	lib1	G0	

### Sample `ltfsee info scans` command output

Example 11-7 shows a sample output of the `ltfsee info scans` command.

Example 11-7 Sample output of the `ltfsee info scans` command

```
[root@ltfs17 ~]# ltfsee info scans
```

Scan ID	Parent	Queued	Succeed	Failed	Canceled	Library	Ctrl	Node
2391414785	-	10	0	0	0	lib1	ltfs17	
2391415297	2391414785	10	0	0	0	lib1	ltfs17	

### Sample ltfsee info files command output

Example 11-8 shows a sample output of the `ltfsee info files` command.

Example 11-8 Sample output of the `ltfsee info files` command

```
[root@ltfs17 MIG_20151116_170428]# ltfsee info files *
Name: file_1M_1.0
Tape id:TY0263L6@lib1:TY0305L6@lib1 Status: migrated
Name: file_1M_1.1
Tape id:TY0263L6@lib1:TY0125L5@lib1 Status: migrated
Name: file_1M_1.2
Tape id:TY0263L6@lib1:TY0125L5@lib1 Status: migrated
Name: file_1M_1.3
Tape id:TY0263L6@lib1:TY0125L5@lib1 Status: migrated
Name: file_1M_1.4
Tape id:TY0263L6@lib1:TY0287L6@lib1 Status: migrated
Name: file_1M_1.5
Tape id:TY0263L6@lib1:TY0125L5@lib1 Status: migrated
Name: file_1M_1.6
Tape id:TY0263L6@lib1:TY0201L6@lib1 Status: migrated
Name: file_1M_1.7
Tape id:TY0263L6@lib1:TY0125L5@lib1 Status: migrated
Name: file_1M_1.8
Tape id:TY0263L6@lib1:TY0288L6@lib1 Status: migrated
Name: file_1M_1.9
Tape id:TY0263L6@lib1:TY0303L6@lib1 Status: migrated
Name: file_1M_1.9
Tape id:TY0263L6@lib1:TY0125L5@lib1 Status: migrated
```

Table 11-5 describes the status codes that can be displayed by the `ltfsee info files` command.

Table 11-5 Status codes for the `ltfsee info files` command

Status code	Description
Resident	The Resident status indicates that the file is resident in the GPFS namespace and is not saved, migrated, or premigrated to a tape.
Migrated	The Migrated status indicates that the file was migrated. The file was copied from GPFS file system to a tape, and exists only as a stub file in the GPFS namespace.
Premigrated	The Premigrated status indicates that the file was premigrated. The file was copied to a tape (or tapes), but the file was not removed from the GPFS namespace.
Saved	The Saved status indicates that the file system object that has no data (a symbolic link, an empty directory, or an empty regular file) was saved. The file system object was copied from GPFS file system to a tape.
Offline	The Offline status indicates that the file was saved or migrated to a tape cartridge and thereafter the tape cartridge was exported offline.
Lost	The Lost status indicates that a file has the migrated status, but it is not accessible from IBM Spectrum Archive EE because the tape cartridge it is supposed to be on is not accessible. The file might be lost because of tape corruption or if the tape cartridge was removed from the system without exporting.
Strayed	The Strayed status indicates that the file was in the Saved or Premigrated status, but is not accessible from the IBM Spectrum Archive EE because the tape cartridge it is supposed to be on is not accessible. The file might be lost because of tape corruption or if the tape cartridge was removed from the system without exporting.

## The ltfsee migrate command

To migrate files to tape storage pools, run the **ltfsee migrate** command, as shown in the following example:

```
ltfsee migrate -s <gpfs_scan_result_file> -p <pool_name1[@library_name1]>
[<pool_name2[@library_name2]> [<pool_name3[@library_name3]>]]
```

The `gpfs_scan_result_file` file includes the list of files to migrate. Each line of this file must end with `-- <filename>`. All files are migrated to the specified target storage pool `<pool_name1>`.

Optionally, redundant copies can be created in up to two additional storage pools (`pool_name2`, `pool_name3`). The pools can be in the same or different tape libraries. In single-library systems, the library name can be omitted in the input parameters. In multiple library systems, it must be specified.

This command does not complete until all files have been migrated.

Here are two examples that use the **ltfsee migrate** command:

```
ltfsee migrate -s ./gpfsscan.txt -p mypool
ltfsee migrate -s ./gpfsscan.txt -p poolA@library1 poolB@library2
```

**Note:** If using policy-based migration, the syntax for the OPTS parameter value is `-p <pool1[@lib1]> [<pool2[@lib2]>]`. Here is an example:

```
RULE EXTERNAL POOL 'rule1' EXEC '/opt/ibm/ltfsee/bin/ltfsee' OPTS '-p pool1@lib1
pool2@lib2'
```

## ltfsee node command

Run the **ltfsee node** command to enable or disable nodes and to list the settings of each node:

```
ltfsee node show -n <nodeid> [-l <libraryname>] [-a <attribute>]
ltfsee node show -A
ltfsee node <up|down> -n <nodeid> [-l <libraryname>]
```

Run the **ltfsee node show** command to show the configuration attributes of a specified node or **ltfsee node show -A** to show the configuration attributes of all nodes in a table.

Run the **ltfsee node <up|down>** command to enable or disable a node. The enabled node can be used as part of the IBM Spectrum Archive EE system, whereas a disabled node does not participate in the system. The command has these attributes:

► **-n <nodeid>**

The ID of the node of which you want to show the properties.

► **-l <libraryname>**

The name of the library with which the node is associated. If only a single library is configured in the system, this option can be omitted.

► **-a <attribute>**

The attribute to show. If this option is omitted, all attributes of the selected node are shown. Valid node attributes are:

- **nodeid:** Unique identifier
- **host:** Host name
- **ip:** IP address

- **port**: Port number for LTFS
- **enabled**: If the node ID is up or down
- **nodegroup**: Associated node group

## ltfsee pool command

To create, delete, add, remove, show, and modify IBM Spectrum Archive EE storage pools, and to format tapes before they are added to storage pools, run the **ltfsee pool** command, as shown in the following example:

```
ltfsee pool <create|delete> -p <poolname> [-g <nodegroup>] -l <libraryname>
[--worm <worm>]
ltfsee pool <add|remove> -p <poolname> -t <tape_id_1> [ ... <tape_id_N>]
[-l <libraryname>] [-f | -F | -e | -c | -d | -r] [-T <format_type>]
ltfsee pool set -p <poolname> [-l <libraryname>] -a <attribute> -v <value>
ltfsee pool show -p <poolname> [-l <libraryname>] [-a attribute]
```

The **pool create** and **pool delete** options are used to create or delete the specified IBM Spectrum Archive EE tape storage pool:

► **-p <poolname>**

The name of the pool to create or delete.

► **-g <nodegroup>**

The node group that the pool is/will be associated with. For creating pools, this option can be omitted if only a single node group is configured for the library. When deleting a pool, the node group can always be omitted.

► **-l <libraryname>**

The name of the library that the pool is, or will be, associated with. If only a single library is configured for the system, this option can be omitted.

► **--worm <worm>**

Specifies whether the new pool is a WORM (Write-Once, Read-Many) pool, or not. The following are the valid options:

- **physical**: Physical WORM
- **no**: No WORM (default)

You must specify a value for <poolname>, as shown in the following example:

```
ltfsee pool create -p mypool --worm no
ltfsee pool create -p mypool -g mynodegroup -l mylibrary
ltfsee pool delete -p mypool -l mylibrary
```

The **pool add** and **pool remove** options are used to add or remove one or more tapes from the specified IBM Spectrum Archive EE storage pool. With these options, you must specify a value for <poolname> and one or more <tape\_ids>.

► **-p <poolname>**

The name of the pool to which the tapes are to be added or removed.

► **-t <tapeid>**

The ID of the tape to add or remove from the pool. Multiple tapes can be added or removed with a single command by specifying the **-t** option multiple times.

► **-l <libraryname>**

The name of the library that the pool is associated with. If only a single library is configured for the system, this option can be omitted.

▶ **-f, --format**

Format the specified tape before adding it to pool.

**Important:** Before you specify the `ltfsee pool add` command with the `-F` option, ensure any tapes that you want to add are empty. Otherwise, replicas might become lost.

▶ **-F, --force\_format**

Force the format of the tape. Required to format an “already formatted” tape.

▶ **-e, --exception\_format**

Exceptionally format the tape. Required to format a “containing migrated file data” tape. This option needs to be used to re-add a tape that has been previously removed by the `ltfsee pool remove` command with the `-E` option.

▶ **-c, --check**

Check and repair the specified tape before adding it to pool.

▶ **-d, --deep\_recovery**

Specify this option instead of `-c` to repair a tape that is missing an EOD mark, `-r`, `--force_remove`.

Try to remove the tape even if there are active files on the tape.

▶ **-E, --empty\_remove**

Remove the tape even if the tape has not been reconciled and still contains stale file data that has already been deleted on GPFS. Note that this option does not perform reconciliation. The removed tape must be formatted with the `-e` option when adding it back to a pool. For more information about the `-E` command, see “Active file check before formatting tape cartridges” on page 151.

Here are some examples of the `ltfsee pool add` command:

```
ltfsee pool add -p mypool -t 327AAQL5 -t 277BAXL5 -t 329AAQL5
ltfsee pool add -p mypool -t 327AAQL5 -f
ltfsee pool add -p mypool -t 327AAQL5 -c
```

The `pool set` command sets the configuration attributes of a pool:

▶ **-p <poolname>**

The name of the pool on which to set the properties.

▶ **-l <libraryname>**

The name of the library with which the pool is associated. If only a single library is configured in the system, this option can be omitted.

▶ **-a <attribute>**

The attribute to set.

► **-v <value>**

The value that you can assign to the attribute. The following attributes can be set, with their valid values:

- **poolname**: 16 characters, ASCII alpha-numeric and underscores.
- **mediarestriction**: none, L5, L6, L7, L8, M8, JB, JC, JD, JY, JZ, JK, or JL
- **format**: E07, E08, or 55F.
- **mountlimit**: Integer; 0 for no limit.
- **lowspacewarningenable**: Yes or no
- **lowspacewarningthreshold**: Integer in Tib, 0 for no threshold
- **nospacewarningenable**: Yes or no

To run this command, you must specify the pool that is to be configured with the appropriate <attribute> and <value> pairs, as shown in the following examples:

```
ltfsee pool set -p mypool -a poolname -v mysecondpool
ltfsee pool set -p mypool -a mountlimit -v 2
```

The **pool show** command shows the configuration attributes of a tape pool. To run this command, you must specify the pool for which the configurations are to be displayed. You can specify one of the following optional parameters with **ltfsee pool show**:

► **-p <poolname>**

The name of the pool to show the properties for.

► **-l <libraryname>**

The name of the library with which the pool is associated. If only a single library is configured in the system, this option can be omitted.

► **-a <attribute>**

The attribute to show. If this option is omitted, all attributes of the pool are shown. The following are the valid pool attributes:

- **poolname**: Human-readable name.
- **poolid**: Unique identifier.
- **devtype**: Device type.
- **mediarestriction**: media type to be used for write
- **format**: Preferred tape format.
- **worm**: Write-once, read-only.
- **nodegroup**: Associated node group.
- **fillpolicy**: Tape fill policy.
- **owner**: Owner.
- **mountlimit**: Maximum number of drives that can be used for migration (0: unlimited).
- **lowspacewarningenable**: SNMP traps are sent when the space of the pool is low
- **lowspacewarningthreshold**: SNMP notification threshold value for free pool size (0: no threshold)
- **nospacewarningenable**: whether SNMP traps are sent when the pool has no space for migration
- **mode**: Operational mode

Here are some examples of the **ltfsee pool show** command:

```
ltfsee pool show -p mypool
ltfsee pool show -p mypool -l mylibrary -a devtype
```

### The **ltfsee premigrate** command

To premigrate files to tape storage pools, run the **ltfsee premigrate** command, as shown in the following example:

```
ltfsee premigrate -s <gpfs_scan_result_file> -p <pool_name1[@library_name1]>
[<pool_name2[@library_name2]> [<pool_name3[@library_name3]>]]
```

The `gpfs_scan_result_file` file includes the list of files to premigrate. Each line of this file must end with `-- <filename>`. All files are premigrated to the specified target storage pool `<pool_name1>`.

Optionally, redundant copies can be created in up to two additional storage pools (`pool_name2`, `pool_name3`). The pools can be in the same or in up to two different tape libraries. In single-library systems, the library name can be omitted in the input parameters. In multiple library systems, it must be specified.

This command does not complete until all files have been premigrated.

Here are two examples of using the **ltfsee premigrate** command:

```
ltfsee premigrate -s ./gpfs_scan.txt -p mypool
ltfsee premigrate -s ./gpfs_scan.txt -p poolA@library1 poolB@library2
```

**Note:** If you are using policy-based premigration, the syntax for the OPTS parameter value is `-p <pool1[@lib1]> [<pool2[@lib2]>]`. Here is an example:

```
RULE EXTERNAL POOL 'rule1' EXEC '/opt/ibm/ltfsee/bin/ltfsee' OPTS '-p
pool1@lib1 pool2@lib2'
```

### The **ltfsee rebuild** command

To rebuild an IBM Spectrum Scale file system into the specified directory with the files and file system objects that are found on the specified tapes, run the **ltfsee rebuild** command:

```
ltfsee rebuild -P <pathName> -p <poolName> -l <libraryName> <-t tapeId1 ... tapeIdN>
```

The specified tape must be in a pool and must not be offline. Internally, the **ltfsee rebuild** command is run by running the **ltfsee import** command.

If multiple versions or generations of a file are found, the latest version or generation is selected. If any of the versions or generations cannot be determined, the file that is most recently imported is renamed and two (or more) versions or generations are rebuilt or recovered.

**Note:** The **ltfsee rebuild** command should be rarely used as it would be in a disaster recovery scenario when the IBM Spectrum Scale file system has been lost.

## The ltfsee recall command

To perform selective bulk recalls on a list of files, run the **ltfsee recall** command:

- ▶ **ltfsee recall [-l library\_name] -s <gpfs\_scan\_result\_file>**

The `gpfs_scan_result_file` file includes the list of files to recall. Each line of this file must end with `-- <filename>`, where there is a space before and after the double dash. In single-library systems, the library name can be omitted in the input parameters. In multiple library systems, it must be specified.

- ▶ **ltfsee recall [-l library\_name] -f <recall\_list\_file>**

The `recall_list_file` file includes the list of files to recall. Each line contains a file name with an absolute path or a relative path based on the working directory. In single-library systems, the library name can be omitted in the input parameters. In multiple library systems, it must be specified. It also is possible to pass the output of another command.

Here are examples of the **ltfsee recall** command using the `-f <recall_list_file>` parameter:

```
ltfsee recall ./recall_list.txt
find . -type f |ltfsee recall
ltfsee recall -l library2 -f ./recall_list.txt
find . -type f |ltfsee recall -l library 2
```

## The ltfsee reclaim command

To free tape space by removing files and content that are not referenced from tapes in a storage pool, run the **ltfsee reclaim** command. The **ltfsee reconcile** command is a prerequisite for efficient tape reclamation, as shown in the following example:

```
ltfsee reclaim -p <poolname> [-l <libraryname>] [-n <number> | -t <tape_id_1
tape_id_2 ... tape_id_N>
[-L[file_number_th]]] [-R[remaining_capacity_th]] [-G[space_gain_th]] [-q]
```

Use this command to start reclamation of one or more tapes within the specified storage pool. Optionally, it can define which tapes to reclaim and remove from the storage pool or how many tapes to reclaim and remove from the storage pool. The set of tapes to be reclaimed can be further refined by setting one or both remaining-capacity and space gain thresholds. Also, the number of files to be reclaimed can be defined, but only when one tape is defined to be reclaimed.

The storage pool to be reclaimed must be specified as the first parameter. If other parameters are not specified, all the tapes from the storage pool are reclaimed and the tapes remain in the storage pool.

The command also includes the following optional parameters:

- ▶ **-p <poolname>**

The name of the tape storage pool.

- ▶ **-l <libraryname>**

The name of the tape library to which the tape pool and tapes to be reclaimed belong. This option can be omitted in single-library systems.

- ▶ **-n <number>**

Reclaim and remove `<number>` tapes from the storage pool. Up to `<number>` tapes from the storage pool are reclaimed and removed from the storage pool. If less than `<number>` tapes are successfully reclaimed, only the reclaimed tapes are removed.

The `-n` and `-t` options are mutually exclusive, but they can be combined with `-R`.

► **-t <tape\_ids>**

Reclaim and remove the listed tapes from the storage pool. The listed tapes from the storage pool are reclaimed and removed from the storage pool. If not all listed tapes are successfully reclaimed, only the reclaimed tapes are removed.

The **-t** and **-n** options are mutually exclusive, but they can be combined with **-R**. All the listed tapes must be members of the storage pool.

► **-L [threshold]**

Defines how many files to reclaim with this command.

It works with the **-t** option when only one **tape\_id** is specified. When this option is used, the tape might not be removed from the storage pool, the storage pool might not gain capacity, while the unreferenced capacity of the tape increases. The tape needs to be reclaimed again to complete the reclaim operation. This option can be used to limit the amount of processing time that the reclaim operation takes.

If this option is specified without a threshold value, the default value is 100,000.

► **-R [threshold]**

Defines the remaining capacity threshold to qualify for reclamation.

The remaining tape capacity threshold defines a percentage of the total tape capacity. If a tape has more free space than this threshold, it will not be reclaimed.

If this option is specified without a threshold value, the default value 10 (10%) is used.

► **-G [threshold]**

Defines the expected gain threshold to qualify for reclamation.

The expected gain threshold defines a percentage of the total tape capacity. If a tape has less expected gain space than this threshold, it will not be reclaimed.

If this option is specified without a threshold value, the default value 75 (75%) is used.

► **-q**

Quick reconcile is performed before reclaim.

All DMAPI enabled GPFS file systems must be mounted to use this option. The **-q** option solves inconsistency of removed files between GPFS and LTFS, but other inconsistencies remain on the files.

To solve inconsistencies that quick reconcile does not handle, a **reconcile** command needs to be run on those files.

Here are some examples of the **ltfsee reclaim** command:

```
ltfsee reclaim -p pool1 -l lib1 -R -G60
ltfsee reclaim -p pool1 -l lib1 -n 2
ltfsee reclaim -p pool1 -l lib1 -t D00279L5 1FA682L5
ltfsee reclaim -p pool1 -l lib1 -t TY0164JC -L10000
```

**Note:** For **-L**, **-R**, and **-G**, when the argument value is provided, it must not be space separated from option. For example, see in the above examples **-G60** and **-L10000**.

## The ltfsee reconcile command

To perform reconciliation tasks between the IBM Spectrum Scale namespace and the LTFS namespace, run the **ltfsee reconcile** command. This command is a prerequisite for efficient tape reclamation, as shown in the following example:

```
reconcile [ -t <tape_id_1 tape_id_2 ..> ] -p <poolName> -l <libraryName> [-P] [-u]
[-w <wait_time>] [-g <gpfs_fs_1 gpfs_fs_2 ..>]
```

Use this command to start reconciliation of all or selected tapes or pools against all or selected GPFS file systems. This command includes the following parameters:

- ▶ **-t <tape\_ids>**  
Reconcile specified tapes.  
Can be combined with **-p** and **-g**.
- ▶ **-p <poolName>**  
The name of the pool to reconcile. If **-t** is not specified, reconcile all tapes in the specified pool.  
Can be combined with **-t** and **-g**.
- ▶ **-l <libraryName>**  
The name of the library that the pool is or will be associated with. If only a single library is configured for the system, this option can be omitted.
- ▶ **-P**  
Partial reconcile. If not all the requested tapes can be reserved for reconcile, reconcile the tapes that can be reserved.
- ▶ **-u**  
Skip reconcile pre-check so that tapes get mounted for reconcile regardless of need.
- ▶ **-w <wait\_time>**  
Maximum time (in seconds) that the reconcile process can spend trying to reserve the requested tapes. Default value is 300 seconds.
- ▶ **-g <gpfs\_fss>**  
Reconcile tapes (defined by other options) against the specified GPFS file systems. Can be combined with any other reconcile option. If not specified, tapes are reconciled against all GPFS file systems.

## The ltfsee recover command

To recover files from a tape or to remove a tape from IBM Spectrum Archive EE when the tape is in a critical state or a write fenced state, run the **ltfsee recover** command. It follows GNU's getopt behavior if two of **-s**, **-c** and **-r** are specified at a same time.

Here is the syntax of the command:

```
ltfsee recover -t <tape_id> [-p <pool_name>] [-l <library_name>] <-s | -r | -c >
```

The command has these attributes:

- ▶ **-t <tape\_id>**  
Specifies the tape ID to recover.
- ▶ **-p <pool\_name>**  
The pool name must be provided.

- ▶ **-l <library\_name>**  
The library name must be provided on multi-library systems, but can be omitted on single-library systems.
- ▶ **-f <list\_file>**  
Specify output file that has a file list to be recovered and a file of non-EE files. This command uses /tmp/[hostname].[pid].[first-mountpoint-GPFS].recoverlist by default.
- ▶ **-s**  
Lists the files that have corresponding migrated stubs or saved objects in GPFS.  
This option is mutually exclusive with **-c** and **-r**.
- ▶ **-c**  
Recover files that have corresponding migrated stubs or saved objects in GPFS.  
This option is mutually exclusive with **-s** and **-r**.
- ▶ **-r**  
Remove the tape from the LTFS inventory if no recoverable file exists on the tape.  
This option is mutually exclusive with **-s** and **-c**.

Here are some examples of running the **ltfsee recover** command:

```
ltfsee recover -t D00279L5 -p pool1 -l library2 -s -f /home/foo/recover_list.txt
ltfsee recover -t D00279L5 -p pool1 -l library2 -c
ltfsee recover -t D00279L5 -p pool1 -l library2 -r
```

**Note:** Older syntax is also supported on single tape library systems.

Here is an example of the older syntax for the **ltfsee recover** command:

```
ltfsee recover <tape_id> <-s | -S | -r | -c | -C [-o <directory>] >
```

## The ltfsee repair command

To repair a file or object in the strayed state by changing it to the resident state when the tape (or tapes) used for premigration, migration, or save are not available, run the **ltfsee repair** command. This command does not check tape availability. This option removes metadata on GPFS that is used for keeping the file/object state. For example:

```
ltfsee repair file123
```

## The ltfsee retrieve command

To trigger the IBM Spectrum Archive EE system to retrieve information about physical resources from the tape library, run the **ltfsee retrieve** command. The IBM Spectrum Archive EE system periodically does this operation automatically, but it can be useful to trigger this operation if the configuration has changed, such as when a drive is added or removed. The command has this option:

- ▶ **-l <library\_name>**  
Retrieve the information for specified library. If the **-l** option is omitted, retrieve the information from all active libraries, for example.

Here are some examples of the **ltfsee retrieve** command:

```
ltfsee retrieve
ltfsee retrieve -l library1
```

## The ltfsee save command

To save file system objects to tape storage pools, run the **ltfsee save** command, as shown in the following example:

```
save -s <gpfs_scan_result_file> -p <target_storage_pool_name>
[<redundant_copy_pool_1> [<redundant_copy_pool_2>]]
```

The `gpfs_scan_result_file` file includes the list of file system objects (empty files, empty directories, and symbolic links) to save. Each line of this file must end with `-- <filename>`. All file system objects are saved to the specified target storage pool. Optionally, redundant copies can be created in up to two extra storage pools. If more than one library is configured, then the pool must be specified as `poolname@libraryname`.

This command does not complete until all file system objects have been saved.

Here are some examples of the **ltfsee save** command:

```
ltfsee save -s ./gpfsscan.txt -p poolA
ltfsee save -s ./gpfsscan.txt -p poolA@library1 poolB@library2
ltfsee save -s ./gpfsscan.txt -p mypool@mylib
```

## The ltfsee start command

To start the IBM Spectrum Archive EE system, run the **ltfsee start** command. The IBM Spectrum Archive LE+ component and HSM component must be running before you can use this command. You can run the **ltfsee start** command on any IBM Spectrum Archive EE node in the cluster.

**Note:** Before you run the **ltfsee start** command, other IBM Spectrum Archive EE components must already be running. For more information about how to run and obtain the statuses for these components, see 7.4.1, “Starting IBM Spectrum Archive EE” on page 143.

**Important:** If the **ltfsee start** command does not return after several minutes, it might be because the firewall is running or unmounting tapes from drives. The firewall service must be disabled on the IBM Spectrum Archive EE nodes. For more information, see 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 60.

If the IBM Spectrum Archive EE system is not already running on a node in the cluster, the **ltfsee start** command starts the IBM Spectrum Archive EE system on the local node, as shown in the following example:

```
ltfsee start [-l <library_name>]
```

This example starts the IBM Spectrum Archive EE system or its part that is related to one library. The `-l` option allows you to define the name of the library for which IBM Spectrum Archive EE will be started. If the `-l` option is omitted, IBM Spectrum Archive EE is started for all the configured libraries, as shown in the following example:

```
ltfsee start -l LIB1
ltfsee start
```

**Note:** On a single library system, the old command syntax can also be used:

```
ltfsee start <gpfs file system>
```

**Important:** If there is more than one managed file system, select one of the file systems to use and always use the same one. Otherwise, the `ltfsee.config` file will be inconsistent.

## The `ltfsee status` command

To identify the node where the multi-tape management module (MMM) service started, run the `ltfsee status` command.

The `status` option returns the IP address and corresponding process ID of the node where the MMM service was started, as shown in the following example:

```
ltfsee status
```

## The `ltfsee stop` command

To stop the IBM Spectrum Archive EE system, run the `ltfsee stop` command. You can run the `ltfsee stop` command on any IBM Spectrum Archive EE node in the cluster:

```
ltfsee stop
```

This option stops the IBM Spectrum Archive EE system or its part that is related to one library. The `-l` option allows you to define the name of the library for which IBM Spectrum Archive EE will be stopped. If the `-l` option is omitted, IBM Spectrum Archive EE is stopped for all the configured libraries, as shown in the following example:

```
ltfsee stop [-l <library_name>]
```

## The `ltfsee tape` command

To move one or more tapes to an import/export (IE) slot in the I/O station, or to a home slot, run the `ltfsee tape move` command, as shown in the following example:

```
ltfsee tape move homeslot | ieslot -t <tape_id_1 tape_id_2 ... tape_id_N> [-p <pool_name>] [-l <library_name>]
```

This command moves the specified tape between its home slot and an IE slot in the I/O station. If the tape belongs to a storage pool and is online (not in the offline state), a request to move it to an IE slot fails. A tape in an IE slot cannot be added to a storage pool. A tape must be moved to its home slot before it can be added.

Here are the options for the `ltfsee tape move` command:

- ▶ **-t <tape\_id>**  
Specifies the tape ID to move.
- ▶ **-p <pool\_name>**  
A pool name must be provided if the target tapes are assigned to a storage pool. The pool name must be omitted for unassigned tapes.
- ▶ **-l <library\_name>**  
The library name, which must be provided on multi-library systems, but can be omitted on single-library systems.

The following examples show the tape move `homeslot | ieslot` command being used:

```
ltfsee tape move homeslot -t D00279L5 1FA682L5 -p pool1 -l library1
ltfsee tape move homeslot -t D00279L5 1FA682L5 -l library1
ltfsee tape move ieslot -t D00279L5 1FA682L5 -p pool1 -l library2
```

**Note:** The older syntax is also supported on single-tape library systems:

```
ltfsee tape move homeslot | ieslot <tape_id_1 tape_id_2 ... tape_id_N>
```

To test the real status of the specified tape that is assigned to a pool but whose status is UNKNOWN, run the **ltfsee tape validate** command as shown in the following example:

```
ltfsee tape validate -p <pool_name> [-l <library_name>] -t <tape_id>
```

Here are the options for the **tape validate** command:

- ▶ **-p <pool\_name>**  
The pool name, which must be provided.
- ▶ **-l <library\_name>**  
The library name, which must be provided on multi-library systems, but can be omitted on single-library systems.
- ▶ **-t <tape\_id>**  
Specifies the tape ID to validate. The tape must be in the UNKNOWN state, and must be assigned to a pool.

An example of the tape validate command is shown below:

```
ltfsee tape validate -l library1 -p pool1 -t D00279L5
```

To show the configuration attributes of a tape, run the **ltfsee tape show** command as shown in the following example:

```
ltfsee tape show -t <tape_id> [-l <libraryname> ] [-a <attribute>]
```

Here are the options of the command:

- ▶ **-t <tape\_id>**  
The names (IDs) of the tape of which to show the properties.
- ▶ **-l <libraryname>**  
The name of the library with which the tape is associated. If only a single library is configured in the system, this option can be omitted.
- ▶ **-a <attribute>**  
The attribute to show. If this option is omitted, all attributes of the tape are shown. These are valid tape attributes:
  - **tapeid**: Name (ID) of the tape.
  - **poolname**: Name of the pool to which the tape is assigned.
  - **offline**: Offline message of the tape that was offline-exported.

Here are some examples of the command:

```
ltfsee tape show -t 327AAQL5  
ltfsee tape show -t 327AAQL5 -l mylibrary -a poolname
```

## The ltfsee threshold command

To specify the limit at which migrations are preferred over recalls, run the **ltfsee threshold** command, as shown in the following example:

```
ltfsee threshold <percentage>
```

The **threshold** parameter determines the limit of the file system (by percentage) at which migrations are preferred over recalls. The default value is 95%.

If this value is passed for one of the managed file systems, recalls are preferred again after the file system usage drops by 5%. For example, if a threshold of 93% is selected, recalls are preferred again when the file system usage is at or below 88%.

If no threshold value is specified, the current value is shown, for example:

```
lrfsee threshold 22
```

## 11.1.2 IBM Spectrum Scale commands

Use these commands to manage GPFS file systems that you use with your IBM Spectrum Archive EE system.

### The `mmapplypolicy` command

To manage migration and replication of the data to and from IBM Spectrum Scale storage pools, run the GPFS `mmapplypolicy` command. It can also be used to delete files from IBM Spectrum Scale.

The node on which the command is run must have the GPFS mounted. The node must be able to run remote shell commands on any other node in the IBM Spectrum Scale cluster without the use of a password and without producing any extraneous messages. For more information, see “Requirements for administering a GPFS file system” in *GPFS: Administration and Programming Reference*, SA23-2221. GPFS documentation is available at IBM Knowledge Center.

For IBM Spectrum Scale:

[https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale\\_welcome.html](https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale_welcome.html)

For GPFS:

[https://www.ibm.com/support/knowledgecenter/SSFKCN/gpfs\\_content.html](https://www.ibm.com/support/knowledgecenter/SSFKCN/gpfs_content.html)

For more information about the `mmapplypolicy` command and about GPFS or IBM Spectrum Scale, see your GPFS or IBM Spectrum Scale documentation.

### The `mmgetstate` command

To display the state of the IBM Spectrum Scale daemon on one or more nodes, run the GPFS `mmgetstate` command.

The node on which the command is run must have the GPFS mounted. The node must be able to run remote shell commands on any other node in the IBM Spectrum Scale cluster without the use of a password and without producing any extraneous messages. For more information, see “Requirements for administering a GPFS file system” in *GPFS: Administration and Programming Reference*, SA23-2221, and the websites that are referenced in “The `mmapplypolicy` command” on page 324.

The GPFS `mmgetstate -a` command displays the nodes where GPFS is active.

## 11.1.3 Tivoli Storage Manager for Space Management commands

There is a subset of Tivoli Storage Manager for Space Management commands that you can use to configure and administer space management for the file systems that are used with your IBM Spectrum Archive EE system.

IBM Spectrum Archive EE provides a limited description of the commands in this subset. For more information about these commands, see the “HSM client command reference” topic and other related topics in “Tivoli Storage Manager for Space Management client options, commands, and reference” at IBM Knowledge Center for IBM Spectrum Protect family:

[http://www.ibm.com/support/knowledgecenter/SSSR2R\\_7.1.3/hsmul/r\\_hsmul\\_ref.html](http://www.ibm.com/support/knowledgecenter/SSSR2R_7.1.3/hsmul/r_hsmul_ref.html)

**Important:** The IBM Spectrum Archive EE license does not entitle customers to use any other Tivoli Storage Manager components or products other than Tivoli Storage Manager for Space Management from the IBM Spectrum Protect family to migrate data to LTFS.

## Compatibility of Tivoli Storage Manager for Space Management commands with IBM Spectrum Archive EE

Only a subset of Tivoli Storage Manager for Space Management commands is compatible with the IBM Spectrum Archive EE environment. Use only compatible Tivoli Storage Manager for Space Management commands. Otherwise, your system might not work correctly and error messages such as the following message might be returned:

```
ANS2172E Command not supported in HSMBACKENDMODE TSMFREE
```

The following Tivoli Storage Manager for Space Management commands are compatible with the IBM Spectrum Archive EE environment:

- ▶ **dsmmigfs**
- ▶ **dsm1s**
- ▶ **dsmreca11**

## The dsmmigfs start, stop, and enablefailover functions

To start or stop HSM daemons, run the **dsmmigfs** command with the start or stop parameter.

**Important:** The daemons are started with the same environment as the **dsmwatchd** watch daemon. Therefore, the **dsm.opt** and **dsm.sys** options files in the Tivoli Storage Manager for Space Management default installation path `/usr/tivoli/tsm/client/ba/bin` are used.

Before you use this command, complete the following steps to verify your IBM Spectrum Scale settings:

1. Verify that the IBM Spectrum Scale is active.
2. Verify that the node that you want to run the command from belongs to the IBM Spectrum Archive EE cluster, and has the file system mounted.

**Important:** To display the nodes where IBM Spectrum Scale is active, run the IBM Spectrum Scale **mmgetstate -a** command. For more information, see 7.2.1, “IBM Spectrum Scale” on page 134.

The **dsmmigfs** command includes the following parameters:

- ▶ **start**  
Starts all HSM daemons on the local client node, except for the watch daemon (**dsmwatchd**).
- ▶ **stop**  
Stops all space management daemons on the local client node, except for the watch daemon (**dsmwatchd**).

► **enablefailover**

Activates the node for failover operations within the IBM Spectrum Scale cluster.

**dsmrecall**

To selectively recall migrated files to the local file system, run the **dsmrecall** command. Space management must be active.

The duration of the recall process depends on the size of the files that are being recalled. It takes time to mount and spool the LTFS tape cartridge, and the data transfer time for a large file can be considerable.

This command includes the `dsmrecall gpfs_path` syntax, as shown in the following example:

```
dsmrecall /ibm/ltfs/filename1.txt
```

## 11.2 Formats for IBM Spectrum Scale to IBM Spectrum Archive EE migration

This section describes the data and metadata formats that IBM Spectrum Archive EE uses for migrating files from an IBM Spectrum Scale / GPFS environment to an IBM Spectrum Archive (LTFS EE) system.

IBM Spectrum Archive EE uses the following data and metadata formats for migrating files from GPFS to LTFS:

► **GPFS data**

Figure 11-1 shows an example of original GPFS data.

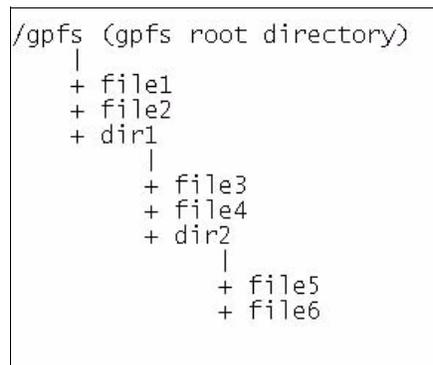


Figure 11-1 Example GPFS data directory structure

► Migrated data and symbolic links in LTFS

Figure 11-2 shows the LTFS data format of tapeA and tapeB after the following migrations of the GPFS files from Figure 11-1 on page 326:

- file1, file2, and file3 to tapeA.
- file1 (which is a replica), file4, file5, and file6 to tapeB.

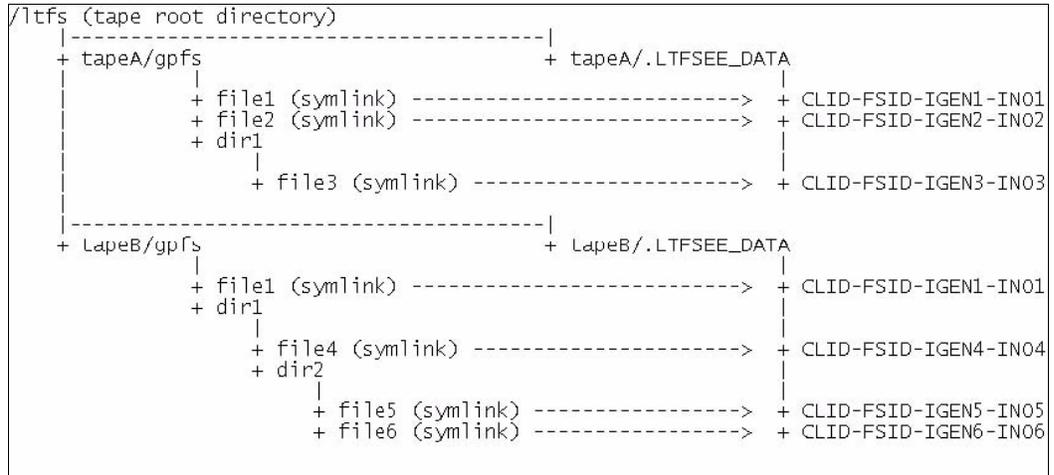


Figure 11-2 Example layout of migrated data and symbolic links on LTFS (on tapes)

► Migrated data

Migrated data is saved as files under the `.LTFSEE_DATA` directory on each tape. The `LTFSEE_DATA` directory is placed directly under the tape root directory on each tape. These data files are stored under unique ID (UID) based file names. A UID-based file name consists of the cluster ID (CLID), file system ID (FSID), inode generation number (IGEN), and inode number (INO). In this example, all of the files have the same CLID and the same FSID because all of the files belong to the same GPFS file system.

► Symbolic links

The GPFS directory structure is rebuilt on each tape under the tape root directory for all the migrated files that the tape contains. For each data file, a symbolic link is created under the original GPFS file name and location, which points to the corresponding data file on the tape. When a symbolic link is created, a relative path to the target is used so that if the IBM Spectrum Archive mount point changes, the link stays correct. For the `file1` example in Figure 11-2, the following symbolic link that corresponds to `file1` is created on tapeA:

```
/ltfs/tapeA/gpfs/file1 → ../.LTFSEE_DATA/CLID-FSID-IGEN1-IN01
```

► GPFS path as metadata in IBM Spectrum Archive

For each data file in IBM Spectrum Archive, an extended attribute is set that contains the original GPFS file path. In the `file1` example in Figure 11-2, the two following LTFS files have the extended attribute `gpfs.path` set to the value (file1 GPFS path) `/gpfs/file1`:

- `/ltfs/tapeA/.LTFSEE_DATA/CLID-FSID-IGEN1-IN01`
- `/ltfs/tapeB/.LTFSEE_DATA/CLID-FSID-IGEN1-IN01`

The saved GPFS path gives you the ability to re-create the original IBM Spectrum Scale namespace by using the reconciliation and export processes followed by the import process. In case of disaster, the approximate IBM Spectrum Scale namespace can be recovered from tapes because without reconciliation, not all file deletions and renames in IBM Spectrum Scale are reflected on the migration target tapes. The saved path is also used for removing stale symbolic links when a file is recalled into resident state.

## 11.3 System calls and IBM tools

This section describes downloading the IBM Tape Diagnostic Tool (ITDT) and using the IBM Linear Tape File System Format Verifier (LFV).

### 11.3.1 Downloading the IBM Tape Diagnostic Tool

The ITDT is an independent tool that provides diagnostic tests on tape drives and libraries. This section describes how to download ITDT and access the related documentation.

#### Before you begin

IBM maintains the latest levels of the ITDT and related documentation on Fix Central. Information about using the ITDT is available in *IBM Tape Device Drivers Installation and User's Guide*, S7002972, which is available on the same website.

#### About this task

To access the Fix Central portal and download the most recent version of the ITDT, complete the following steps:

1. Open the following URL in your web browser:  
<http://www.ibm.com/support/fixcentral>
2. Click **Product Group** → **System Storage**.
3. Click **Product Family** → **Tape systems**.
4. Click **Product Type** → **Tape drivers and software**.
5. Click **Product** → **IBM Tape Diagnostic Tool ITDT**.
6. Select your operating system from the **Platform** menu.
7. Click **Continue**.
8. (Optional) Narrow the search of available downloads according to your criteria.
9. Click **Continue** to view the list of available downloads.
10. Select the version that you want to download.
11. To download the new version, follow the instructions on the Fix Central download page.

## 11.3.2 Using the IBM LTFS Format Verifier

This section describes how to download, install, and run the IBM LFV utility command (**lfv**) to verify media hardware and data compatibility.

### Before you begin

Before installing the LTFS LFV, download the most recent version from the Fix Central website.

To download the most recent version of the LTFS LFV, complete the following steps:

1. Open the following URL in your web browser:

<http://www.ibm.com/support/fixcentral>

2. Click **Product Group** → **System Storage**.
3. Click **Product Family** → **Tape systems**.
4. Click **Product Type** → **Tape drivers and software**.
5. Click **Product** → **Linear Tape File System (LTFS) Format Verifier**.
6. Select your operating system from the **Platform** menu.
7. Click **Continue**.
8. Narrow the search of available downloads according to your criteria (this step can be skipped).
9. Click **Continue** to view the list of available downloads.
10. Select the version that you want to download.
11. To download the new version, follow the instructions on the Fix Central download page.

**Latest levels:** IBM maintains the latest levels of the LTFS LFV and information about using the tool and related documentation on the Fix Central website:

<http://www.ibm.com/support/fixcentral>

### About this task

To install the LTFS LFV, complete the following steps:

1. Download `lfvinst_<version>linuxx86_64` from this website:

<http://www.ibm.com/support/fixcentral>

2. To make `lfvinst_<version><OS><arch>` an executable file, run the following command:

```
chmod 700 lfvinst_<version><OS><arch>
```

3. To complete the installation, run the following command:

```
fvinst_<version><OS><arch>
```

### Verifying media compatibility by using the IBM LTFS Format Verifier

This section describes how to verify media hardware and data compatibility by using the LTFS LFV utility command. This section also describes the options that can be used with this command.

**Important:** LTFS LFV is not shipped with IBM Spectrum Archive EE, but is available as a separate download.

To verify that media are compatible with LTFs, run `1fv` from the command line. Enter one of the following commands:

- ▶ For Linux systems where the IBM Tape Device Driver is installed, `<target device>` should be `/dev/IBMtapeX`, where `X` is the index of the tape device to use, as shown in the following example:

```
./1fv -f /dev/IBMtape1
```

- ▶ For Linux systems where no IBM Tape Device driver is installed, `<target device>` should be `/dev/sgX`, where `X` is the index for the tape device to use, as shown in the following example:

```
./1fv -f /dev/sg0
```

**Important:** The index for the target tape device in the previous examples is shown as 0. If you are unsure which index value to use, run the `./1fv -s` command to scan for all attached tape devices.

The following `1fv` command options are available:

- ▶ **-f <target device>**

The target tape device on which verification is performed.

- ▶ **-h**

Displays help information.

- ▶ **-l**

Specifies the log file name. The default name is `1fv.log`.

- ▶ **-ll [Errors|Warnings|Information|Debug]**

Specifies the log level and the level of logging created. **Errors** is the default value.

- ▶ **-lp**

Specifies the log output directory. The default directory is `./output`.

- ▶ **-s**

Scans the system for tape devices and prints results to the window. This option provides a list of the available devices and can help you identify which drive to use. This option provides the following information:

- Sequential number.
- Driver handle/device file name.
- Drive product name.
- Drive firmware revision.
- Drive serial number (S/N).
- Host (H), bus (B), Target ID (T), and LUN (L) physical address of the drive.

For example, information that is provided by this list appears as shown in the following example:

```
#0 /dev/IBMtape0 -[ULT3580-TD4]-[85V1] S/N:1300000388 H2-B0-T0-L0
#1 /dev/IBMtape1 -[ULT3580-HH5]-[A2SG] S/N:1068000051 H2-B0-T1-L0
```

- ▶ **-v**

Enables verbose verification information.

▶ **-V --version**

Displays the program version.

▶ **-x**

Specifies that the extended verification is performed. The extended verification analyzes the entire tape cartridge and can take up to three hours to complete. Quick verification is the default.

## 11.4 IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products

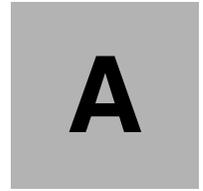
IBM Spectrum Archive EE cannot run concurrently with IBM Spectrum Archive LE. If IBM Spectrum Archive LE is already installed, it must be uninstalled before IBM Spectrum Archive EE is installed. For more information about the uninstallation procedure, see the “Uninstalling LTFS from a Linux system” topic in the IBM Spectrum Archive LE section of IBM Knowledge Center, which is available at this website:

<https://www.ibm.com/support/knowledgecenter/en/STZMZN>

In addition to uninstalling the IBM Spectrum Archive LE package, it is necessary to uninstall the IBM Spectrum Archive LE license module. To uninstall the license, you must run the following command after the other uninstallation commands that are presented in the IBM Spectrum Archive LE IBM Knowledge Center:

```
# rpm -e ltfs-license-2.1.0-[revision]
```





## Additional material

This Redbooks publication refers to additional material that can be downloaded from the Internet as described in the following sections.

### Locating the web material

The web material that is associated with this book is available in softcopy on the Internet from the IBM Redbooks web server. Point your web browser at:

<ftp://www.redbooks.ibm.com/redbooks/SG248333>

Alternatively, you can go to the IBM Redbooks website at:

[ibm.com/redbooks](http://ibm.com/redbooks)

Search for SG248333, select the title, and then click **Additional materials** to open the directory that corresponds with the IBM Redbooks form number.

### Using the web material

The additional web material that accompanies this book includes the following files:

<i>File name</i>	<i>Description</i>
<b>disclaimer.txt</b>	Disclaimer information
<b>relocate_replica.sh</b>	Installation tool

### Downloading and extracting the web material

Create a subdirectory (folder) on your workstation, and extract the contents of the web material `relocate_replica.sh` file into this folder.



# Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this book.

## IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only:

- ▶ *Active Archive Implementation Guide with IBM Spectrum Scale Object and IBM Spectrum Archive*, REDP-5237
- ▶ *IBM Linear Tape File System Installation and Configuration*, SG24-8090
- ▶ *IBM Tape Library Guide for Open Systems*, SG24-5946
- ▶ *IBM TS4500 R4 Tape Library Guide*, SG24-8235

You can search for, view, download, or order these documents and other Redbooks, Redpapers, Web Docs, draft and other materials, at this website:

<http://www.redbooks.ibm.com/>

## Other publications

The following publications are also relevant as further information sources:

- ▶ *IBM Tape Device Drivers Installation and User's Guide*, GC27-2130
- ▶ *IBM Tivoli Storage Manager for Space Management for UNIX and Linux Version 6.4 and Version 7.1 User's Guide*, SC23-9794
- ▶ *IBM Tivoli Storage Manager for UNIX and Linux Backup-Archive Clients Version 6.4 and Version 7.1 Installation and User's Guide*, SC23-9791

## Online resources

The following websites are also relevant as further information sources:

- ▶ IBM Client Demonstration Center  
<https://www.ibm.com/systems/clientcenterdemonstrations>

**Note:** The IBM Client Demonstration Center (for Business Partners, IBMers, and anyone with an IBMid) provides a catalog of remote demonstrations (video or live connection) which consist of self contained material for customer demonstrations of IBM solutions. Most of the demonstrations are provided with predefined scenarios and some also allow for the development of new scenarios. Demonstrations can also be considered as “ready to use” material for enablement or training.

- ▶ IBM Product Support: TS3310:  
<http://www.ibm.com/support/docview.wss?uid=ssg1S7005240>
- ▶ IBM Knowledge Center: TS4300:  
<https://www.ibm.com/support/knowledgecenter/en/STAKKZ>
- ▶ IBM Knowledge Center: TS4500:  
<https://www.ibm.com/support/knowledgecenter/STQRQ9>
- ▶ IBM Spectrum Archive Enterprise Edition at IBM Knowledge Center:  
<http://www.ibm.com/support/knowledgecenter/ST9MBR>
- ▶ IBM Spectrum Archive Library Edition at IBM Knowledge Center:  
<http://www.ibm.com/support/knowledgecenter/STZMZN>
- ▶ IBM Spectrum Archive Support:  
<http://www.ibm.com/systems/storage/tape/lufs/index.html>
- ▶ IBM Spectrum Scale at IBM Knowledge Center:  
[https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale\\_welcome.html](https://www.ibm.com/support/knowledgecenter/en/STXKQY/ibmspectrumscale_welcome.html)
- ▶ IBM GPFS at IBM Knowledge Center:  
[https://www.ibm.com/support/knowledgecenter/SSFKCN/gpfs\\_welcome.html](https://www.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html)
- ▶ IBM System Storage interactive product guide:  
<https://www.ibm.com/common/ssi/cgi-bin/ssialias?htmlfid=TS000364USEN>
- ▶ Linear Tape-Open Program:  
<http://www.lto.org/technology>
- ▶ SNIA Linear Tape File System Format Specification:  
[http://snia.org/sites/default/files/LTFS\\_Format\\_2.2.0\\_Technical\\_Position.pdf](http://snia.org/sites/default/files/LTFS_Format_2.2.0_Technical_Position.pdf)

## Help from IBM

IBM Support and downloads:

<http://www.ibm.com/support>

IBM Global Services:

<http://www.ibm.com/services>

**Redbooks**

**IBM Spectrum Archive Enterprise Edition V1.2.6: Installation and Configuration Guide**

SG24-8333-05

ISBN 0783442844



(0.5" spine)

0.475" x 0.873"

250 x 459 pages







SG24-8333-05

ISBN 0783442844

Printed in U.S.A.

Get connected

