

IBM Spectrum Archive Enterprise Edition V1.3.2.2 Installation and Configuration Guide

Hiroyuki Miyoshi

Khanh Ngo

Arnold Byron Lua

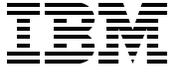
Yuka Sasaki

Yasuhiro Yoshihara

Larry Coyne



Storage



International Technical Support Organization

**IBM Spectrm Archive Enterprise Edition V1.3.2.2:
Installation and Configuration Guide**

March 2022

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

Tenth Edition (March 2022)

This edition applies to Version 1, Release 3, Modification 2, Fix Level 2 of IBM Spectrum Archive Enterprise Edition (product number 5639-LP1).

© Copyright International Business Machines Corporation 2015, 2020, 2022. 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

Notices	ix
Trademarks	x
Preface	xi
Authors	xi
Now you can become a published author, too!	xiii
Comments welcome	xiv
Stay connected to IBM Redbooks	xiv
Summary of changes	xv
December 2021, Tenth Edition Version 1.3.2.2	xv
June 2021, Ninth Edition Version 1.3.1.2	xv
February 2020, Eighth Edition Version 1.3.0.6	xvi
April 2019, Seventh Edition Version 1.3.0	xvii
June 2018, Sixth Edition Version 1.2.6	xvii
January 2018, Fifth Edition Version 1.2.5.1	xviii
August 2017, Fourth Edition Version 1.2.4	xviii
February 2017, Third Edition Version 1.2.2 minor update	xix
January 2017, Third Edition Version 1.2.2	xix
August 2016, Second Edition Version 1.2.1	xx
June 2016, First Edition, Version 1.2	xx
January 2015, Third Edition Version 1.1.1.2	xxi
November 2014, Second Edition Version 1.1.1.1	xxii
Chapter 1. IBM Spectrum Archive Enterprise Edition	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
1.2.1 User Task Reporting	7
1.3 IBM Spectrum Archive EE components	9
1.3.1 IBM Spectrum Archive EE terms	9
1.3.2 Hierarchical Storage Manager	13
1.3.3 Multi-Tape Management Module	13
1.3.4 IBM Spectrum Archive Library Edition component	14
1.4 IBM Spectrum Archive EE cluster configuration introduction	15
Chapter 2. IBM Spectrum Archive overview	17
2.1 Introduction to IBM Spectrum Archive and LTFS	18
2.1.1 Tape media capacity with IBM Spectrum Archive	19
2.1.2 Comparison of the IBM Spectrum Archive products	21
2.1.3 IBM Spectrum Archive Single Drive Edition	22
2.1.4 IBM Spectrum Archive Library Edition	23
2.1.5 IBM Spectrum Archive Enterprise Edition	24
2.2 IBM Spectrum Scale	27
2.2.1 Overview	27
2.2.2 Storage pools	28
2.2.3 Policies and policy rules	29
2.2.4 Migration or premigration	30

2.2.5 Active File Management	30
2.2.6 Scale Out Backup and Restore	32
2.3 OpenStack SwiftHLM	33
2.4 IBM Spectrum Archive EE dashboard	34
2.5 IBM Spectrum Archive EE REST API	36
2.6 Types of archiving	36
Chapter 3. Planning for IBM Spectrum Archive Enterprise Edition	39
3.1 IBM Spectrum Archive EE deployment options	40
3.1.1 On IBM Spectrum Scale Servers	40
3.1.2 As an IBM Spectrum Scale Client	43
3.1.3 As an IBM Elastic Storage Systems IBM Spectrum Scale Client	43
3.1.4 As an IBM Spectrum Scale stretched cluster	44
3.2 Data-access methods	46
3.2.1 Data access using application or users on IBM Spectrum Scale Clients	46
3.2.2 Data access using Protocol Nodes on IBM Spectrum Scale Server Nodes with IBM Spectrum Archive EE	47
3.2.3 Data access using Protocol Nodes integrated with the IBM ESS	49
3.3 System requirements	50
3.3.1 Limitations	51
3.4 Required software for Linux systems	54
3.4.1 Required software packages for Red Hat Enterprise Linux systems	54
3.4.2 Required software to support REST API service on RHEL systems	55
3.4.3 Required software to support a dashboard for IBM Spectrum Archive Enterprise Edition	55
3.4.4 Required software for SwiftHLM	55
3.5 Hardware and software setup	56
3.6 IBM Spectrum Archive deployment examples	57
3.6.1 Deploying on Lenovo servers	57
3.6.2 Deploying on Versastack converged infrastructure	60
3.7 Sizing and settings	61
3.7.1 Redundant copies	62
3.7.2 Planning for LTO-9 Media Initialization/Optimization	63
3.7.3 IBM Spectrum Archive EE Sizing Tool	63
3.7.4 Performance	65
3.7.5 Ports that are used by IBM Spectrum Archive EE	70
3.8 High-level component upgrade steps	70
Chapter 4. Installing IBM Spectrum Archive Enterprise Edition	73
4.1 Installing IBM Spectrum Archive EE on a Linux system	74
4.2 Installation prerequisites for IBM Spectrum Archive EE	74
4.2.1 Installing the host bus adapter and device driver	75
4.3 Installing IBM Spectrum Archive EE	75
4.3.1 Extracting binary rpm files from an installation package	76
4.3.2 Installing, upgrading, or uninstalling IBM Spectrum Archive EE	78
4.4 Installing a RESTful server	83
4.5 Quick installation guide for IBM Spectrum Archive EE	88
4.6 Library replacement	88
4.6.1 Library replacement procedure	89
4.6.2 Pool relocation procedure	91
4.7 Tips when upgrading host operating system	95
Chapter 5. Configuring IBM Spectrum Archive Enterprise Edition	97
5.1 Configuration prerequisites	98

5.1.1	Configuration worksheet tables	98
5.1.2	Obtaining configuration information	101
5.1.3	Configuring key-based login with OpenSSH.	103
5.1.4	Preparing the IBM Spectrum Scale file system for IBM Spectrum Archive EE.	104
5.2	Configuring IBM Spectrum Archive EE	105
5.2.1	The ltfsee_config utility	105
5.2.2	Configuring a single node cluster	108
5.2.3	Configuring a multiple-node cluster.	112
5.2.4	Configuring a multiple-node cluster with two tape libraries.	115
5.2.5	Modifying a multiple-node configuration for control node redundancy	116
5.3	First-time start of IBM Spectrum Archive EE.	118
5.3.1	Configuring IBM Spectrum Archive EE with IBM Spectrum Scale AFM	119
5.3.2	Configuring a Centralized Archive Repository solution.	120
5.3.3	Configuring an Asynchronous Archive Replication solution	123
Chapter 6. Managing daily operations of IBM Spectrum Archive Enterprise Edition		129
6.1	Overview	131
6.1.1	IBM Spectrum Archive EE command summaries	131
6.1.2	Using the command-line interface.	135
6.2	Status information	136
6.2.1	IBM Spectrum Scale	136
6.2.2	IBM Spectrum Archive Library Edition component	137
6.2.3	Hierarchical Space Management	138
6.2.4	IBM Spectrum Archive EE.	140
6.3	Upgrading components	142
6.3.1	IBM Spectrum Scale	142
6.3.2	IBM Spectrum Archive LE component	145
6.3.3	Hierarchical Storage Management	145
6.3.4	IBM Spectrum Archive EE.	145
6.4	Starting and stopping IBM Spectrum Archive EE	146
6.4.1	Starting IBM Spectrum Archive EE.	146
6.4.2	Stopping IBM Spectrum Archive EE	147
6.5	Task command summaries	147
6.5.1	eadm task list	147
6.5.2	eadm task show	148
6.6	IBM Spectrum Archive EE database backup	152
6.7	IBM Spectrum Archive EE automatic node failover	153
6.7.1	IBM Spectrum Archive EE monitoring daemon.	153
6.8	Tape library management	154
6.8.1	Adding tape cartridges	154
6.8.2	Moving tape cartridges	156
6.8.3	Formatting tape cartridges	158
6.8.4	Removing tape drives	160
6.8.5	Adding tape drives	161
6.9	Tape storage pool management	162
6.9.1	Creating tape cartridge pools	162
6.9.2	Deleting tape cartridge pools	163
6.10	Pool capacity monitoring	163
6.11	Migration	165
6.11.1	Managing file migration pools	168
6.11.2	Threshold-based migration	168
6.11.3	Manual migration	173
6.11.4	Replicas and redundant copies.	177

6.11.5	Data Migration	180
6.11.6	Migration hints and tips	182
6.12	Premigration	184
6.12.1	Premigration with the eeadm premigrate command	185
6.12.2	Premigration running the mmapplypolicy command	185
6.13	Preserving file system objects on tape	186
6.13.1	Saving file system objects with the eeadm save command	186
6.13.2	Saving file system objects with policies	187
6.14	Recall	190
6.14.1	Transparent recall	190
6.14.2	Selective recall using the eeadm recall command	191
6.14.3	Read Starts Recalls: Early trigger for recalling a migrated file	192
6.14.4	Recommend Access Order	193
6.15	Recalling files to their resident state	196
6.16	Reconciliation	197
6.17	Reclamation	200
6.17.1	Reclamation considerations	200
6.18	Checking and repairing tapes	201
6.19	Importing and exporting	202
6.19.1	Importing	202
6.19.2	Exporting tape cartridges	204
6.19.3	Offlining tape cartridges	206
6.20	Drive Role settings for task assignment control	207
6.21	Tape drive intermix support.	208
6.21.1	Objective for WORM tape support	209
6.21.2	Function overview for WORM tape support	210
6.21.3	The effects of file operations on immutable and appendOnly files	211
6.22	Obtaining the location of files and data	213
6.23	Obtaining system resources, and tasks information	214
6.24	Monitoring the system with SNMP	216
6.25	Configuring Net-SNMP	218
6.25.1	Starting and stopping the snmpd daemon	219
6.25.2	Example of an SNMP trap	219
6.26	IBM Spectrum Archive REST API	219
6.26.1	Pools endpoint	220
6.26.2	Tapes endpoint	223
6.26.3	Libraries endpoint	225
6.26.4	Nodegroups endpoint	226
6.26.5	Nodes endpoint	227
6.26.6	Drives endpoint	228
6.26.7	Task endpoint	229
6.27	File system migration	231
Chapter 7. Hints, tips, and preferred practices		233
7.1	Preventing migration of the .SPACEMAN and metadata directories.	235
7.2	Maximizing migration performance with redundant copies	235
7.3	Changing the SSH daemon settings	237
7.4	Setting mmapplypolicy options for increased performance.	237
7.5	Preferred inode size for IBM Spectrum Scale file systems	239
7.6	Determining the file states for all files within the GPFS file system.	239
7.7	Memory considerations on the GPFS file system for increased performance	242
7.8	Increasing the default maximum number of inodes in IBM Spectrum Scale	242
7.9	Configuring IBM Spectrum Scale settings for performance improvement.	243

7.10	Use cases for mmapplypolicy	244
7.10.1	Creating a traditional archive system policy	244
7.10.2	Creating active archive system policies	245
7.10.3	IBM Spectrum Archive EE migration policy with AFM.	246
7.11	Capturing a core file on Red Hat Enterprise Linux with the Automatic Bug Reporting Tool	247
7.12	Anti-virus considerations	248
7.13	Automatic email notification with rsyslog.	248
7.14	Overlapping IBM Spectrum Scale policy rules	248
7.15	Storage pool assignment	250
7.16	Tape cartridge removal	250
7.16.1	Reclaiming tape cartridges before you remove or export them	251
7.16.2	Exporting tape cartridges before physically removing them from the library.	251
7.17	Reusing LTFS formatted tape cartridges	251
7.17.1	Reformatting LTFS tape cartridges through eeadm commands	252
7.18	Reusing non-LTFS tape cartridges	253
7.19	Moving tape cartridges between pools	254
7.19.1	Avoiding changing assignments for tape cartridges that contain files.	254
7.19.2	Reclaiming a tape cartridge and changing its assignment	254
7.20	Offline tape cartridges	254
7.20.1	Do not modify the files of offline tape cartridges	254
7.20.2	Solving problems	254
7.21	Scheduling reconciliation and reclamation	255
7.22	License Expiration Handling	255
7.23	Disaster recovery	256
7.23.1	Tiers of disaster recovery	256
7.23.2	Preparing IBM Spectrum Archive EE for a tier 1 disaster recovery strategy (offsite vaulting)	258
7.23.3	IBM Spectrum Archive EE tier 1 DR procedure	259
7.24	IBM Spectrum Archive EE problem determination	261
7.24.1	Rsyslog log suppression by rate-limiting.	262
7.25	Collecting IBM Spectrum Archive EE logs for support	262
7.26	Backing up files within file systems that are managed by IBM Spectrum Archive EE	264
7.26.1	Considerations	265
7.26.2	Backing up a GPFS or IBM Spectrum Scale environment	265
7.27	IBM TS4500 Automated Media Verification with IBM Spectrum Archive EE	266
7.28	How to disable commands on IBM Spectrum Archive EE.	271
7.29	LTO 9 Media Optimization	272
Chapter 8. IBM Spectrum Archive Enterprise Edition use cases		275
8.1	Overview of use cases	276
8.1.1	Use case for archive	276
8.1.2	Use case for tiered and scalable storage	277
8.1.3	Use case data exchange	278
8.2	Media and Entertainment	279
8.3	Media and Entertainment	280
8.4	High-Performance Computing.	281
8.5	Healthcare	283
8.6	Genomics	284
8.7	Archive of research and scientific data for extended periods	285
8.8	University Scientific Data Archive	286
8.9	Oil and gas	287
8.10	S3 Object Interface	288

8.11 AFM use cases	290
8.11.1 Centralized archive repository.	290
8.11.2 Asynchronous archive replication	291
Chapter 9. Troubleshooting IBM Spectrum Archive Enterprise Edition	293
9.1 Overview	294
9.1.1 Quick health check	294
9.1.2 Common startup errors	294
9.2 Hardware.	297
9.2.1 Tape library	297
9.2.2 Tape drives	297
9.3 Recovering data from a write failure tape	301
9.4 Recovering data from a read failure tape	302
9.5 Software	303
9.5.1 Linux	303
9.5.2 IBM Spectrum Scale	306
9.5.3 IBM Spectrum Archive LE component	308
9.5.4 Hierarchical storage management	309
9.5.5 IBM Spectrum Archive EE logs	310
9.6 Recovering from system failures.	312
9.6.1 Power failure	312
9.6.2 Mechanical failure	312
9.6.3 Inventory failure	312
9.6.4 Abnormal termination	313
Chapter 10. Reference	315
10.1 Command-line reference.	316
10.1.1 IBM Spectrum Archive EE help guide for commands	316
10.1.2 Drive status and state codes.	317
10.1.3 Node status codes	319
10.1.4 Tape status codes.	320
10.1.5 IBM Spectrum Scale commands.	323
10.1.6 IBM Spectrum Protect for Space Management commands	323
10.2 Formats for IBM Spectrum Scale to IBM Spectrum Archive EE migration	325
10.3 System calls and IBM tools	327
10.3.1 Downloading the IBM Tape Diagnostic Tool.	327
10.3.2 Using the IBM LTFS Format Verifier.	328
10.4 IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products	330
Related publications	331
IBM Redbooks	331
Other publications	331
Online resources	331
Help from IBM	332

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.

IBM®	Interconnect®	Redbooks®
IBM Cloud®	Operating System/2®	Redbooks (logo)  ®
IBM Elastic Storage®	POWER®	Storwize®
IBM FlashSystem®	PowerVM®	Tivoli®
IBM Spectrum®	ProtectTIER®	

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.

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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.

Red Hat, are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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 Enterprise Edition (EE) Version 1.3.2.2 for the IBM TS4500, IBM TS3500, IBM TS4300, and IBM TS3310 tape libraries.

IBM Spectrum Archive Enterprise Edition enables the use of the LTFS for the policy management of tape as a storage tier in an IBM Spectrum Scale based environment. It also helps encourage the use of tape as a critical tier in the storage environment.

This edition of this publication is the tenth 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 9, 8, 7, 6, and 5 tape drives. and the IBM TS1160, TS1155, TS1150, and TS1140 tape drives.

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 customers, IBM Business Partners, IBM specialist sales representatives, and technical specialists.

Authors

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



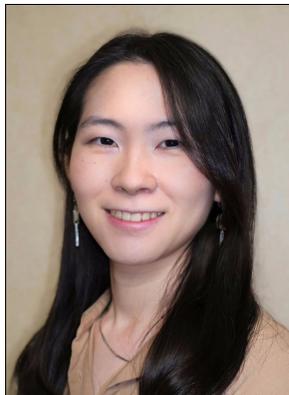
Hiroyuki Miyoshi is a manager of the tape development team in Tokyo. His team is engaged in developing tape drive hardware, linear tape file system (LTFS), and Hierarchical Storage Management (HSM) solution, which is IBM Spectrum Archive. He is a Master Inventor and joined IBM in 2001 with a master's degree in Electrical Science at Waseda University. Before becoming the manager in 2020, he was a developer for various storage products, such as RAID subsystems, BladeCenter storage/switch module, TS7700, SONAS, IBM Spectrum Scale, and IBM Spectrum Archive.



Khanh Ngo is an IBM Senior Technical Staff Member and Master Inventor in Tucson, Arizona. Khanh is in the Storage CTO Office specializing in data integration with IBM Storage products. 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. Because of his design and implementation work with many IBM Spectrum Archive EE customers across multiple industries worldwide, Khanh is often sought out for his expertise to lead, execute, and successfully complete proof of concepts and custom engineering solutions that integrate IBM Spectrum Archive EE into customers' production environments.



Arnold Byron Lua is the IBM Storage Architect and Systems Storage Product Line Manager in the ASEAN region. He has extensive experience in the Technology industry as an IBM Storage Technical Specialist, Systems Architect, and Systems Engineer in various companies over a period of 20 years. Aside from product line management, he designs fit-for-purpose architecture solutions that focus on business value while streamlining costs. Arnold is a Professional Electronics Engineer of the Philippines with a Bachelor of Science degree in Electronics and Communications Engineering.



Yuka Sasaki is a software developer for the IBM Spectrum Archive family of products: IBM Spectrum Archive Enterprise Edition, Library Edition, and Single Drive Edition in the IBM Tokyo Development Lab. Her team is engaged in developing and providing technical support of IBM Spectrum Archive. She joined IBM in 2019 with a Master's degree in Media and Governance degree from Keio University, where she received her Cyber Security and Cyber Informatics Certificate.



Yasuhiro Yoshihara is an Open Tape Subject Matter Expert who works in Yokohama since 2018 supporting IBM Spectrum Archive (LTFS), ProtecTIER®, and Open Tape Libraries/Tape Drives. He joined IBM in 1992 with a bachelor's degree in Commerce (Corporate Accounting). He has many years of experience in software support in charge of IBM Operating System/2® and then, host integration software, such as Personal Communications, Host On-Demand, Host Access Transformation Services, and Communications server. In addition to Host Integration Software, he started supporting Tape Software products, such as LTFS and ProtecTIER since 2012.



Larry Coyne is a Project Leader at the International Technical Support Organization, Tucson Arizona center. He has over 35 years of IBM experience with 23 years 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 customer relationship management, quality assurance, development management, and support management for IBM Storage Software.

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

Atsushi Abe
Hiroshi Itagaki
Takeshi Ishimoto
Osamu Matsumiya
Junta Watanabe
IBM Systems

Joseph Chuen Wei Liew
IBM Cloud® and Cognitive Software

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,
and Markus Schaefer

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

- ▶ 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

- ▶ 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 SG24-8333-09

for IBM Spectrum Archive Enterprise Edition V1.3.2.2: Installation and Configuration Guide
as created or updated on March 10, 2022.

December 2021, Tenth Edition Version 1.3.2.2

New and changed information Version 1.3.2.2

This IBM Redbooks publication includes new functions and updates for Version 1.3.2.0, 1.3.2.1, and 1.3.2.2.

Note: Review the readme and important notice files in [IBM Fix Central](#) for the latest IBM Spectrum Archive Enterprise Edition Version to make sure your upgrade plan will be current with important recommended updates.

These updates include the following support:

- ▶ File system migration support using SOBAR
- ▶ Recommended Access Order (RAO) function in IBM TS11xx and LTO 9 drives

High level modifications for publication include:

- ▶ The IBM tape drive TS1160 model 60S is now supported.

June 2021, Ninth Edition Version 1.3.1.2

New and changed information Version 1.3.1.2

High level updates to this Redbooks publication include:

Updated the following sections in Chapter 3., “Planning for IBM Spectrum Archive Enterprise Edition” on page 39.

- ▶ Added deployment options to “IBM Spectrum Archive EE deployment options” on page 40
- ▶ Updated “Data-access methods” on page 46

Reclaim tape drive usage has been updated in Chapter 6., “Managing daily operations of IBM Spectrum Archive Enterprise Edition” on page 129 to describe how IBM Spectrum Archive supports the parallel reclaim feature.

Added the following new use cases to Chapter 8., “IBM Spectrum Archive Enterprise Edition use cases” on page 275.

- ▶ 8.4, “High-Performance Computing” on page 281
- ▶ 8.10, “S3 Object Interface” on page 288

Updated Red Hat Enterprise Linux Servers and IBM Spectrum Scale levels.

New and changed information Version 1.3.1.0

The following updates were made to the list of supported software:

- ▶ The IBM tape drive TS1160 model 60S is now supported.
- ▶ Updated Red Hat Enterprise Linux Servers and IBM Spectrum Scale levels
- ▶ Updated the required software

New and changed information Version 1.3.0.7

- ▶ Added support for the TS1160 Model 60E tape drive.
- ▶ Added new response data fields to the REST API task information, in support of transparent recall tasks.

February 2020, Eighth Edition Version 1.3.0.6

New and changed information Version 1.3.0.6

- ▶ Supports the enable/disable control of certain tasks by type.
- ▶ Supports the assignment of TCP/UDP ports that IBM Spectrum Archive uses within a certain range.
- ▶ Supports FC/SAS port load balancing.
- ▶ Supports the syslog-ng facility.
- ▶ Supports a multi-node upgrade by using a single “ltfsee_install” command:
`ltfsee_install --upgrade --all`
- ▶ Supports the following new eeadm commands:
 - eeadm cluster set/show
 - eeadm drive up/down
 - eeadm drive set/show
- ▶ Enhanced the tape drive status indication.
- ▶ Adds the following new options to the existing eeadm commands:
 - eeadm MIGRATE --premigrate
 - eeadm LIST --migrate, --premigrate, --save, and --recall
- ▶ Changed the eeadm reclaim command behavior when only a pool is specified as the reclaim target.
- ▶ Renamed some of the eeadm reclaim command options.

New and changed information Version 1.3.0.4

- ▶ Performance improvements of eeadm migrate command and eeadm recall --resident command

New and changed information Version 1.3.0.3

- ▶ Improved the eeadm cluster stop command to complete more quickly
- ▶ IBM Spectrum Archive Enterprise Edition can be configured to use the IBM Spectrum Scale admin network, rather than the daemon network.
- ▶ IBM Spectrum Archive Enterprise Edition can be configured without SAN zoning in a multi-server configuration.

- ▶ Enhanced the eeadm migrate, eeadm premigrate and eeadm save commands to accept a list of files from stdin or from an input file.
- ▶ Updated Red Hat Enterprise Linux Servers and IBM Spectrum Scale levels
- ▶ Updated the required software
- ▶ Added the description of ports used by IBM Spectrum Archive
- ▶ Added the description of disabling/enabling the transparent recalls
- ▶ Updated recommended inode size
- ▶ Added limitation of file path length
- ▶ Updated REST API definition
- ▶ Renamed LE+ component to LE component
- ▶ All examples and references have been updated with the new CLI syntax changes and output. Added online help description for the eeadm CLI command.
- ▶ Removed chapter about upgrading from 1.1.x

April 2019, Seventh Edition Version 1.3.0

New information Version 1.3.0

- ▶ User Task Control and Reporting: Usability enhancements with new command-line interface (CLI) with more support for monitoring the progress and results of user operations and for tape maintenance
 - Active/Completed task listing including detailed information and output of command
 - Task results including file state transition results
 - Ability to run the command in background, with `--async` option
- ▶ Supports the Storage Networking Industry Association's LTFS format specification 2.4
- ▶ Expanded storage capacity with the TS1160 tape drive
- ▶ Supports the IBM Spectrum Scale backup function (mmbackup) for the same file system managed by IBM Spectrum Archive
- ▶ Bundles the open source package for external monitoring of IBM Spectrum Archive through a GUI/dashboard
- ▶ Usage of `/dev/sgX` device (`lin_tape` device driver is no longer required to be installed)

Changed Information

- ▶ Updated Red Hat Enterprise Linux Servers and IBM Spectrum Scale levels
- ▶ All examples and references have been updated with the new CLI syntax changes and output

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 Red Hat Enterprise Linux Servers 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.
- ▶ Updated Red Hat Enterprise Linux Servers and IBM Spectrum Scale levels

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

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 3.7.2, “Planning for LTO-9 Media Initialization/Optimization” on page 63.

Note: For more information about migration performance, see *IBM Spectrum Archive Enterprise Edition v1.2.2 Performance* [here](#).

- ▶ Updated the following examples: Example 6-60 on page 169, Example 6-62 on page 170, Example 6-63 on page 171, Example 6-64 on page 175, Example 6-66 on page 177, Example 7-7 on page 244, Example 7-8 on page 245, Example 7-9 on page 246, Example 7-19 on page 259.

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 6.16, “Reconciliation” on page 197
- ▶ Updated the `ltfsee recover` command in 10.1, “Command-line reference” on page 316
- ▶ 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
- ▶ Updated Red Hat Enterprise Linux Servers
- ▶ 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 78

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 added 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 more 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
- ▶ More operating system 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):

- ▶ 1.1, “Introduction” on page 2
- ▶ 1.2, “IBM Spectrum Archive EE functions” on page 6
- ▶ 1.3, “IBM Spectrum Archive EE components” on page 9
- ▶ 1.4, “IBM Spectrum Archive EE cluster configuration introduction” on page 15

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 more 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.

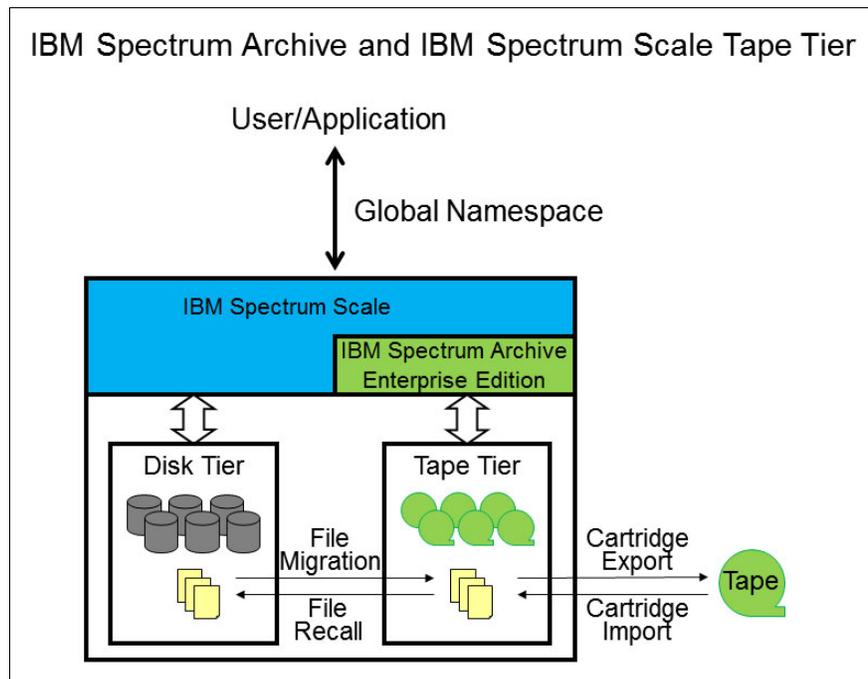


Figure 1-1 High-level overview of an IBM Spectrum Archive EE archive solution

IBM Spectrum Archive EE uses the IBM Spectrum Archive LE 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: To learn more about IBM Spectrum Archive EE and try it in a virtual tape library environment, see [IBM Spectrum Archive Enterprise Edition Fundamentals & Lab Access](#).

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.

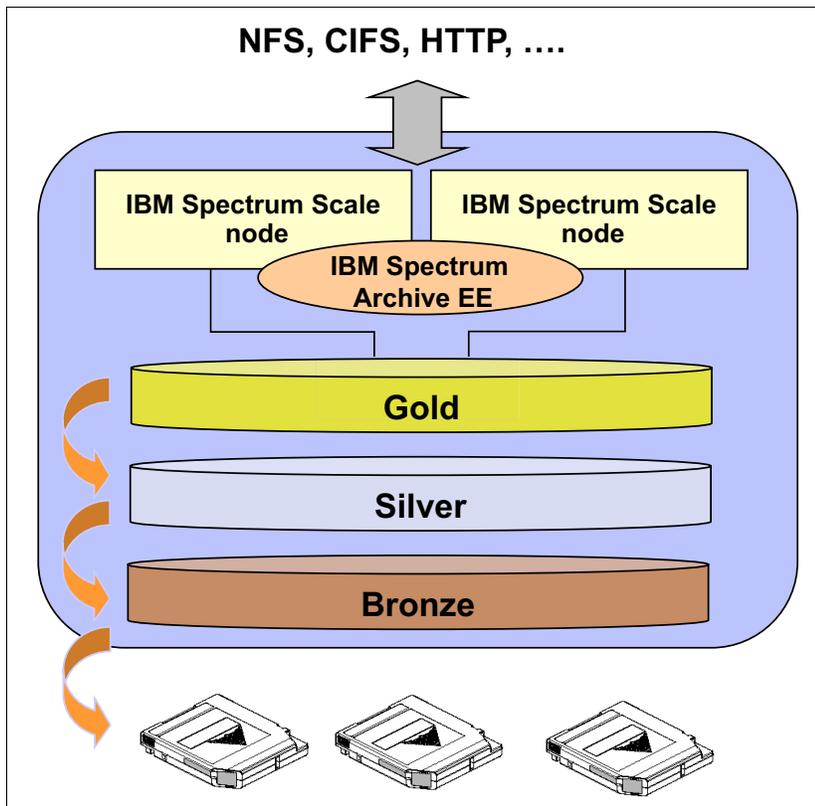


Figure 1-2 Tiered operational storage with IBM Spectrum Archive EE managing the tape tier

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 7.10.1, “Creating a traditional archive system policy” on page 244.

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 LTFS 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 LTFS 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.

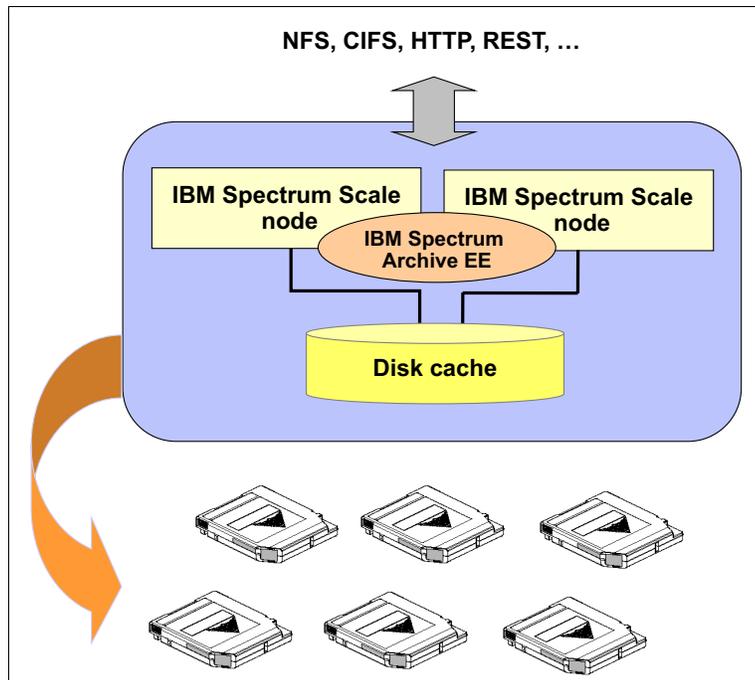


Figure 1-3 Archive storage management with IBM Spectrum Archive EE

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 7.10.2, “Creating active archive system policies” on page 245.

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.

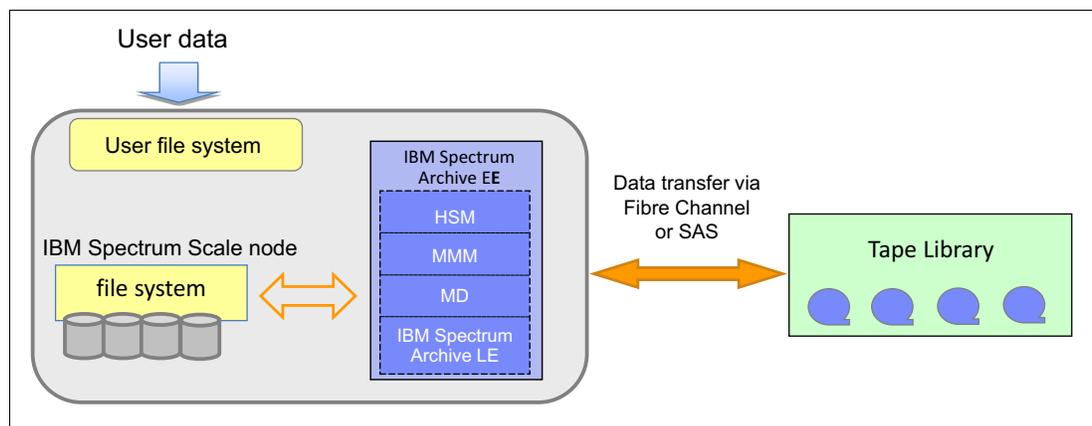


Figure 1-4 IBM Spectrum Archive EE integration with IBM Spectrum Scale and IBM Spectrum Archive LE

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, task status of your IBM Spectrum Archive EE solution.

1.2.1 User Task Reporting

Versions prior to IBM Spectrum Archive EE v1.3.0.0 reported any running or pending operations through the `ltfsee info scans` and `ltfsee info jobs` commands. However any successes or failures of prior commands, history of operations, or output of commands cannot be seen except through examination of the trace logs. Furthermore, the result of any file state changes from any migration, premigration, or recall commands cannot be viewed.

IBM Spectrum Archive EE v1.3.0.0 introduced the concept of user tasks. A user task is defined as:

- ▶ User initiated command or operation where each is considered a single task.
- ▶ Each task is uniquely identified by a task ID in the range of 1,000 to 99,999 with the next task ID after 99,999 starting back at 1,000.
- ▶ A maximum number of tasks that can be accepted is defined as follows:
 - 128 transparent recalls per each active control node
 - For tasks other than transparent recall, 512 tasks for each active control node
- ▶ If a task cannot be accepted due to this limit, the command is rejected immediately without assigning a task ID.
- ▶ Tasks can be manually cleared for any completed tasks.
- ▶ The last 10 tasks with errors will always be preserved unless manually cleared.

The following commands and operations create user tasks:

- ▶ `eeadm migrate`
- ▶ `eeadm premigrate`
- ▶ `eeadm recall`
- ▶ reading of a file (transparent recall)
- ▶ `eeadm save`
- ▶ `eeadm drive down`
- ▶ `eeadm drive unassign`
- ▶ `eeadm drive up`
- ▶ `eeadm tape assign`
- ▶ `eeadm tape datamigrate`
- ▶ `eeadm tape export`
- ▶ `eeadm tape import`
- ▶ `eeadm tape move`
- ▶ `eeadm tape offline`
- ▶ `eeadm tape online`
- ▶ `eeadm tape reclaim`
- ▶ `eeadm tape reconcile`
- ▶ `eeadm tape replace`
- ▶ `eeadm tape unassign`
- ▶ `eeadm tape validate`
- ▶ `eeadm task clearhistory`

The current user task features are:

- ▶ A task which is currently being processed or will be processed is defined as an “active” task. An “active” task status can be in one of the following:
 - interrupted: The task was running but currently not running due to other higher priority tasks such as recall tasks
 - running: The task is running using at least 1 tape cartridge in 1 tape drive resource
 - waiting: The task is created but is not running yet (pending to run)

- ▶ A task which has been processed is defined as a “completed” task. A “completed” task status can be in one of the following:
 - aborted: The task was running but a control node failover occurred or the cluster was stopped forcefully by 'eeadm cluster stop -f'. The task needs to be manually resubmitted
 - failed: The task completed with one or more errors
 - succeeded: The task completed with a success (no errors)
- ▶ List active tasks and completed tasks including the following information:
 - Task ID
 - Type
 - Priority
 - Status
 - Number of drives used for the task
 - Created time
 - Started time
- ▶ Show details of active tasks and completed tasks within the last 3 months including the following information:
 - Task ID
 - Type
 - Command and parameters
 - Status
 - Result
 - Accepted time
 - Started time
 - Completed time
 - In-use resources for pools, tape drives, and tape cartridges
 - Workload
 - Progress
- ▶ Show file results (success or failure) for any completed migration, premigration, or recall tasks
- ▶ Support an `--async` option which allows the task to be started asynchronously (i.e. when specified, the command returns immediately with the task ID after the request has been issued and the status of the command can be monitored later). It should be used on longer running commands where the administrator does not want to wait for it's completion. Using the `--async` option, the administrator can start the command, later query the status of the command using the `eeadm task list` command and finally review detailed information using the `eeadm task show` command. The following commands support this `--async` option:
 - eeadm drive down
 - eeadm drive unassign
 - eeadm drive up
 - eeadm tape export
 - eeadm tape import
 - eeadm tape assign
 - eeadm tape unassign
 - eeadm tape datamigrate
 - eeadm tape reclaim
 - eeadm tape replace
 - eeadm tape reconcile
 - eeadm tape validate
 - eeadm tape move

- eeadm task clearhistory
- eeadm task cancel

Note: The other commands (`eeadm migrate`, `eeadm premigrate`, `eeadm recall`, and `eeadm save`) does not fit well with the “`--async`” option when used through the `mmapplypolicy` command because it will cause `mmapplypolicy` command to continuously submit the same files repeatedly.

1.3 IBM Spectrum Archive EE components

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

- ▶ EE components: Multi-tape management module (MMM) and Monitoring Daemon (MD)
- ▶ Library Edition (LE) component
- ▶ Hierarchical Storage Management (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 10 and Figure 1-6 on page 11.

1.3.1 IBM Spectrum Archive EE terms

An IBM Spectrum Archive EE solution consists of the following components:

- ▶ IBM Spectrum Archive EE node

An x86_64 or ppc64le 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, SAS, or RoCE connection. One EE node cannot be connected to more than one logical library.

- ▶ IBM Spectrum Archive EE Cluster

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.

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.

- ▶ IBM Spectrum Scale Cluster

IBM Spectrum Scale servers (non-EE nodes) and EE nodes.

- ▶ IBM Spectrum Scale only Node

An IBM Spectrum Scale server that is running on a supported platform, such as Linux or Microsoft Windows, without IBM Spectrum Archive EE.

- ▶ Tape Pool

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 does not span across multiple tape libraries.

A tape pool is assigned to only one node group.

► 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.

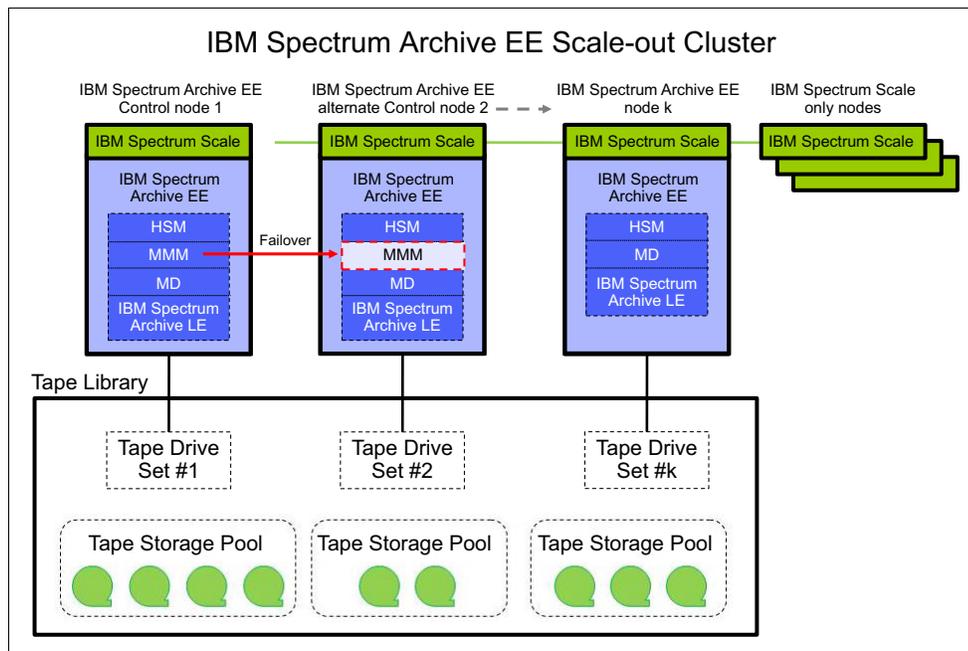


Figure 1-5 Components of IBM Spectrum Archive EE with separate IBM Spectrum Scale nodes

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

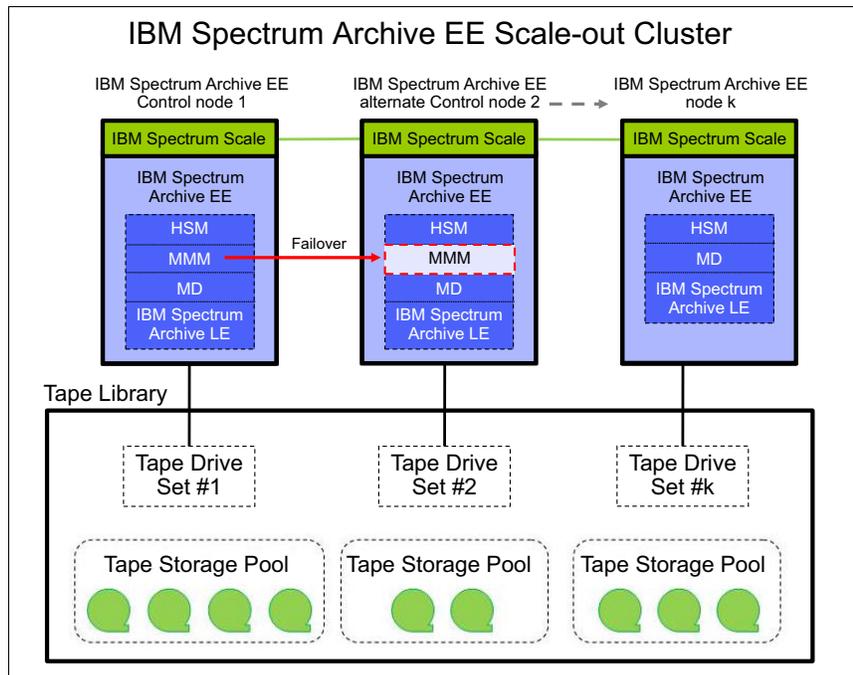


Figure 1-6 Components of IBM Spectrum Archive EE with no separate IBM Spectrum nodes

A second tape library can be added to the configuration, which expands the storage capacity and offers the opportunity to add 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).

For IBM Spectrum Scale V5.0.0 and later, distances up to 300 km (186.4 miles) are supported for stretched cluster with synchronous mirroring using block-level replication.

For more information, see [IBM Documentation](#).

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 [IBM Documentation](#).

Important: Stretched cluster is available for distances of 100 km (62 miles) - 300 km (186.4 miles). 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 is with two different IBM Spectrum Scale clusters with one instance of IBM Spectrum Archive at each site.

For more information about IBM Spectrum Scale AFM, see 2.2.5, “Active File Management” on page 30.

For more information, see [IBM Documentation](#).

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

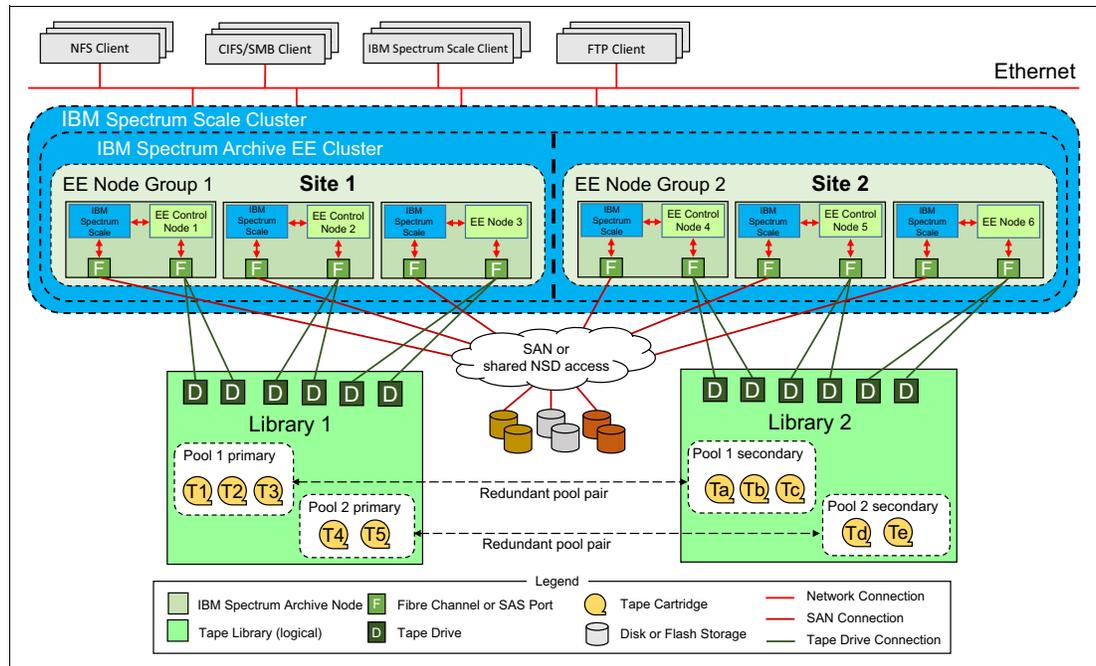


Figure 1-7 IBM Spectrum Archive EE

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 task requests, such as migration and recall requests, which are fulfilled by using available system nodes and tape resources. The following tasks are performed by the scheduler:

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

The control node also manages the creation of replicas across multiple tape libraries. For example, when the `eeadm 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 task queue for the primary tape library, then passes the secondary copy task to the control node for the second tape library. The second control node puts the copy task in the task queue for the second tape library.

When the scheduler component of the control node selects a tape cartridge and tape drive for a migration task, 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 tasks 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 task

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 task 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 **eeadm cluster start** command. The MMM service runs on only one IBM Spectrum Archive EE control node, for each library, 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 6.4, “Starting and stopping IBM Spectrum Archive EE” on page 146, and 6.24, “Monitoring the system with SNMP” on page 216.

Important: If the **eeadm cluster 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 78.

The **eeadm cluster 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 component

The IBM Spectrum Archive Library Edition (LE) component is the IBM Spectrum Archive tape tier of IBM Spectrum Archive EE. The LE is configured 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 v1.3.0.0 and later, IBM Spectrum Archive LE is started automatically when running the **eeadm cluster start** command. If errors occur during start of the IBM Spectrum Archive EE system, run the **eeadm node list** command to display which component failed to start.

For more information about the updated **eeadm node list** command, see 6.7, “IBM Spectrum Archive EE automatic node failover” on page 153.

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 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.

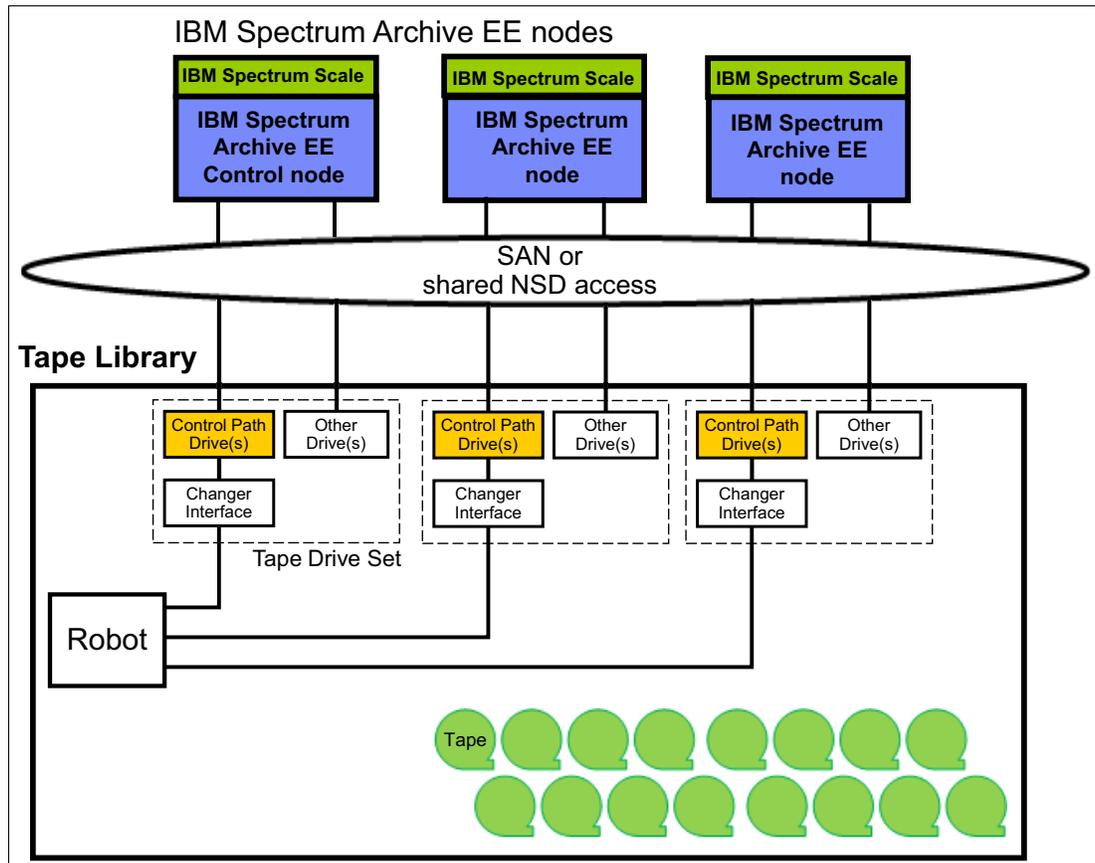


Figure 1-8 Single-library multiple-node access setup

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 physical changes on the cartridges are handled correctly.



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:

- ▶ 2.1, “Introduction to IBM Spectrum Archive and LTFS” on page 18
- ▶ 2.2, “IBM Spectrum Scale” on page 27
- ▶ 2.3, “OpenStack SwiftHLM” on page 33
- ▶ 2.4, “IBM Spectrum Archive EE dashboard” on page 34
- ▶ 2.5, “IBM Spectrum Archive EE REST API” on page 36
- ▶ 2.6, “Types of archiving” on page 36

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^a
- ▶ IBM TS4300 tape library
- ▶ IBM TS3500 tape library^{ab}
- ▶ IBM TS3310 tape library
- ▶ IBM LTO Ultrium 9, 8, 7, 6, or 5 tape drives, and IBM TS1160 (3592 60E, 60F, 60G, or 60S), TS1155 (3592 55F or 55G), TS1150, or TS1140 tape drives
- ▶ LTO Ultrium 9, 8, M8^c, 7, 6, and 5, and 3592 JB, JC, JD, JE, JK, JL, JM, JV, JY, and JZ tape media

For more information about the latest system requirements, see [IBM Documentation](#).

For the latest IBM Spectrum Archive Library Edition Support Matrix (supported tape library and tape drive firmware levels, see IBM Support's [Fix Central web page](#).

- a. IBM TS1160, TS1155, TS1150, and IBM TS1140 support on IBM TS4500 and IBM TS3500 tape libraries only.
- b. TS3500 does not support the LTO-9 tape drive.
- c. 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 another 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 10+ years 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 v1.2, corresponding Write Once, Read Many (WORM) tape cartridges are supported in an IBM Spectrum Archive EE solution that operates supported IBM Enterprise tape drives. With the same release, tape drives in mixed configurations are supported. For more information, see 6.21, “Tape drive intermix support” on page 208.

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 to the server. 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 ^a	Native capacity ^{b, c}	LTFS data partition size ^{cd}
IBM TS1160 tape drive ^{e, f}	Advanced type E data (JE)	20000 GB (18626 GiB)	19485 GB (18147 GiB)
	Advanced type E WORM data (JV)	20000 GB (18626 GiB)	19485 GB (18147 GiB)
	Advanced type D data (JD)	15000 GB (13969 GiB)	14562 GB (13562 GiB)

Tape drive	Tape media^a	Native capacity^{b, c}	LTFS data partition size^{cd}
	Advanced Type D WORM data (JZ)	15000 GB (13969 GiB)	14562 GB (13562 GiB)
	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)
	Advanced Type E economy tape (JM)	5000 GB (4656 GiB)	4870 GB (4536 GiB)
	Advanced Type D economy tape (JL)	3000 GB (1862 GiB)	2912 GB (1804 GiB)
	Advanced Type C economy tape (JK)	900 GB (838 GiB)	869 GB (809 GiB)
IBM TS1155 tape drive ^{e, f}	Advanced type D data (JD)	15000 GB (13969 GiB)	14562 GB (13562 GiB)
	Advanced type D WORM data (JZ)	15000 GB (14969 GiB)	14562 GB (13562 GiB)
	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)
	Advanced type D economy data (JL)	3000 GB (1862 GiB)	2912 GB (1804 GiB)
	Advanced type C economy data (JK)	900 GB (838 GiB)	869 GB (809 GiB)
IBM TS1150 tape drive ^{e, f}	Advanced type D data (JD)	10,000 GB (9313 GiB)	9687 GB (9022 GiB)
	Advanced type D WORM data (JZ)	10,000 GB (9313 GiB)	9687 GB (9022 GiB)
	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)
	Advanced type D economy data (JL)	3000 GB (2794 GiB)	2912 GB (2712 GiB)
	Advanced type C economy data (JK)	900 GB (838 GiB)	869 GB (809 GiB)

Tape drive	Tape media ^a	Native capacity ^{b, c}	LTFS data partition size ^{cd}
IBM TS1140 tape drive ^{e, f}	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)
	Advanced 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 9 tape drive	LTO 9	18000 GB (16764 GiB)	17549 GB (16344 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)
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 [here](#).
d. Values that are given are the default size of the LTFS data partition, unless otherwise indicated.
e. TS1160, 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.

2.1.2 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 was discontinued 05/01/2020.

Table 2-2 LTFS product comparison

Name	License required	Market	Tape Library support	Integrates with IBM Spectrum Scale
IBM Spectrum Archive Single Drive Edition (SDE)	No	Entry - Midrange	No	No
IBM Spectrum Archive Library Edition (LE)	ILAN license (free) or commercial license (for support from IBM)	Midrange - Enterprise	Yes	No
IBM Spectrum Archive Enterprise Edition (EE)	Yes	Enterprise	Yes	Yes

2.1.3 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.

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

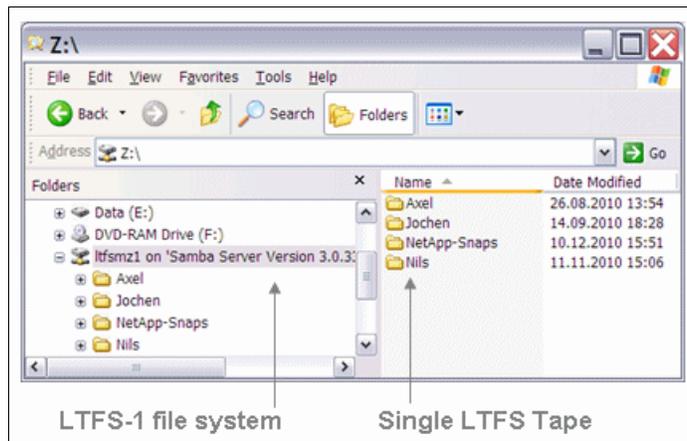


Figure 2-1 IBM Spectrum Archive SDE user view

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 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 Documentation web page](#).

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 *Mounting media by using the ltfscmd* topic at [IBM Documentation](#).

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 *Assigning a drive letter to a tape drive* topic at [IBM Documentation](#).

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](#)) [web page](#).

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 IBM Spectrum Archive SDE software must be uninstalled before IBM Spectrum Archive EE is installed. Follow the uninstallation procedure that is documented in the *IBM Linear Tape File System Installation and Configuration*, 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 6.19.1, “Importing” on page 202. Tapes that were formatted in LTFS 1.0 format by older versions of IBM Spectrum Archive are automatically upgraded to LTFS 2.4 format on first write.

2.1.4 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 or IBM Enterprise supported tape drives.

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

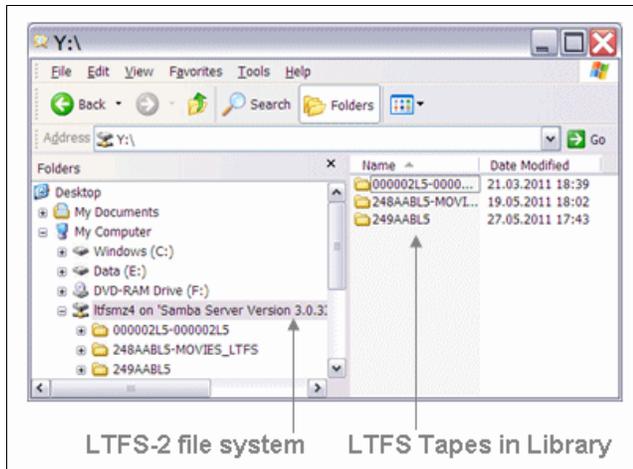


Figure 2-2 IBM Spectrum Archive LE User view of multiple LTFS tape cartridges

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 Documentation website](#).

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](#).

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 IBM Spectrum Archive LE software must be uninstalled before IBM Spectrum Archive EE is installed. Follow the uninstall procedure that is described at this [IBM Documentation web page](#).

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 6.19.1, “Importing” on page 202.

2.1.5 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.

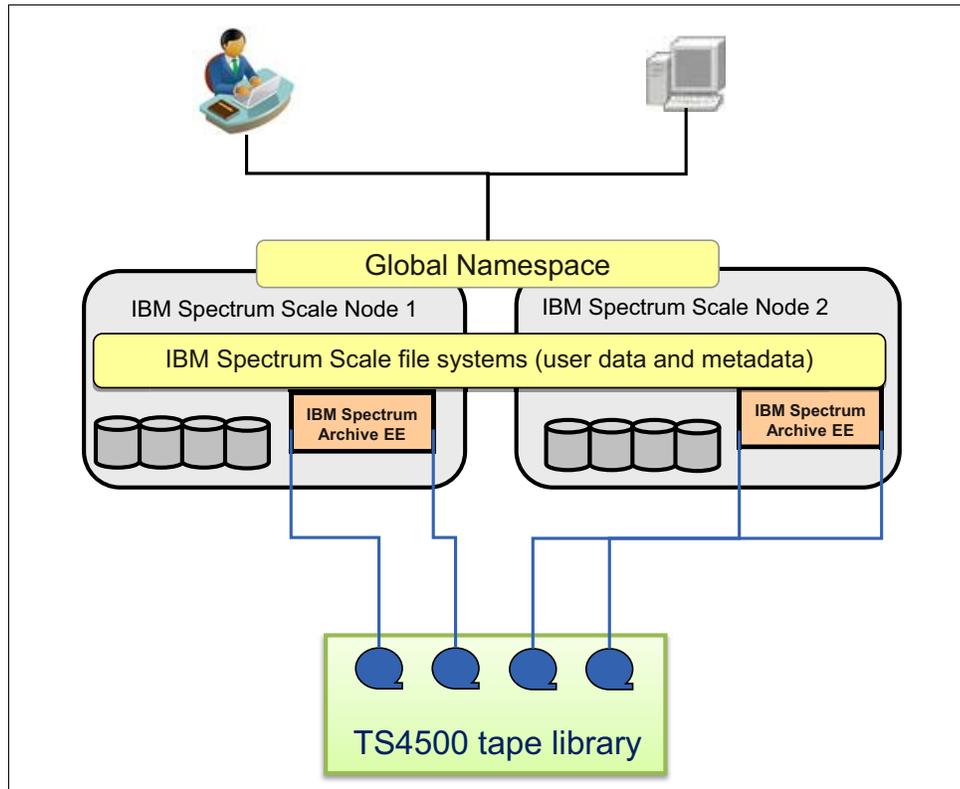


Figure 2-3 IBM Spectrum Archive EE system view

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 (over 10 years).

- ▶ Provides efficient recalls of files from tape with Recommended Access Order (RAO) supported tape drives and media. For more information, see 6.14.4, “Recommend Access Order” on page 193, and Table 3-1, “Linux system requirements” on page 50.
- ▶ 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.

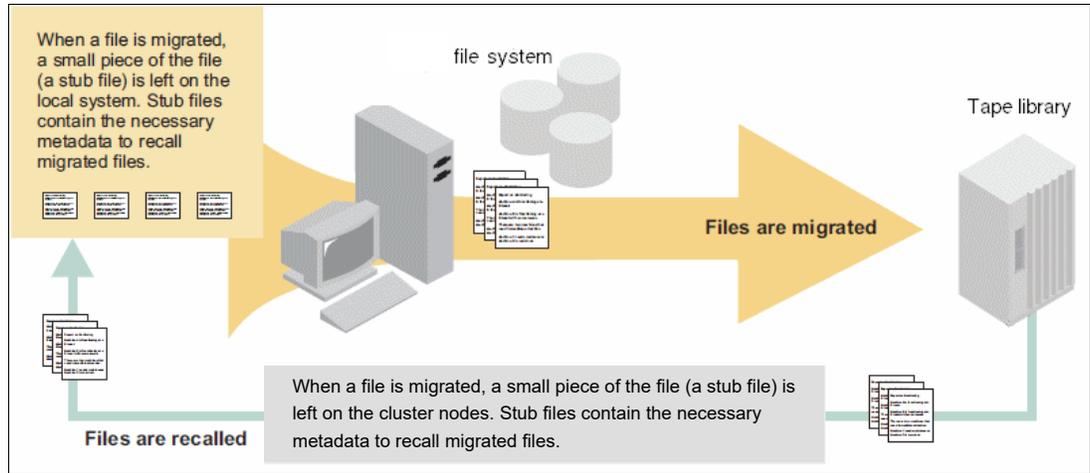


Figure 2-4 IBM Spectrum Archive EE data flow

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 storage 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.

2.2 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 5.1.2. The first version with the name IBM Spectrum Scale is Version 4.1.1. 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.

For more information about current documentation and publications for IBM Spectrum Scale, see this [IBM Documentation web page](#).

2.2.1 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.

2.2.2 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

2.2.3 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 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 RAIDs 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 [IBM Documentation web page](#).

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 6.11, “Migration” on page 165.

2.2.4 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 **eeadm migrate** or **eeadm 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 6.11, “Migration” on page 165.

2.2.5 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 31 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 supports only the independent-writer (IW) cache mode.

Independent-writer

The IW cache mode of AFM 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.

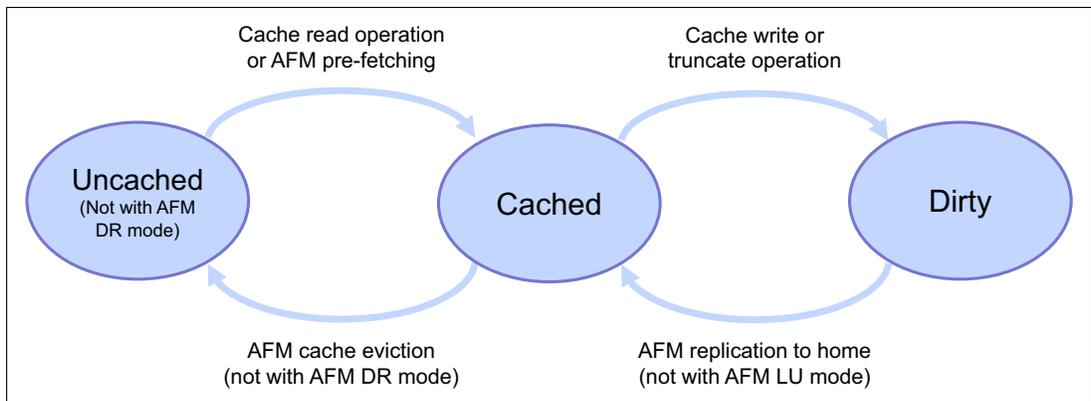


Figure 2-5 AFM file state transitions

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 7.10.3, “IBM Spectrum Archive EE migration policy with AFM” on page 246. For use cases, see Figure 8-11, “IBM Spectrum Archive use case for university Scientific Data Archive” on page 286.

2.2.6 Scale Out Backup and Restore

Scale Out Backup and Restore (SOBAR) is an IBM Spectrum Scale data protection mechanism for Disaster Recovery (DR) incidents. SOBAR is used to back up and restore IBM Spectrum Scale files that are managed by IBM Spectrum Protect for Space Management.

The idea of SOBAR is to have all the file data pre-migrated or migrated to the IBM Spectrum Protect server and periodically back up the file system configuration and metadata of all the directories and files into the IBM Spectrum Protect server.

When the file system recovery is needed, the file system is restored by using the configuration, and all the directories or files are restored by using the metadata. After the recovery, the files are in stub format (HSM migrated state) and can be recalled from the IBM Spectrum Protect server on-demand.

Because the recovery processes only the metadata and does not involve recalling or copying the file data, the recovery is fast with significantly reduced recovery time objective (RTO). This is especially true for large file systems.

For more information, see this [IBM Documentation web page](#).

As of this writing, SOBAR is not supported on the IBM Spectrum Archive managed file systems for DR. However, from version 1.3.1.2, IBM Spectrum Archive supports a procedure that uses SOBAR to performed a planned data migration between two file systems.

By using SOBAR, a file system can be migrated without recalling the file data from tapes. For more information, see 6.27, “File system migration” on page 231.

2.3 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 34.

For more information, see *Implementing OpenStack SwiftHLM with IBM Spectrum Archive EE or IBM Spectrum Protect for Space Management*, [REDP-5430](#).

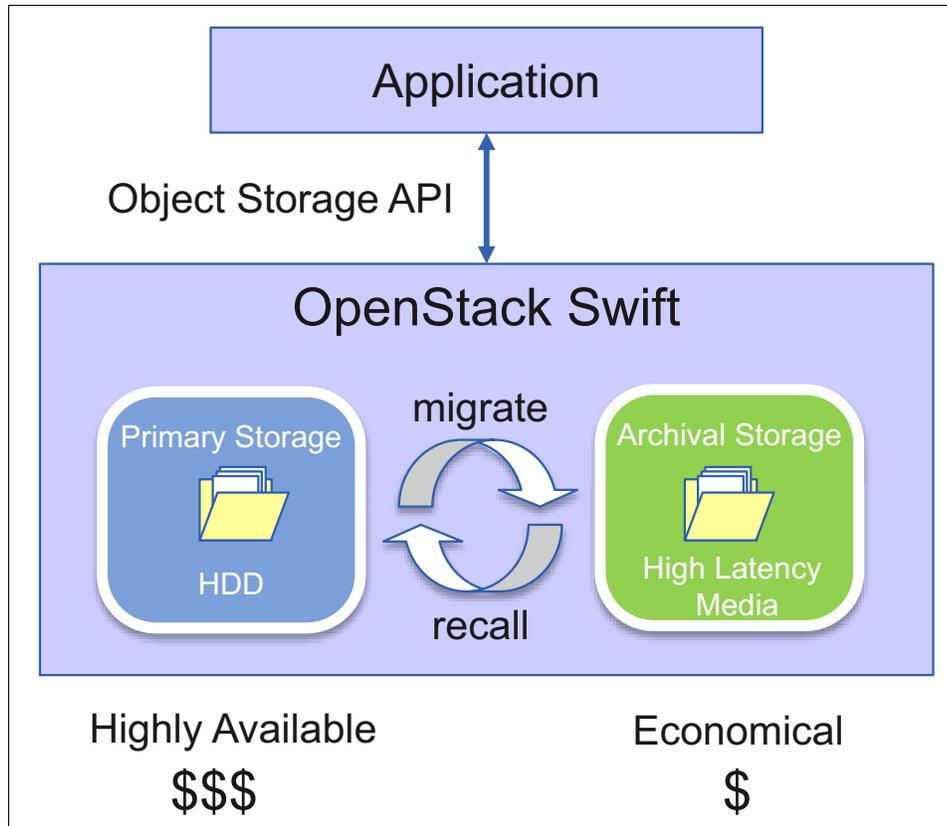


Figure 2-6 High-level SwiftHLM architecture

2.4 IBM Spectrum Archive EE dashboard

IBM Spectrum Archive EE provides 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.
- ▶ See current running/waiting tasks

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.

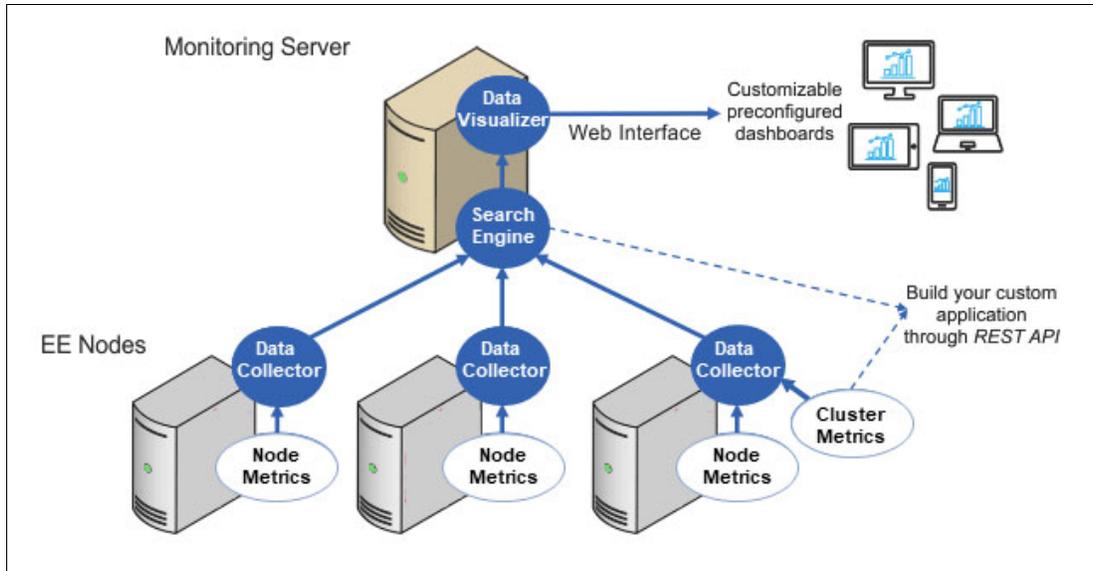


Figure 2-7 IBM Spectrum Archive Dashboard architecture

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

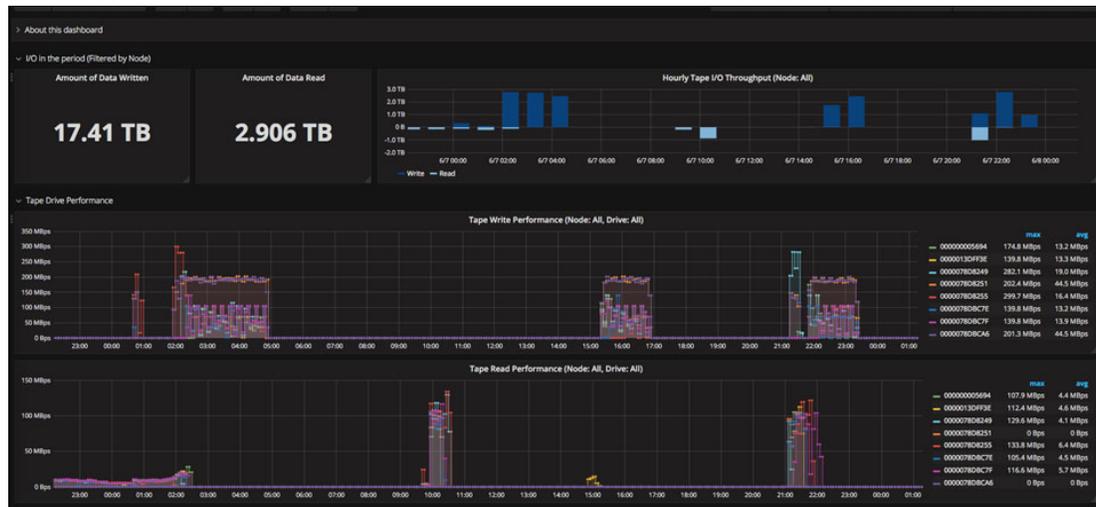


Figure 2-8 IBM Spectrum Archive EE Dashboard activity view

For more information on configuring the Dashboard within your environment, see the IBM Spectrum Archive Enterprise Edition Dashboard Deployment Guide, which is available at this [IBM Documentation web page](#).

2.5 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 `eeadm task list` and `eeadm task 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 83.

For usage examples including commonly used parameters, see 6.26, “IBM Spectrum Archive REST API” on page 219.

2.6 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 for IBM Spectrum Archive Enterprise Edition

This chapter provides planning information that is related to the IBM Spectrum Archive Enterprise Edition (EE). Review the Planning section in the IBM Spectrum Archive EE that is available at this [IBM Documentation web page](#).

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:

- ▶ 3.1, “IBM Spectrum Archive EE deployment options” on page 40
- ▶ 3.2, “Data-access methods” on page 46
- ▶ 3.3, “System requirements” on page 50
- ▶ 3.4, “Required software for Linux systems” on page 54
- ▶ 3.5, “Hardware and software setup” on page 56
- ▶ 3.6, “IBM Spectrum Archive deployment examples” on page 57
- ▶ 3.7, “Sizing and settings” on page 61
- ▶ 3.8, “High-level component upgrade steps” on page 70

3.1 IBM Spectrum Archive EE deployment options

It is important to understand the target environment which IBM Spectrum Archive Enterprise Edition will be deployed. Resources can be used as part of your deployment. Configuration flexibility is one of the many advantages of IBM software-defined storage solutions that provide high business-value while streamlining long-term data-retention storage costs.

IBM Spectrum Archive Enterprise Edition is to be installed and configured on one or more IBM Spectrum Scale nodes. IBM Spectrum Archive EE installed nodes make up an IBM Spectrum Archive EE cluster within an IBM Spectrum Scale cluster. The IBM Spectrum Archive EE nodes need to be attached to a supported tape library facility. A maximum of two tape libraries are supported for an IBM Spectrum Archive EE cluster.

From an IBM Spectrum Scale architecture perspective, the IBM Spectrum Archive EE software-defined storage can be installed and configured on an IBM Spectrum Scale Client node, an IBM Spectrum Scale Server node, or an IBM Spectrum Scale Protocol node. The node that you choose must have the supported requirements for IBM Spectrum Archive EE.

Note: IBM Spectrum Archive EE can NOT be installed on IBM Elastic Storage® Server (IBM ESS) nodes. See 3.3.1, “Limitations” on page 51.

This section describes some of the typical deployment configurations for IBM Spectrum Archive EE and provides insights to help plan for its implementation.

3.1.1 On IBM Spectrum Scale Servers

On this architecture, the IBM Spectrum Archive EE is deployed within the same server as that of the IBM Spectrum Scale. This approach is common in entry- to mid-level solutions that consolidate the architecture to a minimum solution. A common use case for this architecture is for Active Archive solutions, which require only a relatively small amount of flash or disk staging tiers for data migration to the tape tier.

Figure 3-1 shows the physical and its equivalent logical diagram that show the minimum architecture approach for IBM Spectrum Archive EE.

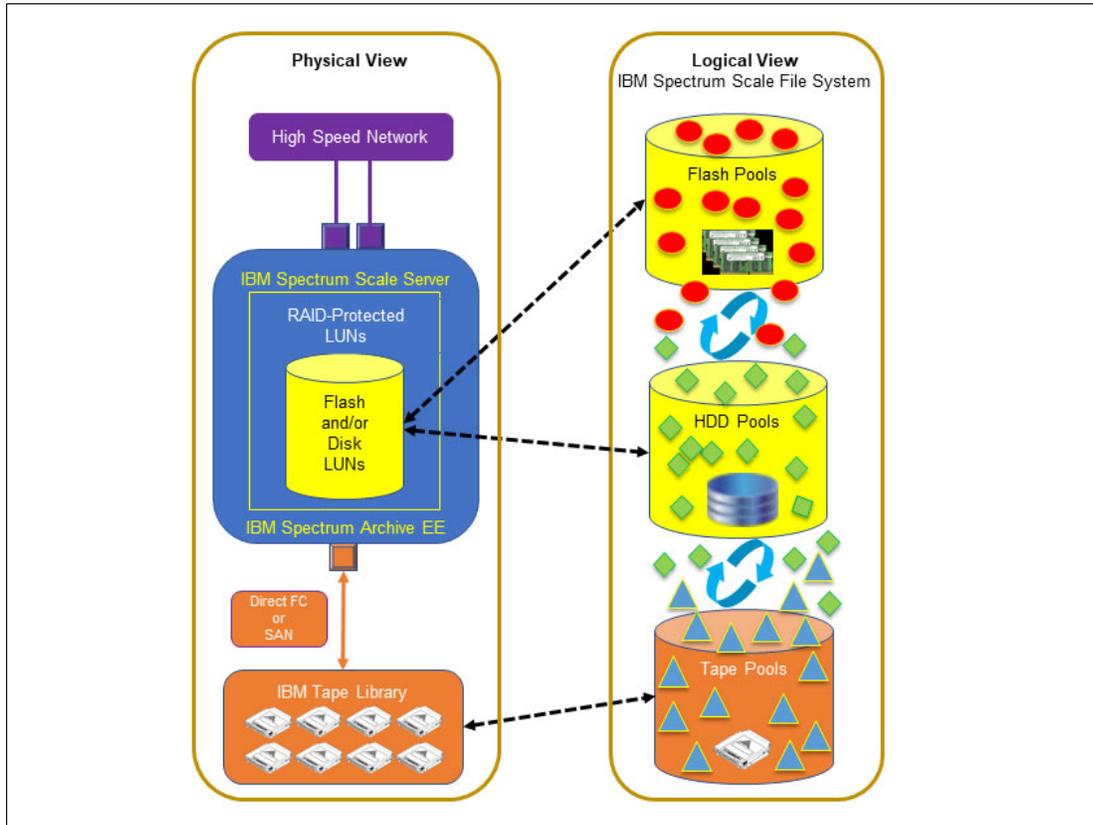


Figure 3-1 Deployed on IBM Spectrum Scale Server

The required software components of IBM Spectrum Scale and IBM Spectrum Archive EE are installed in one physical server either directly Fibre Channel-attached to an IBM tape library or using a storage area network. The RAID-protected LUNs and IBM tape library are configured in IBM Spectrum Scale and IBM Spectrum Archive EE as pools - flash pools, spinning disk pools and tape pools. This architecture deployment option enables automated data tiering based on user-defined policies with minimum solution components.

This architecture is also known as Long-Term Archive Retention (LTAR). However, as this solution is usually deployed for long term data archive retention, injecting high availability is highly recommended.

Figure 3-2 shows a high availability architecture approach to the minimum solution.

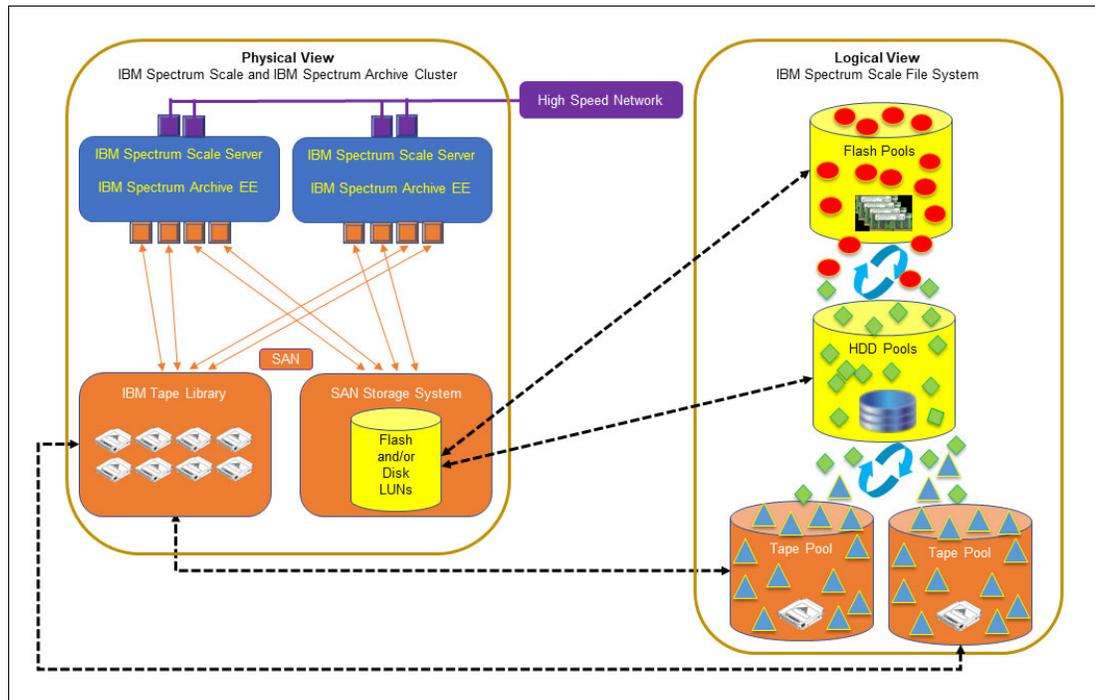


Figure 3-2 High availability approach IBM Spectrum Archive EE on IBM Spectrum Scale Servers

Two physical servers are deployed each with IBM Spectrum Scale and IBM Spectrum Archive EE software components. Both servers are connected to a storage area network which provides connectivity to a shared storage area network (SAN) storage system and an IBM tape library. This approach ensures that both servers are connected to the storage systems by way of a high-speed Fibre Channel network. The connectivity of the cluster to a shared flash or disk storage system provides high performance and high availability in the event of a server outage. For more information about configuring a multiple-node cluster, refer to 5.2.3, “Configuring a multiple-node cluster” on page 112.

Another approach to implementing high availability for data on tapes is the deployment of two or more tape pools. Policies can be configured to create up to three copies of data in the additional tape pools.

Note: Multiple tape pools also can be deployed on the minimum solution, as shown in Figure 3-1 on page 41.

3.1.2 As an IBM Spectrum Scale Client

On this architecture, the IBM Spectrum Archive EE node is deployed as a separate software-defined storage module integrating with the IBM Spectrum Scale single namespace file system with its own Storage Area Network interface to the tape platform.

Figure 3-3 shows this unique approach that leverages software-defined storage for flexibility of scaling the system.

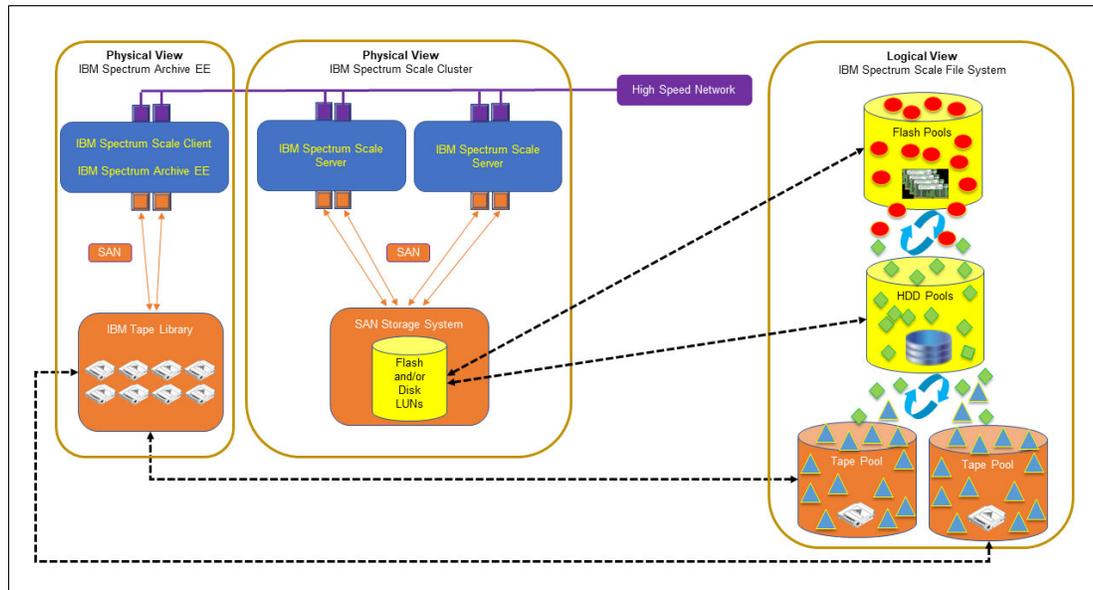


Figure 3-3 IBM Spectrum Archive EE node as an IBM Spectrum Scale Client

The diagram shows a single node IBM Spectrum Archive EE as an IBM Spectrum Scale Client with a single tape library attached using SAN. However, it is easy to add IBM Spectrum Archive EE nodes that are attached to the same library using SAN for high-availability or higher tape-access throughput.

This is the solution approach for customers with IBM Spectrum Scale environments who want to use tapes to stream line costs for long-term archive retention. This architecture provides a platform to scale both IBM Spectrum Scale and IBM Spectrum Archive EE independently, that is by adding nodes to each of the clusters when needed. The integration of the server nodes with the storage systems can be done over the same SAN fabric considering the appropriate SAN zoning implementations. This way, scaling the IBM Spectrum Scale flash- or disk-storage capacity is accomplished by merely adding flash or disk drives to the storage server.

3.1.3 As an IBM Elastic Storage Systems IBM Spectrum Scale Client

This architecture is similar to the architecture described in 3.1.2, “As an IBM Spectrum Scale Client” on page 43. However, instead of an IBM Spectrum Scale cluster of nodes, the IBM ESS appliance is deployed. The IBM Elastic Storage Systems (ESS) is a pre-integrated, pre-tested appliance that implements the IBM Spectrum Scale file system in the form of building blocks. The IBM Spectrum Archive EE tightly integrates with the IBM ESS, enabling automated data-tiering of the file system to and from tapes based on policies.

Figure 3-4 shows this architecture wherein an IBM Spectrum Archive EE node is integrated with an IBM Spectrum Scale architecture deployed on an IBM ESS system.

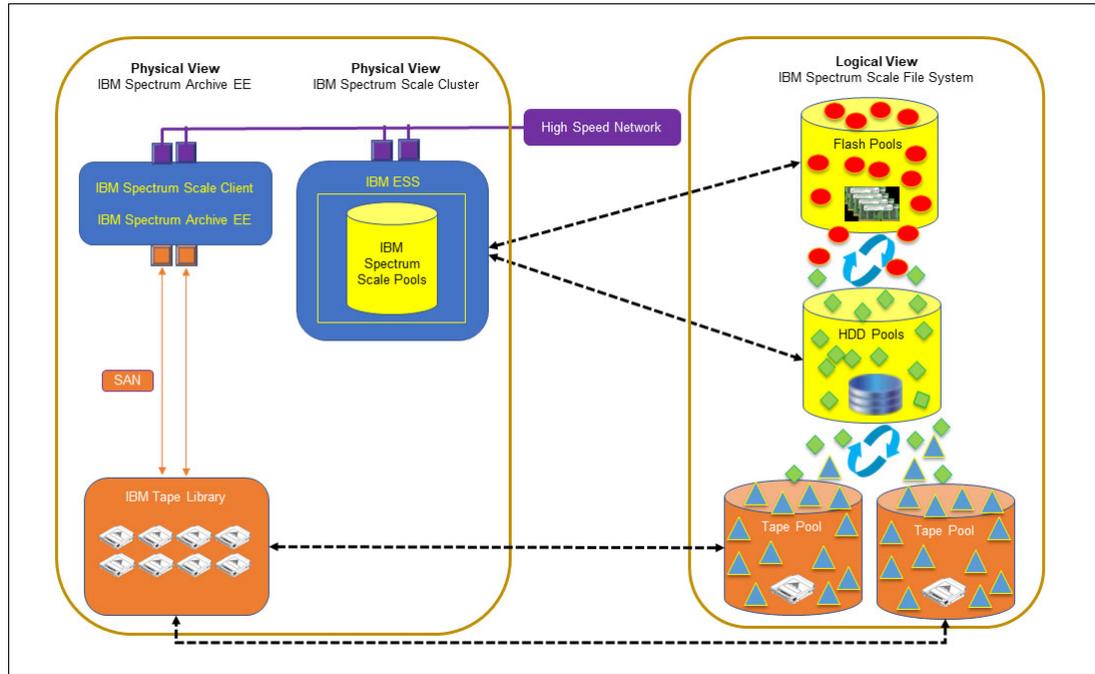


Figure 3-4 IBM Spectrum Archive EE node as an IBM Spectrum Scale Client in an IBM ESS

This architecture allows the same easy scalability feature as the IBM Spectrum Scale Client described in 3.1.2, “As an IBM Spectrum Scale Client” on page 43; With this scalability feature, you can add nodes when needed. In Figure 3-4, you can add IBM ESS modules to expand the IBM Spectrum Scale cluster; or you can add IBM Spectrum Archive EE nodes to expand the file system tape facility channels. Needless to say, adding tape drives and tape cartridges to address growing tape access and storage requirements is also simple and nondisruptive.

3.1.4 As an IBM Spectrum Scale stretched cluster

An IBM Spectrum Scale stretched cluster is a high-availability solution that can be deployed across two sites over metropolitan distances of up to 300 km apart. Each site will have both IBM Spectrum Scale and IBM Spectrum Archive EE clusters integrated with a tape library. These two sites form a single namespace, thus simplifying data-access methods. Data created from one site is replicated to the other site and might be migrated to tapes based on the user-defined policies to protect data. This way, if one site failure occurs, data will still be accessible on the other site.

Figure 3-5 shows a single namespace file system IBM Spectrum Scale stretched cluster architecture deployed across two sites.

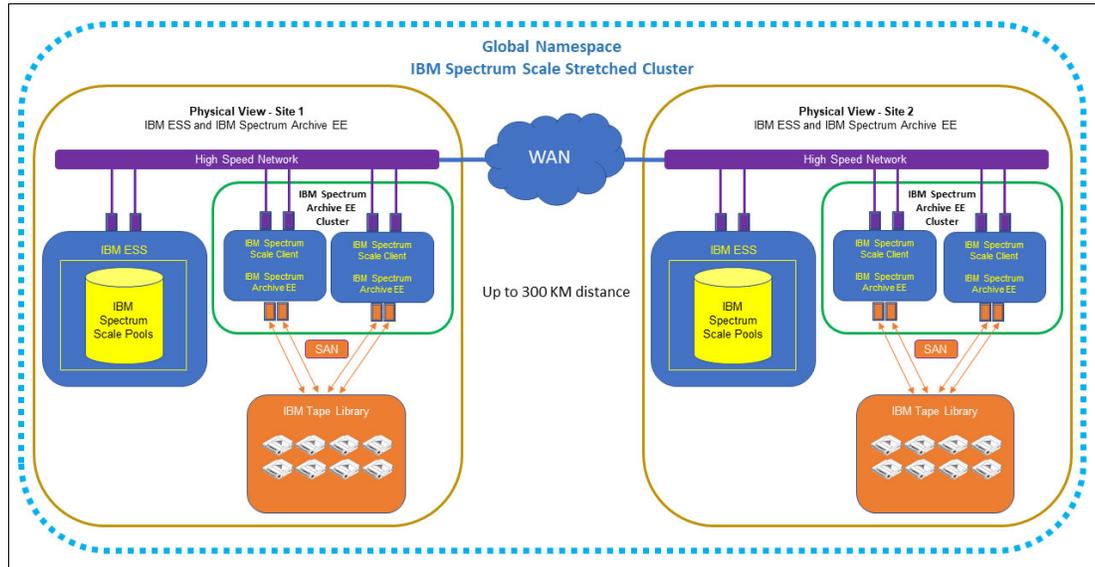


Figure 3-5 IBM Spectrum Scale Stretched Cluster across metropolitan distances of up to 300 KM

This architecture shows a good use case for deploying two tape libraries in one IBM Spectrum Scale cluster. For sites that are geographically separated by greater than 300 KM, Active File Management (AFM) might be deployed to replicate data across the sites. However, both sites are independent IBM Spectrum Scale clusters as compared to the single namespace file system over a stretched cluster.

For more information about stretched clusters, see 5.2.4, “Configuring a multiple-node cluster with two tape libraries” on page 115.

For example use cases on stretched clusters and AFM, see 8.8, “University Scientific Data Archive” on page 286 and 8.11, “AFM use cases” on page 290.

Note: Figure 3-5 shows the IBM Spectrum Scale cluster deployed as an IBM ESS appliance. The cluster might also be a software-defined storage deployment on IBM Spectrum Scale Server nodes as discussed in 3.1.1, “On IBM Spectrum Scale Servers” on page 40 and 3.1.2, “As an IBM Spectrum Scale Client” on page 43.

3.2 Data-access methods

This section describes how applications or users can access data in the IBM Spectrum Scale file system that has an integrated IBM Spectrum Archive EE facility. It is important to note that IBM Spectrum Scale is a single namespace file system by which data can be automatically tiered based on user-defined policies; and IBM Spectrum Archive EE is the channel of IBM Spectrum Scale to utilize the tape platform as a storage tier of the single namespace file system.

There are three ways of accessing data on an IBM Spectrum Archive EE deployed IBM Spectrum Scale system:

3.2.1 Data access using application or users on IBM Spectrum Scale Clients

This access method refers to compute nodes directly accessing the IBM Spectrum Scale System. Each compute node on which the applications are running, needs to have an IBM Spectrum Scale Client installed for high-speed data access to the IBM Spectrum Scale Cluster.

Figure 3-6 shows that the IBM Spectrum Scale Client is installed on the compute nodes where the applications are running.

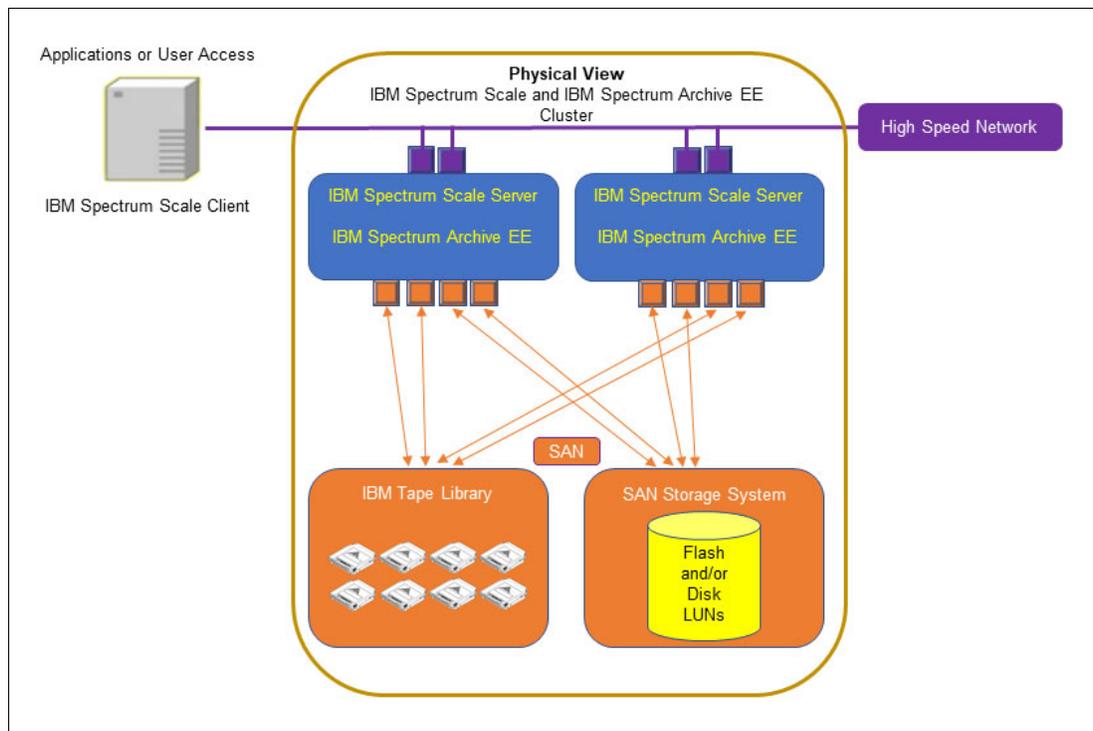


Figure 3-6 Access using IBM Spectrum Scale Client on compute node

Note: Each compute node must have IBM Spectrum Scale Clients installed to access the IBM Spectrum Scale Cluster. This allows high-performance (network dependent) access to the cluster.

3.2.2 Data access using Protocol Nodes on IBM Spectrum Scale Server Nodes with IBM Spectrum Archive EE

This access method can be more “application-friendly” because alterations or software modules are not required to be installed on the application compute-nodes. Access to the IBM Spectrum Scale Cluster is gained by using the Protocol Nodes which offers other choices of data interface. The following data interfaces might be used by the application compute-nodes through the IBM Spectrum Scale Protocol Nodes:

- ▶ Network File System (NFS)
- ▶ Server Message Block (SMB)
- ▶ Hadoop Distributed File System (HDFS)
- ▶ Object

Figure 3-7 shows this architecture where Protocol Nodes are deployed on the same servers as that of the IBM Spectrum Scale Server and IBM Spectrum Archive EE nodes.

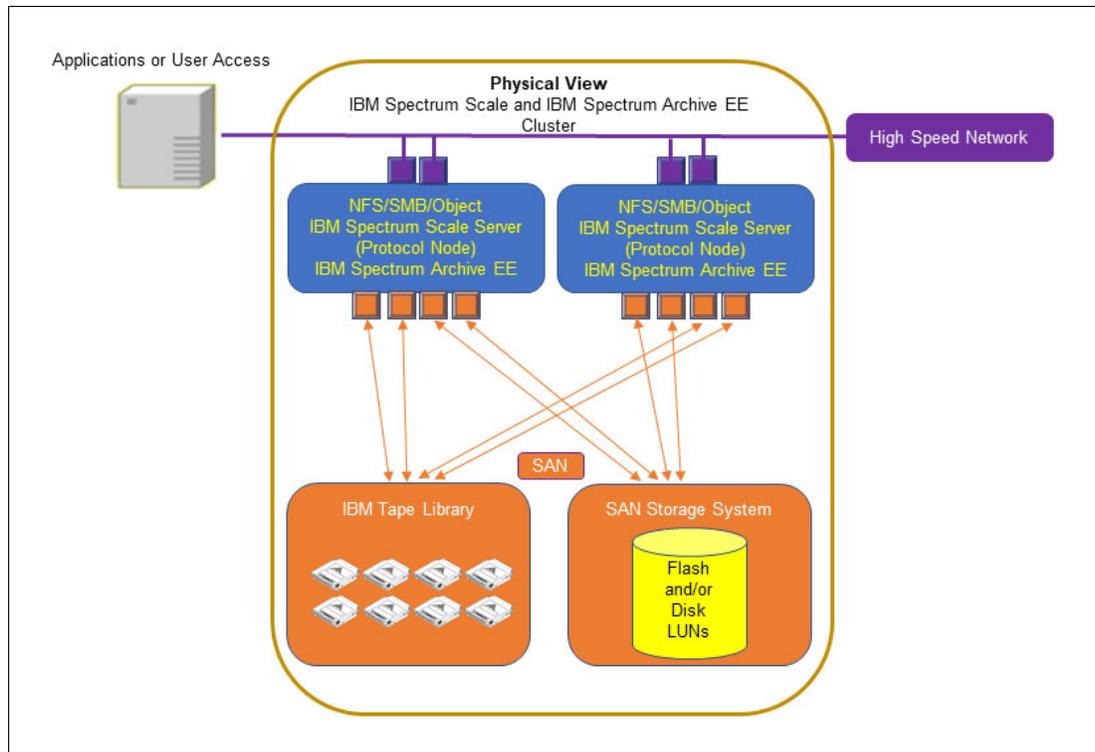


Figure 3-7 Access using Protocol Nodes on IBM Spectrum Scale and IBM Spectrum Archive EE nodes

This architecture is usually deployed on Operational Archive and Active Archive use cases where each cluster host runs IBM Spectrum Scale and IBM Spectrum Archive EE software-defined storage modules. The scalability potential of each entity (application compute-nodes, and IBM Spectrum Scale or IBM Spectrum Archive EE nodes) remains independent of each other. This architecture might provide streamlined costs for long-term archive retention solution requirements.

Figure 3-8 shows another version of this access method, where Application Compute Nodes access the IBM Spectrum Scale Cluster using Protocol Nodes. However, this architecture deploys the Protocol Nodes in separate hosts or virtual machines. This way, scaling is much simpler, that is, you can add the required nodes when needed.

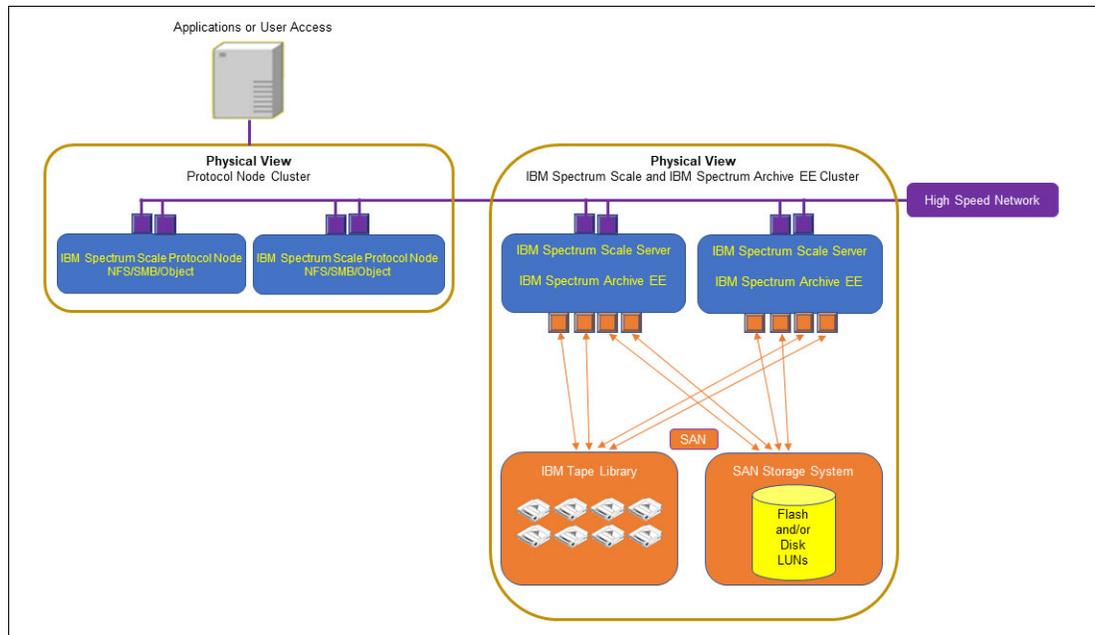


Figure 3-8 Access using Protocol Nodes on separate hosts

3.2.3 Data access using Protocol Nodes integrated with the IBM ESS

This access method is similar to Figure 3-8. However, instead of software-defined storage modules implemented on servers or hosts, the IBM Spectrum Scale system is deployed on an IBM ESS environment. The IBM Spectrum Archive EE environment is also deployed independently as an IBM Spectrum Scale Client, which provides flexible scalability for all components.

Figure 3-9 shows an IBM ESS system integrated with an IBM Spectrum Archive EE node. On this architecture, a separate cluster of Protocol Nodes are deployed to provide applications and users with access to the IBM Spectrum Scale in the IBM ESS system.

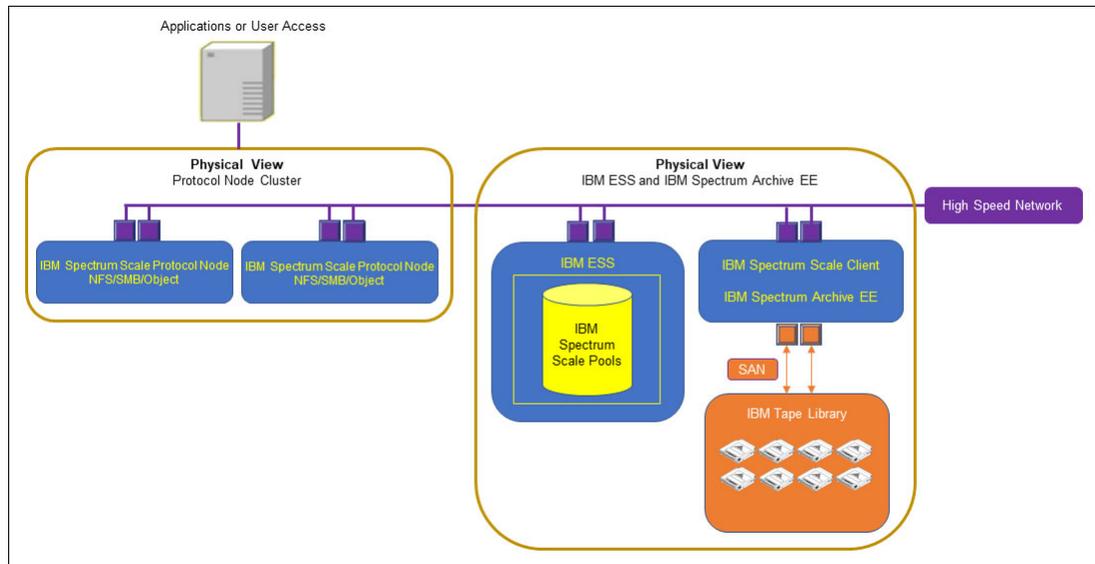


Figure 3-9 Access using Protocol Node cluster to IBM ESS

This architecture uses the pre-tested and pre-installed IBM Spectrum Scale environment as an IBM ESS appliance. Scaling this architecture is simplified, where upgrades can be deployed to the independent clusters (which make up the IBM Spectrum Scale single namespace filesystem). Also, note that IBM Spectrum Archive EE nodes can be added to this architecture for high availability or to increase tape access performance.

3.3 System requirements

IBM Spectrum Archive EE supports the Linux operating systems and hardware platforms that are shown in Table 3-1.

Table 3-1 Linux system requirements

Linux computers	
Supported operating systems (x86_64)	Red Hat Enterprise Linux Server 7.4, 7.5, 7.6, 7.7, 7.8, 7.9, 8.2, 8.3, and 8.4
Supported operating systems (ppc64le)	Red Hat Enterprise Linux Server 7.4, 7.5, 7.6, 7.7, 7.8, 7.9, 8.2, 8.3, and 8.4
Supported tape libraries	<p>IBM Spectrum Archive Enterprise Edition supports up to two tape libraries of the following types:</p> <ul style="list-style-type: none"> ▶ IBM TS4500 tape library (IBM LTO 5 and later generation LTO tape drives, TS1140 and later generation 3592 tape drives) ▶ IBM TS4300 tape library (IBM LTO 6 and later generation LTO tape drives) ▶ IBM TS3500 tape library (IBM LTO 5 through LTO 8 generation LTO tape drives, TS1140 and later generation 3592 tape drives) ▶ IBM TS3310 tape library (IBM LTO 5 and later generation LTO tape drives)
Supported tape drives	IBM TS1140 tape drive
	IBM TS1150 tape drive
	IBM TS1155 tape drive
	IBM TS1160 tape drive
	LTO-5 tape drive
	LTO-6 tape drive
	LTO-7 tape drive
	LTO-8 tape drive
	LTO-9 tape drive
Supported tape media	TS1140 media: JB, JC, JK, and JY TS1155/TS1150 media: JC, JD, JK, JL, JY, and JZ TS1160 media: JC, JD, JE, JK, JL, JM, JV, JY, and JZ
	LTO media: LTO 9, LTO 8, LTO 7 Type M8, LTO 7, LTO 6, and LTO 5
Recommended Access Order (RAO) support	IBM TS1140, IBM TS1150, IBM TS1155, IBM TS1160, and LTO-9
Server	
Processor	<ul style="list-style-type: none"> ▶ One of the following servers: <ul style="list-style-type: none"> – x86_64 processor (physical server only) – IBM POWER® ppc64le processor (either physical server or IBM PowerVM® partition) ▶ Minimum: A x86_64 processor ▶ Preferred: Dual socket server with the latest chipset

Linux computers	
Memory	<ul style="list-style-type: none"> ▶ Minimum: $2 \times (d) \times (f) + 1$ GB of RAM available for the IBM Spectrum Archive EE program: <ul style="list-style-type: none"> – d: Number of tape drives – f: Number of millions of files/directories on each tape cartridge in the system 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 each tape cartridge. The minimum required RAM is 37 GB ($2 \times 6 \times 3 + 1 = 37$). ▶ Preferred: 64 GB RAM and greater
Host Bus Adapter (HBA) ^a , RoCE	<ul style="list-style-type: none"> ▶ Minimum: Fibre Channel Host Bus Adapter supported by TS1160, TS1155, TS1150, TS1140, LTO-9, 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
Disk device for LTFS EE tape file system metadata	For more information, see 3.7, “Sizing and settings” on page 61.
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.

a. For more information about HBA interoperability, see the IBM System Storage Interoperation Center (SSIC) [web page](#).

3.3.1 Limitations

In this section, we describe the limitations of IBM Spectrum Archive EE.

Limitations on supported files

This section describes the limitations of IBM Spectrum Archive EE on file attributes when migrating supported files.

Maximum file size

IBM Spectrum Archive EE cannot split a file into multiple sections and distribute the sections across more than one tape. Therefore, it cannot migrate a file that is larger than the data partition size of a single tape. For more information about media types and partition sizes, see 2.1.1, “Tape media capacity with IBM Spectrum Archive” on page 19.

Maximum file name length

Consider the following points:

- ▶ A file cannot be migrated if its full path name (file name length plus path length) exceeds 1024 bytes.
- ▶ A file cannot be saved if its full path name (file name length plus path length) exceeds 1022 bytes.

Minimum file size

Consider the following points:

- ▶ A nonzero length regular file can be migrated to tape with the `eeadm migrate` or `eeadm premigrate` commands.
- ▶ The name of an empty (zero length) file, and a symbolic link and an empty directory, can be stored on tape with the `eeadm save` command.

File encryption and file compression

When the file encryption and file compression function of GPFS is used, the in-flight data from disk-storage to tape is in the decrypted and uncompressed form. The data is re-encrypted and recompressed by using the tape hardware function. (Tape encryption requires the library to be set up with key manager software.)

Hard links on files

Migration of files with hard links to tape is discouraged in IBM Spectrum Archive EE.

The hard link creates an alias for the data with different name or directory path without copying it. The original file and other hard links all point to the same data on disk. The view from other hard links is affected if the file that is linked from one hard link is changed.

If multiple hard links are necessary on the IBM Spectrum Archive EE managed file system, those files must be excluded from policy files; otherwise, commands can have unexpected results for the user. The following consequences can occur on running the commands:

- ▶ Migration, Recall, and Premigrate commands
If one of the hard links is specified in a migration task, all hard links of that file are shown as migrated. If multiple hard links are specified in a single migration task, the migration results in a “duplicate” error. If one of the hard links is specified in a recall task, all hard links of that file are shown as recalled. If multiple hard links are specified in a single recall task, the recall results in a “duplicate error”.
- ▶ Save commands
The hard link information is not saved with the save command because the objects contain data.
- ▶ File behavior on tape export and imports
The tape export process can cause unintentional unlinks from I-nodes. Even if all hard links indicate that they are migrated to tape, multiple hard links for a single I-node is not created during the tape import.

Limitations with IBM Spectrum Protect

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.
Specify the `--mmbackup` option of the `eeadm migrate` command to ensure the files are backed up first before migration. For more information, see 7.26.2, “Backing up a GPFS or IBM Spectrum Scale environment” on page 265.
- ▶ 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.

Limitations with IBM Spectrum Scale

This section describes the limitations of IBM Spectrum Archive EE with specific IBM Spectrum Scale functions.

Limitations on IBM Spectrum Scale features

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) for DR purposes
- ▶ Advanced File Management (AFM), in a mode other than Independent Writer mode
- ▶ Transparent Cloud Tiering (TCT)

Limitations with the Snapshot function

To prevent an unexpected massive recall of files from tapes, it is *not* recommended to use IBM Spectrum Scale file system snapshots or file set snapshots when the files are managed by IBM Spectrum Archive EE. The massive recall of files happens if a snapshot is taken after the files are migrated to tapes and then later, the user deletes the migrated files before the snapshot.

Cluster Network File System (CNFS)

To prevent a Network File System (NFS) failover from happening on an IBM Spectrum Archive EE node, do not install IBM Spectrum Archive EE on Cluster NFS (CNFS) member nodes.

IBM Spectrum Archive EE with IBM ESS

IBM Spectrum Archive EE cannot be installed on ESS IO nodes or ESS Protocol Nodes due to the HBA requirement for tape hardware connectivity.

IBM Spectrum Archive EE on an IBM Spectrum Scale Server

IBM Spectrum Archive EE can be installed on IBM Spectrum Scale NSD servers. Evaluate workload and availability needs on the combined IBM Spectrum Scale NSD and IBM Spectrum Archive EE to ensure resource requirements are sufficient.

Security-Enhanced Linux (SELinux)

The SELinux setting must be in permissive mode or disabled when deploying with IBM Spectrum Scale 5.0.4 or earlier. Starting from IBM Spectrum Scale 5.0.5, SELinux also can be set to enforcing mode.

Remote Mount

IBM Spectrum Archive EE handles only file systems that belong to the local (home) IBM Spectrum Scale cluster, but not file systems that are installed remotely.

3.4 Required software for Linux systems

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 v1.3.2.2.

3.4.1 Required software packages for Red Hat Enterprise Linux systems

The following software modules are required to deploy IBM Spectrum Archive Enterprise Edition on Red Hat Enterprise Linux systems:

- ▶ The most current fix pack for IBM Spectrum Scale: 5.0.2, 5.0.3, 5.0.4, 5.0.5, 5.1.0, 5.1.1, 5.1.2, or subsequent releases. The following operating system software:
 - attr
 - Boost.Date_Time (boost-date-time)
 - Boost.Filesystem (boost-filesystem)
 - Boost.Program_options (boost-program-options)
 - Boost.Serialization (boost-serialization)
 - Boost.Thread (boost-thread)
 - boost_regex
 - FUSE
 - fuse
 - fuse-libs
 - gperftools-libs
 - Java virtual machine (JVM)
 - libxml2
 - libuuid
 - libicu
 - lsof
 - net-snmp
 - nss-softokn-freebl
 - openssl
 - Python 2.4 or later, but earlier than 3.0
 - python3-pyxattr (required with Red Hat Enterprise Linux 8.x. Requires access to the Red Hat CodeReady Linux Builder repository)
 - pyxattr (required with Red Hat Enterprise Linux 7.x)
 - rpcbind
 - sqlite
- ▶ HBA device driver (if one is provided by the HBA manufacturer) for the host bus adapter connecting to the tape library and attach to the tape drives.

Note: Java virtual machine (JVM) 1.7 or later must be installed before the “Extracting binary rpm files from an installation package” step for IBM Spectrum Archive on a Linux system during the installation of IBM Spectrum Archive Enterprise Edition.

3.4.2 Required software to support REST API service on RHEL systems

The optional REST API service for IBM Spectrum Archive Enterprise Edition supports RHEL systems. The following software for the REST API support must be installed on the RHEL 7.x system:

- ▶ httpd
- ▶ mod_ssl
- ▶ mod_wsgi
- ▶ Flask 0.12

Install Flask by using one of the following methods:

- Use pip: `pip install Flask==0.12`
- [Download from pypi](#)

Note: The REST API must be installed on one of the IBM Spectrum Archive Enterprise Edition nodes.

3.4.3 Required software to support a dashboard for IBM Spectrum Archive Enterprise Edition

IBM Spectrum Archive Enterprise Edition supports a dashboard monitor system performance, statistics, and configuration. The following software is included in the IBM Spectrum Archive EE installation package:

- ▶ Logstash 5.6.8, to collect data
- ▶ Elasticsearch 5.6.8, to store the data
- ▶ Grafana 5.0.4-1, to visualize data

Note: Consider the following points:

- ▶ The dashboard requires EE node and monitoring server to run on x86_64 platform. For more information about other requirements, see this [IBM Support web page](#).
- ▶ Support for the open source packages can be acquired for a fee by contacting a third-party provider. It is *not* covered by the IBM Spectrum Archive Enterprise Edition license and support contract.

3.4.4 Required software for SwiftHLM

IBM Spectrum Archive Enterprise Edition can use SwiftHLM functions for suitable integration with Swift. The Swift High Latency Media (HLM) project can create a high-latency storage backend that makes it easier for users to perform bulk operations of data tiering within a Swift data ring.

SwiftHLM 0.2.1 is required for the optional use of SwiftHLM.

3.5 Hardware and software setup

Valid combinations of IBM Spectrum Archive EE components in an IBM Spectrum Scale cluster are listed 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 Management (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 and between the host and the disk storage to create a storage area network (SAN) to further expand storage needs as required.

3.6 IBM Spectrum Archive deployment examples

This section describes some examples of IBM Spectrum Archive deployment on Lenovo servers and on a Versastack converged infrastructure.

3.6.1 Deploying on Lenovo servers

This section describes IBM Spectrum Archive deployment examples with Lenovo servers.

Minimum Lenovo server deployment

Figure 3-10 shows the minimum deployment option as shown in Figure 3-1 on page 41. It is a high level diagram of a minimum deployment where all required software components of IBM Spectrum Scale and IBM Spectrum Archive EE are installed on a Lenovo ThinkSystem SR650 server with a direct attached IBM TS4300 using Fibre-Channel connections.

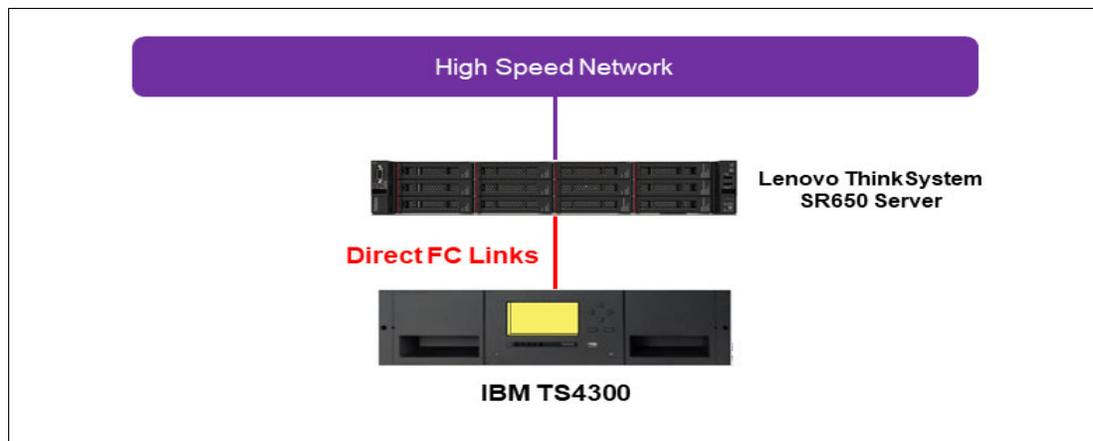


Figure 3-10 Minimum deployment on Lenovo ThinkSystem SR650 server with IBM TS4300

The following example server configuration is used for a Lenovo ThinkSystem SR650 rack server:

- ▶ Forty cores (Two sockets each with 20-cores) Intel Xeon Processors
- ▶ 512 GB RAM
- ▶ RAID controller
- ▶ Nine units 10 TB 3.5-inch 7.2KRPM NL SAS (RAID 6 = 6D+P+Q+Spare)
- ▶ Five units 3.84 TB 3.5-inch SATA SSD (RAID 5 = 3D+P+Spare)
- ▶ Two units 480 GB SATA Non-Hot Swap SSD for operating system
- ▶ Red Hat Enterprise Linux with Lenovo Support for Virtual Data Centers x2 socket licenses
- ▶ One unit Dual-port 16 Gbit FC HBA for tape interface
- ▶ One unit Dual-port 10/25 GbE SFP28 PCIe Ethernet adapter for network interface
- ▶ One available PCIe slot might be configured as other FC or GbE connectivity

After RAID protection of the internal flash and disk resources, the estimated usable capacity is approximately 55 TiB for the nearline SAS disks and 10 TiB for the SSDs for a total of 65 TiB. A portion of the SSD capacity can be allocated for the IBM Spectrum Scale metadata and the rest can be deployed as pools to the IBM Spectrum Scale file system.

IBM Spectrum Scale can be licensed on a per-terabyte basis with the following three license options:

- ▶ Data Access Edition
- ▶ Data Management Edition
- ▶ Erasure Code Edition.

The storage capacity to be licensed is the capacity in terabytes (TiB) from the network shared disk (NSD) in the IBM Spectrum Scale Cluster. This means that less capacity can be allocated to IBM Spectrum Scale if the required capacity of the storage tier required is less than 65 TiB, thereby reducing the IBM Spectrum Scale capacity license.

For more information about IBM Spectrum Scale licenses and licensing models, see the following IBM Documentation web pages:

- ▶ [IBM Spectrum Scale license designation](#)
- ▶ [Capacity-based licensing](#)

The IBM TS4300 tape library is configured with two units of half-high LTO-8 tape drives, which directly connect to the Lenovo ThinkSystem SR650 server using 8Gbit Fibre-Channel connections.

IBM Spectrum Archive EE is licensed per node regardless of the capacity migrated to tapes. With node-based licensing, storing data on tapes is virtually unlimited provided if enough tape resources are available in the tape-storage tier. This node-based licensing can help realize significant savings on storage and operational costs.

In this example, we need only one IBM Spectrum Archive EE license, which is for one node that is installed with IBM Spectrum Scale. If more tape access channels are needed, more IBM Spectrum Scale with IBM Spectrum Archive EE nodes can be added to the cluster.

High-availability minimum deployment architecture on Lenovo servers

Figure 3-11 on page 59 shows the injection of high availability with the minimum deployment architecture as shown in Figure 3-2 on page 42.

Figure 3-11 on page 59 shows the high-level architecture of a high-availability minimum solution using Lenovo ThinkSystem SR650 servers. Both servers will be installed with the same software stack as that of the minimum deployment solution described in “Minimum Lenovo server deployment” on page 57.

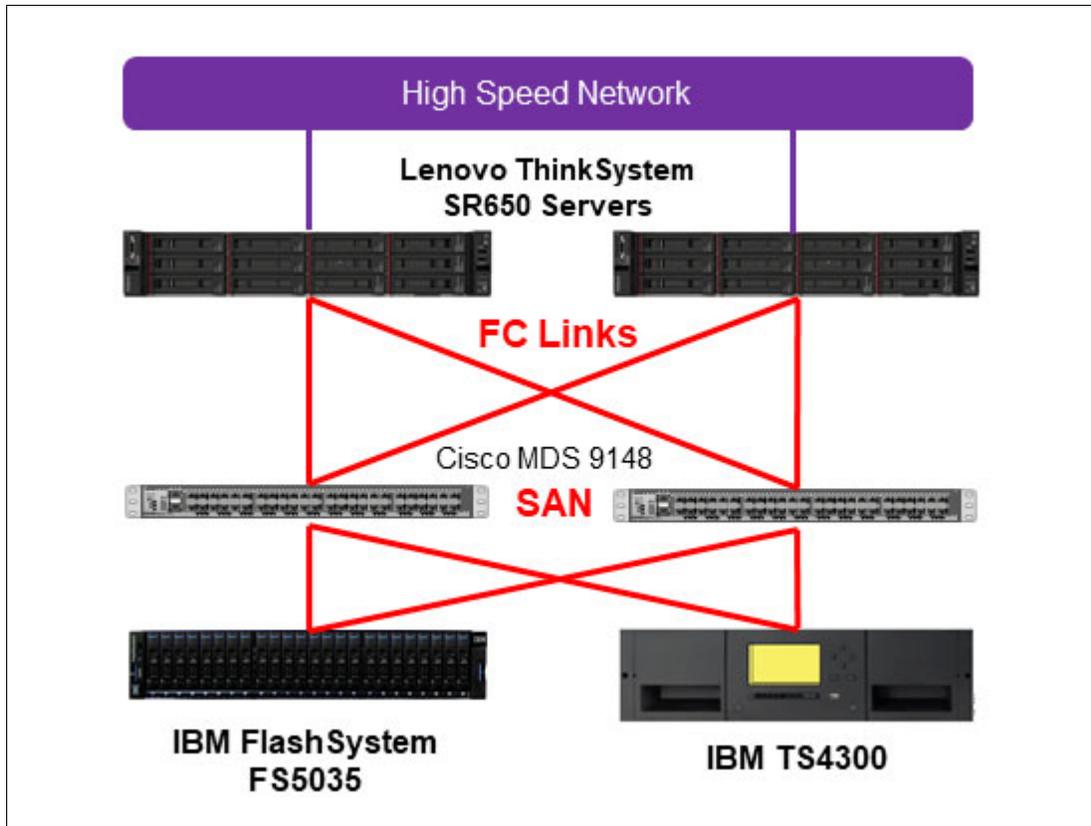


Figure 3-11 High availability on minimum deployment IBM Spectrum Scale

Figure 3-11 shows the following high level components:

- ▶ Lenovo ThinkSystem SR650 servers
- ▶ IBM FlashSystems 5035
- ▶ IBM Storage Networking SAN48C-6 SAN switches
- ▶ IBM TS4300 tape library with LTO 8 tape drives

The difference of this architecture is that instead of internal storage for the IBM Spectrum Scale with IBM Spectrum Archive EE servers, storage capacity will now be on a shared storage system using redundant SAN Switches. In Figure 3-11, the IBM FlashSystem® FS5035 is deployed with a pair of IBM SAN48-C SAN switches for this example.

The IBM Spectrum Scale license for this architecture can be either Data Access Edition or Data Management Edition. Erasure Code Edition is not applicable because of the shared storage architecture. In Figure 3-11, both servers with IBM Spectrum Archive EE need to be licensed.

3.6.2 Deploying on Versastack converged infrastructure

Another example of deployment is to use the deployment of IBM Spectrum Archive EE system on Versastack. VersaStack is a converged infrastructure solution of network, compute and storage designed for quick deployment and rapid time to value. The solution includes Cisco Unified Computing System (UCS) integrated infrastructure together with IBM software-defined storage solutions to deliver extraordinary levels of agility and efficiency. VersaStack is backed by Cisco Validated Designs and IBM Redbooks application guides for faster infrastructure delivery and workload/application deployment.

VersaStack integrates network, compute, and storage, as follows:

- ▶ On the network and compute side, VersaStack solutions use the power of Cisco UCS integrated infrastructure. This infrastructure includes the cutting-edge Cisco UCS and the consolidated Cisco Intersight delivering simple, integrated management, orchestration and workload assurance.
- ▶ On the storage side, IBM offers storage solutions for VersaStack based on IBM Spectrum Virtualize technologies.

To set up IBM Spectrum Archive EE on the Versastack-converged infrastructure, deploy the required software components on properly-sized Cisco UCS blade servers in the UCS chassis and integrate an IBM TS4500 tape library using Cisco Multilayer Director Switch (MDS) SAN switches.

The high-level components that formulate this architecture are as follows:

- ▶ Cisco UCS 5108 Chassis
- ▶ Cisco UCS B200 M5 Servers as sized and configured respectively with IBM Spectrum Scale and IBM Spectrum Archive EE requirements
- ▶ Cisco UCS 2408 Fabric Extenders
- ▶ Cisco UCS 6454 Fabric Interconnect
- ▶ Cisco Nexus 93180 Top of Rack switches
- ▶ IBM FlashSystems 5035
- ▶ IBM Storage Networking SAN48C-6
- ▶ IBM TS4500 tape library with LTO 8 tape drives

Figure 3-12 shows a high-level overview of an IBM Spectrum Archive EE architecture deployed on Versastack.

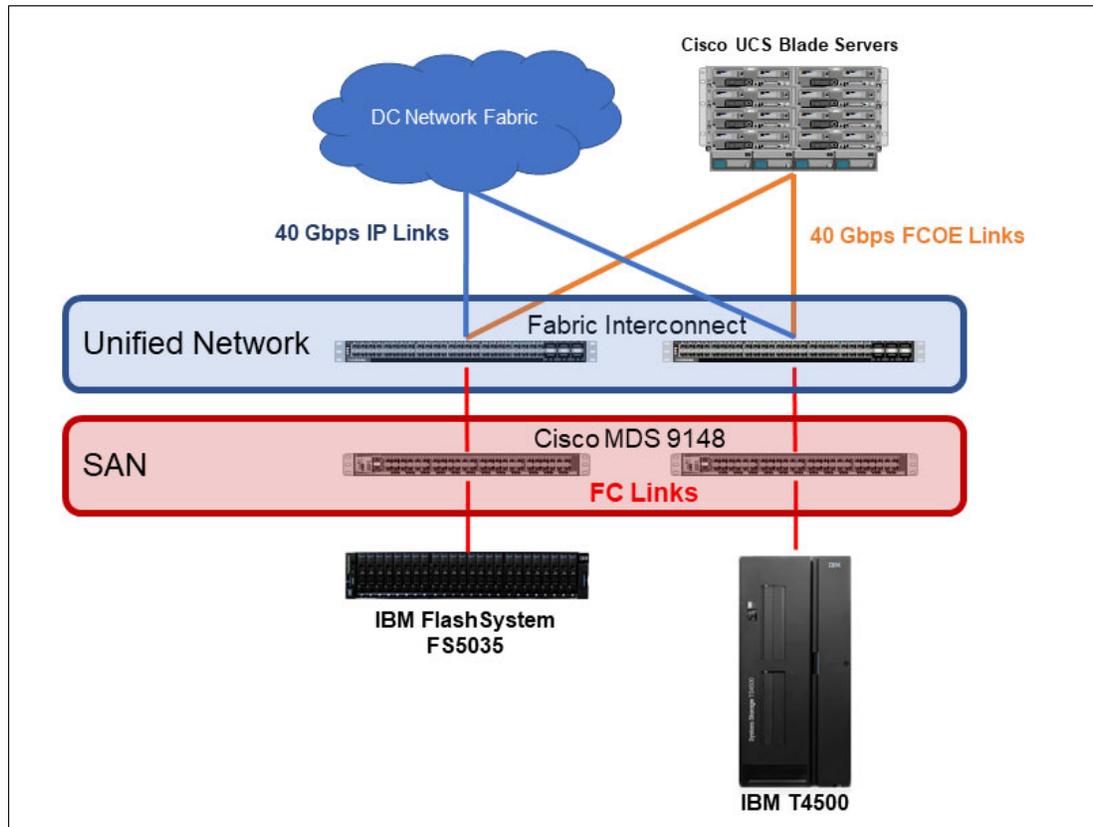


Figure 3-12 IBM Spectrum Archive EE high level architecture on Versastack

Note: You must collaborate with your in-country Cisco Data Center Product Specialists for the sizing of the compute, network, and services requirements.

3.7 Sizing and settings

Several items must be considered 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. This section describes the IBM Spectrum Scale file-system aspects to help avoid the need to make changes later. This section also describes the IBM Spectrum Archive EE Sizing Tool that aids in the design of the solution architecture.

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. Therefore, it is suggested that you use 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 Storage.

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

3.7.1 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 6.11.4, “Replicas and redundant copies” on page 177.

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

3.7.2 Planning for LTO-9 Media Initialization/Optimization

The deployment of LTO-9 tape drives will require a one-time media optimization for all new LTO-9 tape cartridges. Because the process of media optimization can take an average of 40 minutes up to two hours per LTO-9 tape cartridge, refer to the suggested strategies to mitigate disruptions to normal tape operations:

- ▶ Small tape library deployments are recommended to perform media optimization to all media before beginning normal operations.
- ▶ Large tape library deployments are recommended to use as many drives to perform media optimization as possible. Then dedicate one drive for every seventeen drives to continuously load new media for optimization during normal operations.

For more information about LTO-9 media initialization and optimization, see 7.29, “LTO 9 Media Optimization” on page 272, and [this Ultrium LTO web page](#).

Note: The media optimization process cannot be interrupted; otherwise, the process must restart from the beginning.

3.7.3 IBM Spectrum Archive EE Sizing Tool

Sizing the IBM Spectrum Archive EE system requires consideration of workload and data-access aspects to deploy the proper system resources. For this section, download the latest IBM Spectrum Archive EE Sizing Tool package, which is available at this [IBM Support web page](#).

Before simulating configurations with the IBM Spectrum Archive EE Sizing Calculator, the workload data-characteristics first must be studied.

Understanding the workload data file sizes is key to the designing of the IBM Spectrum Scale file system. The file size of most of the data determines the suitable Data and Metadata block sizes during initial IBM Spectrum Scale deployment for new systems.

For more information about Data and Metadata, see [IBM Spectrum Scale Version 5.1.2 Concepts, Planning, and Installation Guide](#).

For small to medium implementations like the architecture depicted in figure 3.1.1, “On IBM Spectrum Scale Servers” on page 40, it is recommended to consider flash storage for the System Pool (which is where the Metadata is stored in the file system) for performance. More Storage Pools can be designed by using flash or spinning disks depending on data tiering requirements.

Note: Consider the following points:

- ▶ Both Data and Metadata block sizes are defined during IBM Spectrum Scale File System creation and CAN'T be changed after.
- ▶ For workloads with extremely small file sizes in the kilobytes range, it is highly recommended to consolidate these files into zip or tar files prior to tape migration.

IBM Spectrum Archive EE Sizing Tool INPUT

The IBM Spectrum Archive EE Sizing Calculator is an intuitive tool that provides guidance on data entry and output explanations. The following data input sections are included:

- ▶ **Data amount and access patterns:** This section considers the average data size of the workload files and ingest and recall estimates. The minimum data size input is 1MB.
- ▶ **Archive Retention Requirements:** This section requires the data retention duration on disk (also known as Cache) before migration or premigration to the tape tier.
- ▶ **Archive Capacity Requirements:** This section will ask for the total projected data archive size on tapes and file replicas on tape.
- ▶ **Tape Technology Selection:** This section allows the user to input the target tape infrastructure to be deployed and will provide estimations on the required tape performance, quantity of tape drives and cartridges, and the number of IBM Spectrum Archive EE nodes.

IBM Spectrum Archive EE Sizing Tool OUTPUT

The tool generates the target configuration components that accommodate the data flow and retention input by the user:

- ▶ **IBM Spectrum Scale Licenses:** IBM Spectrum Scale is a prerequisite for IBM Spectrum Archive EE solutions. As described in the previous sections on deployment options, there are many approaches to setting up the IBM Spectrum Scale environment. Capacity-based licenses simplify this solution component, that is, configure the license with respect to the Disk Storage output of the IBM Spectrum Archive EE Sizing Calculator. However, it is recommended to add TB licenses as a buffer.
- ▶ **IBM Spectrum Archive EE Licenses;** Based on tape-data access throughput, this determines the required IBM Spectrum Archive EE licenses. IBM Spectrum Archive EE is licensed on a per-node basis. If the solution requires redundancy, a minimum of two nodes licenses are required.
- ▶ **Disk Storage:** The output for disk storage specifies the usable disk capacity requirements. It does not indicate the storage type (flash or spinning disks). However, it is recommended to deploy flash storage for Metadata. The stub files are also recommended to be in flash storage but might also be on spinning disks depending on the access patterns of the users. An external-shared storage system is recommended (but not required) for easy scaling and high availability for installations with more than one server.
- ▶ **External storage servers** might deploy different storage technologies such as Flash, SAS, and Nearline SAS disks that set the stage for the IBM Spectrum Scale Storage Pools. Capacity-based licensing of IBM Spectrum Scale also makes it easier to design the system, that is, license only the required disk capacity to be allocated to the IBM Spectrum Scale file system. The performance of shared external-storage systems can also be tuned by deploying more flash or disk spindles.
- ▶ **Tape Storage:** This output section shows the average data transfer tape to and from tape and the quantity of tape drives and tape cartridges with respect to the target total tape-archive capacity.
- ▶ **IBM Spectrum Archive EE server minimum configuration:** This output section indicates the minimum configuration for the IBM Spectrum Archive EE server. For optimal performance, it is recommended that each IBM Spectrum Archive EE server have at least 128GB of memory even if the sizing tool indicates any less. This section also shows the required CPU Socket requirements and Fibre Channel Host Bus Adapters or Disk Connection Adapters (for IBM ESS) architectures.

Note: For converged solution architectures where IBM Spectrum Scale and IBM Spectrum Archive EE are installed in one physical server, it is recommended to assign at least 128GB of memory for the IBM Spectrum Archive EE and respectively add the IBM Spectrum Scale requirements that will be deployed on the same server. Apply the same for CPU and connectivity requirements.

3.7.4 Performance

Performance planning is an important aspect of 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 [IBM Spectrum Archive Enterprise Edition v1.2.2 Performance](#).

The configuration that is shown in Figure 3-13 was used to run the lab performance test in this section.

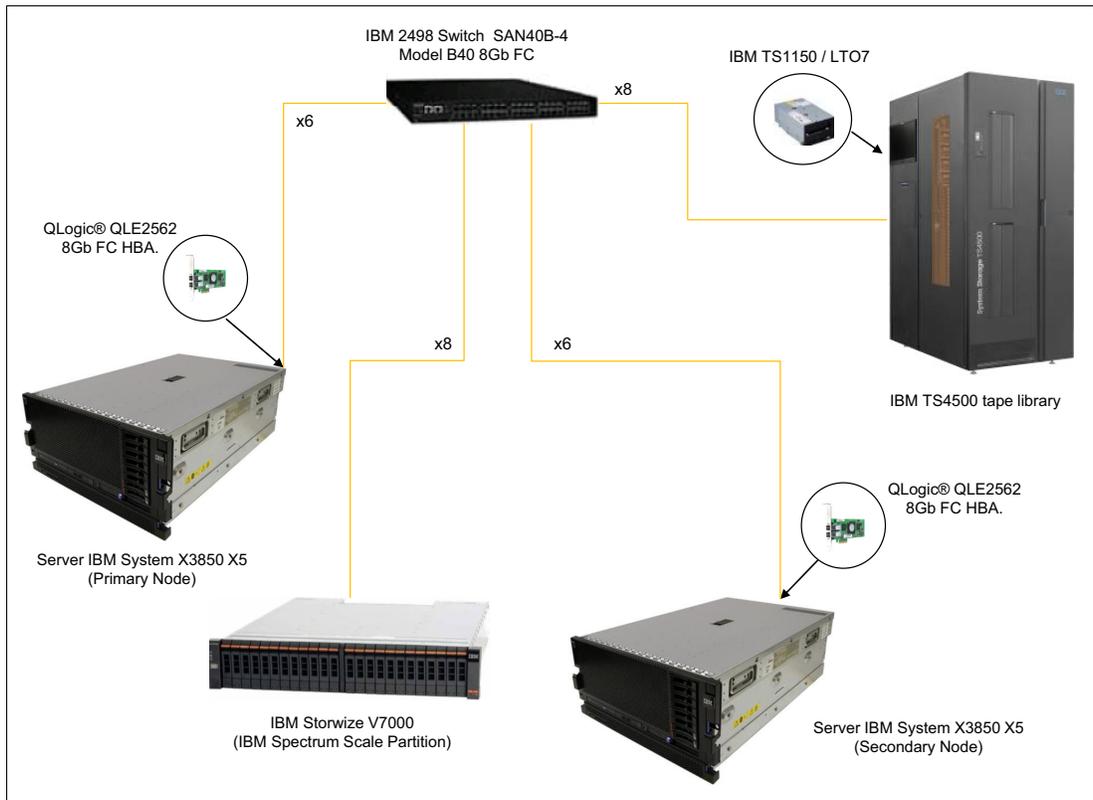


Figure 3-13 IBM Spectrum Archive EE configuration used in lab performance

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 3-13 shows the example configuration. This configuration is only one possibility; yours might be different.

Figure 3-14 shows a basic internal configuration of Figure 3-13 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.

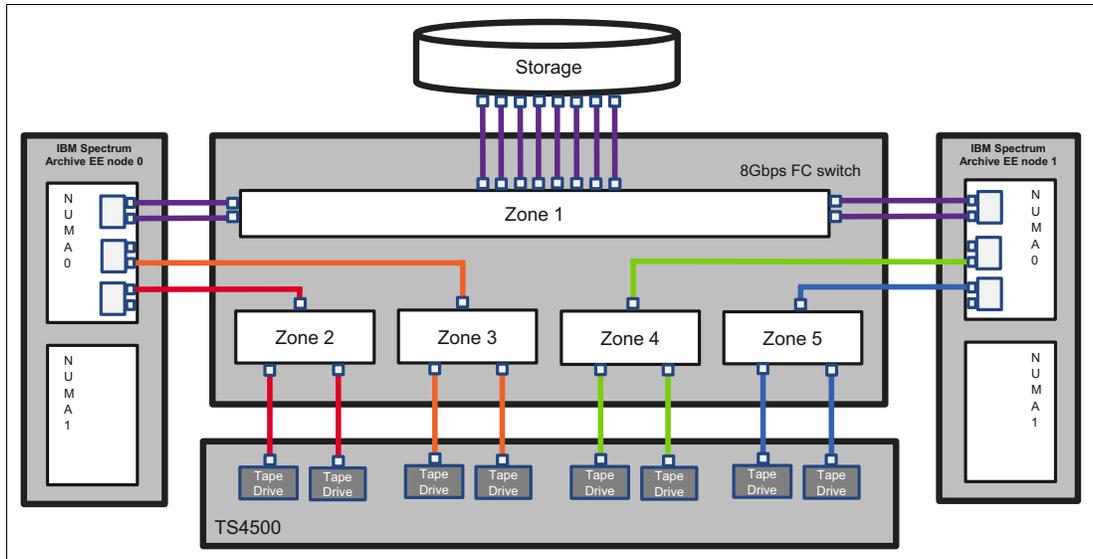


Figure 3-14 Internal diagram of two nodes running off NUMA 0

Note: Figure 3-14 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.

Figure 3-14 is one of many possible configurations and can be used as a guide for how you set up your own environment for best performances. For example, if you have more drives, you can add tape drives per zone, or add Host bus adapters (HBAs) on each server and create a zone.

Table 3-3 and Table 3-4 list 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 lists 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 3-15 shows a comparison line graph of the raw performance data obtain for TS1150 drives and LTO7 drives that use the same configurations.

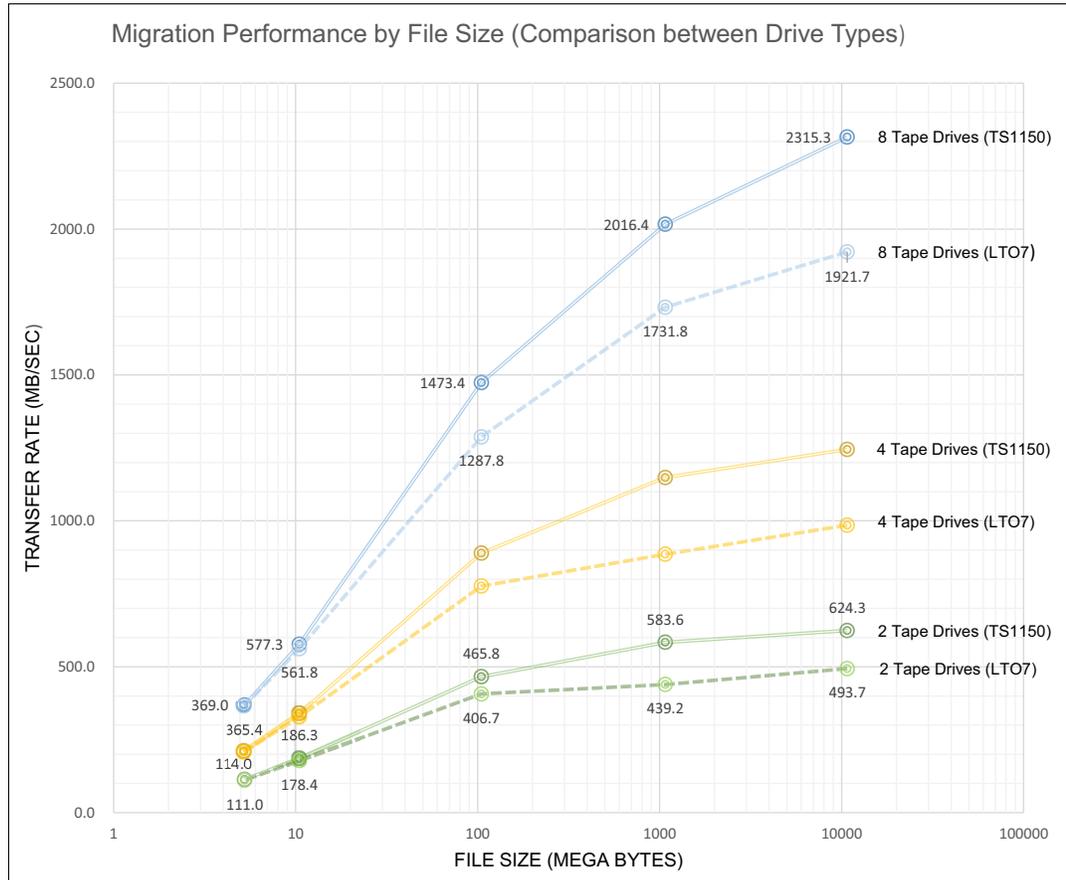


Figure 3-15 Comparison between TS1150 and LTO-7 drives using multiple node/drive configurations

It is clear that by using small files, the comparison between the two types of drives is minimal. However, when migrating file sizes of 1 GiB and larger, a noticeable difference exists. Comparing the biggest configuration of 2N4D8T, the LTO peaks at a total combined transfer rate of 1921.7 MiBps. With the same configuration but with TS1150 drives, it peaks at a total combined transfer rate of 2315.3 MiBps.

3.7.5 Ports that are used by IBM Spectrum Archive EE

In addition to SSH, the software components of IBM Spectrum Archive Enterprise Edition use several TCP/UDP ports for interprocess communication within and among nodes.

The used ports are listed in Table 3-5.

Table 3-5 Ports used by IBM Spectrum Archive EE

Used port number	Required node	Component using the ports
7600	All IBM Spectrum Archive EE nodes	LE
7610	All IBM Spectrum Archive EE nodes	MD
The range from 7620 to (7630 + number of EE nodes)	The IBM Spectrum Archive EE active control node	MMM
RPC bind daemon port (for example, 111)	The IBM Spectrum Archive EE active control node	MMM

Several other ports are used by HSM. For more information, see this [IBM Documentation web page](#).

3.8 High-level component upgrade steps

This section describes some high-level examples of IBM Spectrum Archive EE component upgrades as the need arises (see Table 3-6). Upgrading IBM Spectrum Archive EE components is easy and flexible and addresses scalability and technology refresh requirements.

Table 3-6 High level component upgrade options

Hardware upgrade	Upgrade procedure
Node server	<p>IBM Spectrum Scale Node Follow IBM Spectrum Scale procedures for node add and node delete. For more information, see:</p> <ul style="list-style-type: none"> – Node add – Node delete <p>IBM Spectrum Archive EE node Use IBM Spectrum Archive <code>ltfsee_config</code> command with remove_node and add_node options. Refer to 5.2.1, “The <code>ltfsee_config</code> utility” on page 105.</p>
IBM Spectrum Scale disk storage (internal and external)	<p>Method 1. Preserve GPFS cluster and GPFS file systems Follow IBM Spectrum Scale procedure for NSD add, NSD restripe and NSD delete. For more information, see:</p> <ul style="list-style-type: none"> – NSD add – NSD restripe – NSD delete <p>Method 2. Create a GPFS cluster or new GPFS file systems Migrate the old file system to the new file system using IBM Spectrum Archive procedure “system migration using SOBAR”. Refer to 2.2.6, “Scale Out Backup and Restore” on page 32, 6.27, “File system migration” on page 231, and IBM Documentation Scale Out Backup and Restore.</p>

Hardware upgrade	Upgrade procedure
Tape drives/cartridges technology upgrade	<p>Method 1. IBM Spectrum Archive EE eeadm pool set command to modify the media_restriction attribute Use IBM Spectrum Archive media_restriction attribute of a pool to gradually move data to the specified newer tape technology within a pool.</p> <p>Method 2. IBM Spectrum Archive EE eeadm tape datamigrate command Use IBM Spectrum Archive eeadm tape datamigrate command to move data from an old pool with old tape generation resources to a new pool with new tape generation resources.</p> <p>For more information, see 6.11.5, "Data Migration" on page 180 for details.</p>
Tape library	For more information, see 4.6, "Library replacement" on page 88.

Note: Component upgrades must be reviewed with IBM.



Installing IBM Spectrum Archive Enterprise Edition

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.
- ▶ 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:

- ▶ 4.1, “Installing IBM Spectrum Archive EE on a Linux system” on page 74
- ▶ 4.2, “Installation prerequisites for IBM Spectrum Archive EE” on page 74
- ▶ 4.3, “Installing IBM Spectrum Archive EE” on page 75
- ▶ 4.4, “Installing a RESTful server” on page 83
- ▶ 4.5, “Quick installation guide for IBM Spectrum Archive EE” on page 88
- ▶ 4.6, “Library replacement” on page 88
- ▶ 4.7, “Tips when upgrading host operating system” on page 95

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 74)

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 75)

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 for IBM Spectrum Archive Enterprise Edition” on page 39.
- ▶ 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 for IBM Spectrum Archive Enterprise Edition” on page 39.
- ▶ 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 10.4, “IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products” on page 330.
- ▶ Ensure that all prerequisite software is installed, as described in 4.2.1, “Installing the host bus adapter and device driver” on page 75.
- ▶ 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 75.

- ▶ Determine the distribution package for IBM Spectrum Archive EE that is required for your system.
- ▶ For IBM Spectrum Scale prerequisites, see 3.5, “Hardware and software setup” on page 56.

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 IBM Support’s [Fix Central web page](#).

For more information about HBA interoperability, see this [IBM System Storage Interoperation Center \(SSIC\) web page](#).

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 points:

- ▶ 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 Documentation.
- ▶ The IBM Spectrum Archive EE license does not entitle customers to use any other IBM IBM Spectrum Protect 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 10.4, “IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products” on page 330.

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-[version]-[buildlevel].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-[version]-[buildlevel].bin` by running the following command:

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

```
# chmod +x 1tfsee-[version]-[buildlevel].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 79 or 4.5, “Quick installation guide for IBM Spectrum Archive EE” on page 88 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 installation package 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-[version]-[buildlevel].bin
```
 - If your operating system is running on the GUI (X Window System), run the following command:

```
#./1tfsee-[version]-[buildlevel].bin -i console
```

The messages that are shown in Example 4-1 are displayed.

Example 4-1 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.[version]_[buildlevel]` folder; for example, the `~/LTFSEE/rpm.1322_52823/` subfolder.

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

Example 4-2 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 4-3. You can monitor the progress by watching the text-animated progress bars.

Example 4-3 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 `~/LTFSEE/rpm.[version]_[buildlevel]` to find the rpms. 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.[version]_[buildlevel]` and `README_EE.[version]_[buildlevel]`, 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 79 to complete the installation. If you prefer to install manually, see 4.5, “Quick installation guide for IBM Spectrum Archive EE” on page 88.

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 76.

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 and 8, managing system services are done by using 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.

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/share/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.[version]_[buildlevel]` directory with the IBM Spectrum Archive EE rpm files (such as `/root/LTFSEE/rpm.1322_52823/`).

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 79.

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

- ▶ IBM Spectrum Archive LE component
- ▶ Integrated customized IBM Spectrum Protect 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 output:

The prerequisites checking is completed successfully.

Example 4-4 shows the complete output.

Example 4-4 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 88, except for installing optional TIVsm language packages (if 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 4-5 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 4-5 Output for the ltfsee_install --install command

```
[root@ltfsServer rpm.1320_52814]# ./ltfsee_install --install  
Checking the software prerequisites on ltfsServer.tuc.stglabs.ibm.com.
```

No error found during the prerequisite check.

Starting RPM installation.

Installing HSM modules.
 Created symlink from
 /etc/systemd/system/multi-user.target.wants/hsm.service to
 /usr/lib/systemd/system/hsm.service.
 Installing EE modules.
 Starting HSM.
 Completed RPM installation.
 Running the post-installation steps.
 All rpm packages are installed successfully.
 The installed IBM Spectrum Archive EE version: 1.3.2.0_52814.
 Complete the configuration by using the /opt/ibm/ltfsee/bin/ltfsee_config
 command.

**** ATTENTION ****

For problem determination, it is strongly recommended that you disable
 log suppression and set up the abrtd daemon to capture the Spectrum
 Archive core dumps.

Refer to the Troubleshooting section of the IBM Spectrum Archive
 Enterprise Edition documentation in the IBM Knowledge Center.

-
- iii. Complete the configuration by running the `/opt/ibm/ltfsee/bin/ltfsee_config` command, as described in 5.2, “Configuring IBM Spectrum Archive EE” on page 105.

- b. Upgrading the rpm files to the latest versions:

Note: In a multi-node environment, run the `ltfsee_install --upgrade --all` command on one node to automatically update all nodes in the cluster. Running the `ltfsee_install --upgrade` command on each node is *not* recommended because doing so might cause unexpected failures during a multi-node system upgrade.

- i. Run the `eadm cluster stop` command.
- ii. Run the `pidof mmm` command on all active control nodes and wait until there are no processes returned.
- iii. Run the `pidof ltf` command on every EE node and wait until there are no processes returned.

To perform the upgrade process, run the `ltfsee_install --upgrade` command or the `ltfsee_install --upgrade --all` command for multi node environment. An example for a single node upgrade is shown in Example 4-6 on page 81.

Example 4-6 Running ltfsee_install --upgrade command

```
[root@ltfsServer rpm.1320_52814]# ./ltfsee_install --upgrade
The upgrade installation option is selected.
This option upgrades the IBM Spectrum Archive EE software
  From: Version 1.3.1.0-xxxxx
  To:   Version 1.3.2.0-52814
Do you want to continue? [Y/n]: Y
Preparing for the software upgrade.
Checking the software prerequisites on ltfsServer.tuc.stglabs.ibm.com.
No error found during the prerequisite check.

**** Upgrading the software on ltfsServer.tuc.stglabs.ibm.com. ****
Uninstalling RPMs.
```

```
Removed symlink /etc/systemd/system/multi-user.target.wants/hsm.service.  
Starting RPM installation.  
Installing HSM modules.  
Created symlink from  
/etc/systemd/system/multi-user.target.wants/hsm.service to  
/usr/lib/systemd/system/hsm.service.  
Installing EE modules.  
Starting HSM.  
Completed RPM installation.  
Running the post-installation steps.  
All rpm packages are upgraded successfully on the local node,  
ltfsServer.tuc.stglabs.ibm.com.
```

The following files were updated during the installation. (The original files were copied to the /opt/ibm/ltfsee/bin directory)

- /etc/rsyslog.d/ibmsa-rsyslog.conf
- /etc/logrotate.d/ibmsa-logrotate

**** ATTENTION for upgrade ****

If you have installed the IBM Spectrum Archive EE dashboard, restart the logstash on all of the IBM Spectrum Archive EE nodes.

**** ATTENTION ****

For problem determination, it is strongly recommended that you disable log suppression and set up the abrtd daemon to capture the Spectrum Archive core dumps.

Refer to the Troubleshooting section of the IBM Spectrum Archive Enterprise Edition documentation in the IBM Knowledge Center.

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 (see Example 4-7) was successful, the following message is shown:

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

Example 4-7 Successful ltfsee_install --verify command

```
[root@ltfsServer rpm.1320_52814]# ./ltfsee_install --verify  
Checking the software prerequisites on  
ltfsServer.ltfServer.tuc.stglabs.ibm.com.
```

```
No error found during the prerequisite check.
```

```
An IBM Spectrum Archive EE package ltfs-license is installed.
```

An IBM Spectrum Archive EE package `ltfsle-library` is installed.
An IBM Spectrum Archive EE package `ltfsle-library-plus` is installed.
An IBM Spectrum Archive EE package `ltfsle` is installed.
An IBM Spectrum Archive EE package `ltfs-admin-utils` is installed.
An IBM Spectrum Archive EE package `gskcrypt64` is installed.
An IBM Spectrum Archive EE package `gskssl64` is installed.
An IBM Spectrum Archive EE package `TIVsm-API64` is installed.
An IBM Spectrum Archive EE package `TIVsm-BA` is installed.
An IBM Spectrum Archive EE package `TIVsm-HSM` is installed.
An IBM Spectrum Archive EE package `ltfs-mig` is installed.
All IBM Spectrum Archive EE packages are installed.

The module installation check 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 76 and installing IBM Spectrum Archive EE as described in 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 78. This section must 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 software applications are required:

- ▶ IBM Spectrum Archive EE v1.2.4 or later
- ▶ `httpd`
- ▶ `mod_ssl`
- ▶ `mod_wsgi`
- ▶ Python 2.4 or later, but earlier than 3.0
- ▶ Flask

Example 4-8 shows how to perform the installation of the required software to run REST.

Example 4-8 Required software for rest installation

```
[root@ltfseesrv1 ~]# 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 version 0.12.2:

```

pip install flask==0.12.2

```

Example 4-9 shows how to install flask.

Example 4-9 Install flask V0.12

```

[root@l1fseesrv1 ~]# pip install flask==0.12.2
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-[version]-[buildlevel].x86_64.rpm`. To install the restful service, run the `yum install` command on this file, as shown in Example 4-10.

Example 4-10 Installing IBM Spectrum Archive Rest service

```

[root@ltfseesrv1 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
```

Note: If REST is already installed on cluster updating IBM Spectrum Archive EE will automatically update the REST interface, however a manual restart of httpd is required after starting IBM Spectrum Archive EE.

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 4-11 shows using a **test curl** command to see whether the installation was successful.

Example 4-11 Test curl command

```

[root@l1tfseesrv1 ~]# 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 6.26, “IBM Spectrum Archive REST API” on page 219.

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

```
./l1tfsee-1.3.2.2_52823.bin
```

4. Install IBM Spectrum Archive EE automatically by using the `l1tfsee_install` tool. Use `--check` for preinstallation check, `--install` for the installation, and `--verify` for a postinstallation verification, as shown in the following commands:

```
- ~/LTFSEE/rpm.[version]_[buildlevel]/l1tfsee_install --check
- ~/LTFSEE/rpm.[version]_[buildlevel]/l1tfsee_install --install
- ~/LTFSEE/rpm.[version]_[buildlevel]/l1tfsee_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 6.11.5, “Data Migration” on page 180.

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 must 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 TS3500 tape library with a new TS4500 tape library.

To do so, complete the following steps:

1. Identify the library name and serial number that you intend to replace, for use in the subsequent steps. Example 4-12 shows how to get a list of libraries and serial numbers by running the `ltfsee_config -m LIST_LIBRARIES` command.

Example 4-12 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 `eeadm drive unassign` command. Refer to “IBM Spectrum Scale commands” on page 323 for more information about the command.
3. Stop IBM Spectrum Archive EE by executing the `eeadm cluster 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 4-13 shows the output from running the `ltfsee_config -m REPLACE_LIBRARY` command to associate the new library with the old library serial number.

Example 4-13 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 `eeadm cluster start` command.
10. If you are using new drives in the new library, configure the drives with the `eeadm drive assign` command. See “The `eeadm <resource type> --help` command” on page 316 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 `eeadm pool list` command.
 2. Stop IBM Spectrum Archive EE.

3. Prepare the selected pool for relocation with the `ltfsee_config -m PREPARE_MOVE_POOL` command. Example 4-14 shows the output from running the command on pool1 from library lib1 to library lib2.

Example 4-14 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 4-15.

Example 4-15 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 **eadm cluster 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 **eadm tape list** command.
- Move tapes within the selected pool to the IE slot to remove from the previous library using the **eadm tape move <tape1[tape2 tape3 ...] -L ieslot -p <pool> -l <library>** command.
- Stop IBM Spectrum Archive.
- Activate the pool for relocation by running the **ltfsee_config -m ACTIVATE_MOVE_POOL** command. Example 4-16 illustrates activating the relocated pool after all tapes have been physically moved to the new tape library.

Example 4-16 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 4-17.

Example 4-17 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 **eeadm cluster start** command.

12. Verify that the pool and tapes are valid, as shown in Example 4-18.

Example 4-18 Verify pool and tapes in new tape library

```
[root@server1 ~]# eeadm pool list
```

Pool Name	Usable(TiB)	Used(TiB)	Available(TiB)	Reclaimable%	Tapes	Type
pool1	2.6	1.3	1.2	1%	4	LTO
lib2	GO					
pool2	1.3	0.9	0.4	0%	4	LTO
lib2	GO					


```
[root@server1 ~]# eeadm tape list
```

Tape ID	Status	State	Usable(GiB)	Used(GiB)	Available(GiB)
Reclaimable%	Pool	Library	Location	Task ID	
D00369L5	ok	appendable	1327	835	491
1%	pool1	lib2	homeslot	-	
2MA133L5	ok	appendable	1327	763	564
2%	pool1	lib2	homeslot	-	
2FC181L5	ok	appendable	1327	543	783
0%	pool1	lib2	homeslot	-	
2MA139L5	ok	appendable	1327	655	672
0%	pool2	lib2	homeslot	-	
2MA144L5	ok	appendable	1327	655	672
0%	pool1	lib2	homeslot	-	
1IA134L5	ok	appendable	1327	655	672
0%	pool2	lib2	homeslot	-	
1SA180L5	ok	appendable	1327	0	1327
0%	pool2	lib2	homeslot	-	
1FB922L5	ok	appendable	1327	885	441
0%	pool2	lib2	homeslot	-	

4.7 Tips when upgrading host operating system

IBM Spectrum Archive EE cluster down time can be decreased on a multi-node environment by actively removing and adding nodes to the cluster.

If upgrading a single node cluster or if the environment does not follow the Validity Checklist, see [Upgrading Operating System from RH7.x to RH8.x](#).

Similar node refresh techniques are described in 3.8, “High-level component upgrade steps” on page 70. This information provides more references when planning or performing host operating system upgrades.

Planning the upgrade

The upgrade must be planned with care because the total performance of IBM Spectrum Archive EE decreases along with the nodes that are removed during this procedure.

Requests to IBM Spectrum Archive EE must not exceed the total computing capabilities of the current state of the cluster. Be sure to consider the expected workload that IBM Spectrum Archive EE needs during the time when the nodes are removed for the operating system upgrade. If available, adding a temporary node to the cluster can help to maintain needed computing capabilities.

Validity checklist

Use this check list to confirm whether the upgrade procedure is fully planned. All requirements *must* met to perform the upgrade:

- The cluster consists of more than one node per library.
- All nodes in the cluster are upgraded to 1.3.1.2 or later, and have matching release versions.
- The cluster can handle requests as the nodes are removed.
- Each node that is left in the cluster can reach at least one Control Path Drive on a library.

Performing the upgrade

Caution: Ensure that IBM Spectrum Scale and IBM Spectrum Protect are upgraded to ensure interoperability.

Complete the following steps to upgrade the host operating system:

1. Ensure that the “Validity checklist” is fully met.
2. Stop the cluster to temporarily remove the nodes that is planned for the operating system upgrade.
3. Remove the nodes from the cluster by using `ltfsee_config REMOVE_NODE` command (see 5.2.1, “The ltfsee_config utility” on page 105). At least one control node *must* left in the cluster. After the planned nodes are removed, restart the cluster.
4. Perform the operating system upgrades to the nodes that were rolled out.
5. After the operating system upgrade is finished, stop the cluster temporarily to add the upgraded nodes.
6. Configure the upgraded node back into the cluster by using the `ltfsee_config ADD_NODE` or `ltfsee_config ADD_CTRL_NODE` command.
7. Repeat the process with other nodes until all nodes are upgraded

Note: These operating system upgrade steps refer only to the upgrade steps of IBM Spectrum Archive EE. The operating system upgrade *must* be carefully planned to meet all hardware and other software compatibility requirements. For operating system update assistance, contact IBM.



Configuring IBM Spectrum Archive Enterprise Edition

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:

- ▶ 5.1, “Configuration prerequisites” on page 98
- ▶ 5.2, “Configuring IBM Spectrum Archive EE” on page 105
- ▶ 5.3, “First-time start of IBM Spectrum Archive EE” on page 118

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 3 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.

5.1 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 5.2, “Configuring IBM Spectrum Archive EE” on page 105.

- ▶ The Configuration worksheet is completed and available during the configuration process.

Tip: Table 5-1 on page 98, Table 5-2 on page 99, Table 5-3 on page 99, Table 5-4 on page 100, Table 5-5 on page 100, and Table 5-6 on page 101 provide a set of sample configuration worksheets. You can print and use these samples 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.

5.1.1 Configuration worksheet tables

Print Table 5-1 on page 98, Table 5-2 on page 99, Table 5-3 on page 99, Table 5-4 on page 100, Table 5-5 on page 100, and Table 5-6 on page 101 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 5.1.2, “Obtaining configuration information” on page 101 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 5-4 on page 100, Table 5-5 on page 100, and Table 5-6 on page 101 with the required information and refer to this information as necessary during the configuration process, as described in 5.2, “Configuring IBM Spectrum Archive EE” on page 105.

Table 5-1, Table 5-2, and Table 5-3 show example configuration worksheets with the parameters completed for the lab setup that was used to write this book.

Table 5-1 lists the file systems.

Table 5-1 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 5-2 lists the logical tape library information.

Table 5-2 Example logical tape library

Logical Tape library			
Tape library information			
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	
Tape Drive information			
Drive Serial number	Assigned IBM Spectrum Scale node	CP? (Yes or No)	Linux device name in the node
9A700M0029	htohru9	YES	/dev/sgXX
1068000073	htohru9	NO	/dev/sgYY

Table 5-3 lists the nodes.

Table 5-3 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 5-1 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 by using as shown in Table 5-2 on page 99. 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	769 - 1023

Figure 5-1 Obtain the starting SCSI element address of a TS3500 logical library

Table 5-4 shows an example of a blank file systems worksheet.

Table 5-4 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 5-5 shows an example of a blank logical tape library worksheet.

Table 5-5 Example logical tape library

Logical Tape library			
Tape library information			
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)			
Tape Drive information			
Drive Serial number	Assigned IBM Spectrum Scale node	CP? (Yes or No)	Linux device name in the node

Table 5-6 shows an example a blank nodes worksheet.

Table 5-6 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)

5.1.2 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 5-1 on page 102):

```
# mmstartup -a
```
3. Mount GPFS (if it is not already mounted) by running the following command (see Example 5-1 on page 102):

```
# 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 5-1 on page 102):

```
# mmlsfs all
```
5. Go to the Configuration worksheet (provided in 5.1.3, “Configuring key-based login with OpenSSH” on page 103) 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 5.1.4, “Preparing the IBM Spectrum Scale file system for IBM Spectrum Archive EE” on page 104.
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 5-1 on page 98.
8. Obtain a list of all IBM Spectrum Scale nodes in the IBM Spectrum Scale cluster by running the following command (see Example 5-1 on page 102):

```
# mmlsnode
```
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 5-2 on page 99. For more information and support, see the IBM Documentation website 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 Documentation website 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 5-1 shows how to obtain the information for the worksheet.

Example 5-1 Obtain the IBM Spectrum Scale required information for the configuration worksheet

```
[root@l1tfs97 ~]# mmstartup -a
Fri Apr 5 14:02:32 JST 2013: mmstartup: Starting GPFS ...
htohru9.ltd.sd1: The GPFS subsystem is already active.

[root@l1tfs97 ~]# mmmount all
Fri Apr 5 14:02:50 JST 2013: mmmount: Mounting file systems ...

[root@l1tfs97 ~]# mmlsfs all

File system attributes for /dev/gpfs:
=====
flag                value                description
-----
-f                  8192                 Minimum fragment (subblock) size in bytes
-i                  4096                 Inode size in bytes
-I                  32768                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                  32                   Estimated number of nodes that will mount file
system
-B                  4194304              Block size
-Q                  none                  Quotas accounting enabled
                   none                  Quotas enforced
                   none                  Default quotas enabled
--perfilesset-quota no                    Per-filesset quota enforcement
--filesetdf         no                    Filesset df enabled?
-V                  20.01 (5.0.2.0)      File system version
--create-time       Thu Sep 20 14:19:23 2018 File system creation time
-z                  yes                   Is DMAPAPI enabled?
-L                  33554432              Logfile size
```

```

-E          yes          Exact mtime mount option
-S          relatime     Suppress atime mount option
-K          whenpossible Strict replica allocation option
--fastea   yes          Fast external attributes enabled?
--encryption no         Encryption enabled?
--inode-limit 600000512 Maximum number of inodes
--log-replicas 0        Number of log replicas
--is4KAligned yes       is4KAligned?
--rapid-repair yes      rapidRepair enabled?
--write-cache-threshold 0 HAWC Threshold (max 65536)
--subblocks-per-full-block 512 Number of subblocks per full block
-P          system      Disk storage pools in file system
--file-audit-log no     File Audit Logging enabled?
--maintenance-mode no  Maintenance Mode enabled?
-d          nsd208_209_1;nsd208_209_2;nsd208_209_3;nsd208_209_4 Disks in file
system
-A          yes          Automatic mount option
-o          none         Additional mount options
-T          /ibm/glues Default mount point
--mount-priority 0      Mount priority

[root@l1tfs97 ~]# mmlsnode
GPFS nodeset      Node list
-----
      htohru9      htohru9

```

5.1.3 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 76 and 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 78.

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 pressing `enter` 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.

5.1.4 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 [IBM Documentation](#).

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 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 5.1.1, “Configuration worksheet tables” on page 98.

```
# 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; therefore, the `-z` option should not be specified. The preferred configuration is to have one file system with DMAPI-enabled.

For more information about the inode size, see 7.5, “Preferred inode size for IBM Spectrum Scale file systems” on page 239.

4. Mount the GPFS file system by running the following command:

```
# mmmount gpfs -a
```

For more information about the `mmmount` command, see the following resources:

- ▶ *General Parallel File System Version 4 Release 1.0.4 Advanced Administration Guide*, SC23-7032
- ▶ *IBM Spectrum Scale: Administration Guide*, which is available at [IBM Documentation](#).

5.2 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.

5.2.1 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 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 cluster.

- **ADD_CTRL_NODE [-g | -c | -a]**
Adds the local node as the control (MMM) node to a tape library in an 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 | -a]**
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.
- **UPDATE_FS_INFO**
Apply the current information of IBM Spectrum Scale (GPFS) file system to IBM Spectrum Archive EE configuration.
- **REMOVE_NODE [-N <node_id>] [-f]**
Removes the node and the drives configured for that node from the configuration.
- **REMOVE_NODEGROUP -l <library> -G <removed_nodegroup>**
Removes the node group that is no longer used.
- **INFO**
Shows the current configuration of this cluster.
- **LIST_LIBRARIES**
Shows the serial numbers of the tape libraries that are configured in the cluster.
- **REPLACE_LIBRARY [-b]**
Sets the serial number detected by the node to that of the configured library.
- **LIST_MOVE_POOLS**
Shows the pool translation table.
- **PREPARE_MOVE_POOL -p <pool_name> -s <source_library> -d <destination_library> [-G <node_group>] [-b]**
Prepares the pool translation table information for pool relocations between libraries.
- **CANCEL_MOVE_POOL -p <pool_name> -s <source_library> [-b]**
Cancels the PREPARE_MOVE_POOL operation for pool translation.
- **ACTIVATE_MOVE_POOL -p <pool_name> -s <source_library> -d <destination_library> [-b]**
Activates the pool that was relocated to a different library.
- **RECREATE_STATESAVE**
Delete and initialize the whole of statesave. By using this command, all history and running task information are removed.

► Options:

- **-a**
Assigns the IP address of the Admin node name as the control node (MMM), or as the local node. If -a is not used, the IP address of the Daemon node name is assigned.
- **-c**
Check and show the 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.
- **-G**
Specifies the node group to assign to the pool in the destination library during translation between libraries, or the node group to be removed.
- **-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**).

Important: When the active control node is removed by use of -f option, the library and the pool information stored in the internal database will be in-validated. If files that are migrated only to that library are left in the system, recalls of those files will no longer be possible.

- **-b**
Skips restarting the HSM daemon as a post process of the operation.
- **-p**
Specifies the name of the pool to be relocated to a different library.
- **-P**
Specifies the directory path storing SOBAR Support for System Migration result.
- **-s**
Specifies the name of the source library for a pool relocation procedure.
- **-d**
Specifies the destination library for a pool relocation procedure.
- **-l**
Specifies the library name in which to remove a node group with the -N option.

5.2.2 Configuring a single node cluster

Before you begin this procedure, ensure that all of the tasks that are described in 5.1, “Configuration prerequisites” on page 98 are met. Figure 5-2 shows a single-node configuration that is described in this section.

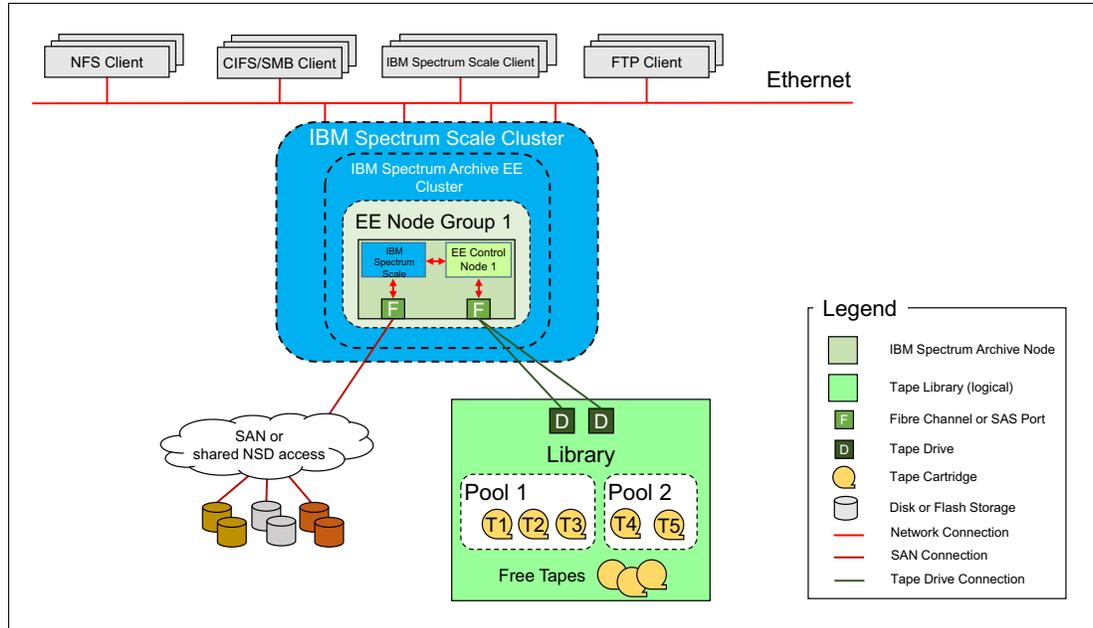


Figure 5-2 IBM Spectrum Archive single-node configuration

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 5.2.3, “Configuring a multiple-node cluster” on page 112.

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 5-2 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. The example stops the process and will not proceed to the `ADD_CTRL_NODE`, though you may answer [y] to continue to `ADD_CTRL_NODE` (step 5).

Example 5-2 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.

Then do you want to perform the ADD_CTRL_NODE mode? [Y/n]: n

Important: During the first run of the `ltfsee_config -m CLUSTER` command, if 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 5.1.4, “Preparing the IBM Spectrum Scale file system for IBM Spectrum Archive EE” on page 104. 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 5-3 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 5-3 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_ltfsm1
Added this node (ltfsm1.tuc.stglabs.ibm.com, node id 2) to library lib_ltfsm1
as its control node.
```

```
## 6. Add this node to a node group ##
Added this node (ltfsm1.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 (ltfsm1.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 5.2.3, “Configuring a multiple-node cluster” on page 112.

5.2.3 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 6.7, “IBM Spectrum Archive EE automatic node failover” on page 153.

Figure 5-3 shows a multiple-node cluster configuration that is described in this section.

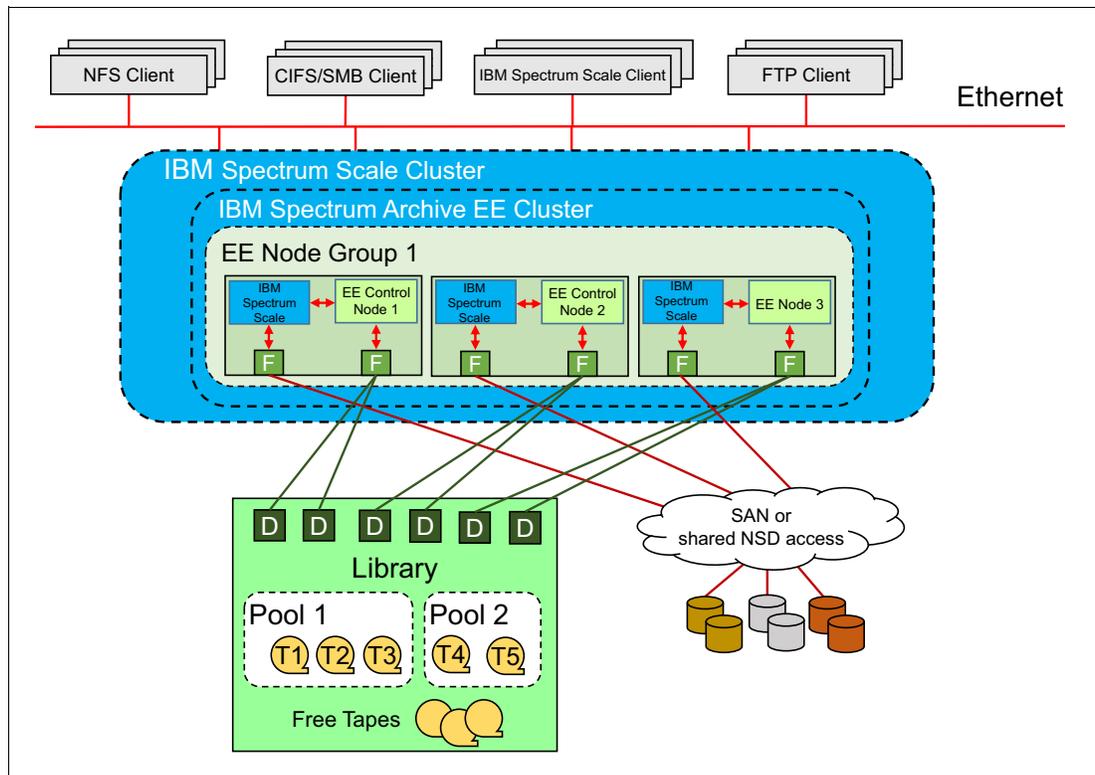


Figure 5-3 IBM Spectrum Archive multiple-node cluster configuration

Before configuring more nodes, ensure that all tasks that are described in 5.1, “Configuration prerequisites” on page 98 are completed and that the first node of the cluster environment is configured, as described in 5.2.2, “Configuring a single node cluster” on page 108.

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 5.2.2, “Configuring a single node cluster” on page 108.

Example 5-4 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 5-4 Adding secondary node as a redundant control node

```
[root@l1tfsm12 ~]# 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) preformance 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 5.2.2, “Configuring a single node cluster” on page 108:

- ▶ `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 5-3 on page 112, 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 5.2.4, “Configuring a multiple-node cluster with two tape libraries” on page 115.

5.2.4 Configuring a multiple-node cluster with two tape libraries

Starting with IBM Spectrum Archive V1.2, 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 [IBM Documentation](#).

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 information about IBM Spectrum Scale AFM support, see 2.2.5, “Active File Management” on page 30.

To add nodes to form a multiple-node cluster configuration with two tape libraries after the first node is configured, complete this task. Figure 5-4 shows the configuration with two tape libraries.

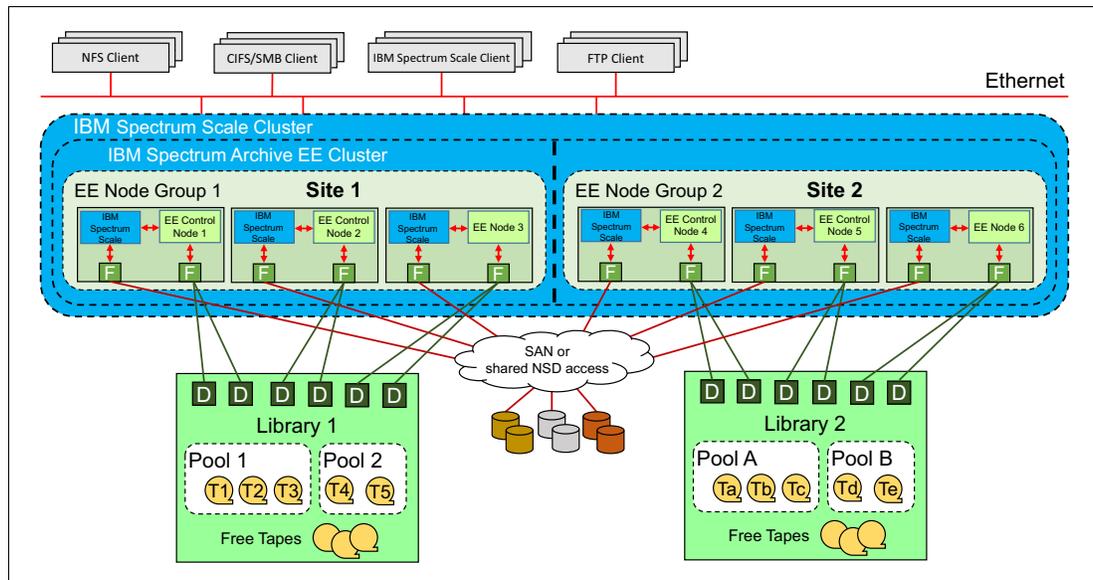


Figure 5-4 IBM Spectrum Archive multiple-node cluster configuration across two tape libraries

Before configuring more nodes, ensure that all tasks that are described in 5.1, “Configuration prerequisites” on page 98 are completed and that the first node of the cluster environment is configured, as described in 5.2.2, “Configuring a single node cluster” on page 108.

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 5-4 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 5-4 as an example, the `ltfsee_config -m ADD_NODE` command is run on EE Node 5 and EE Node 6.

5.2.5 Modifying a multiple-node configuration for control node redundancy

If users are upgrading to the IBM Spectrum Archive EE V1.3.0.0 from a previous version and have a multiple-node configuration with no redundant control nodes, users must manually set a redundant control node. See 4.3.2, “Installing, upgrading, or uninstalling IBM Spectrum Archive EE” on page 78 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 `eadm cluster stop`, to stop IBM Spectrum Archive EE and LE+. 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 5-5 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 5-5 Setting 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 `eeadm cluster start` command. Then, run `eeadm node list` 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 5-6 shows output from a multiple-node configuration with two tape libraries

Example 5-6 eeadm node list

```
[root@l1fsm11 ~]# eeadm node list
```

Node ID	State	Node IP	Drives	Ctrl Node	Library	Node Group
4	Available	9.11.120.224	2	Yes	l1fsee_lib2	G0
l1fsm14						
3	Available	9.11.120.207	2	Yes(Active)	l1fsee_lib2	G0
l1fsm13						
2	Available	9.11.120.201	2	Yes	l1fsee_lib1	G0
l1fsm12						
1	Available	9.11.120.198	2	Yes(Active)	l1fsee_lib1	G0
l1fsm11						

5.3 First-time start of IBM Spectrum Archive EE

To start IBM Spectrum Archive EE the first time, complete the following steps:

1. Check that the following embedded, customized IBM Tivoli® Storage Manager for Space Management (HSM) client components are running on each IBM Spectrum Archive EE node:

```
# ps -ef|grep dsm
```

If HSM is already running on system, the output of this command should show the dsm processes as shown in Example 5-7. HSM is not running if no output is shown. Information about how to start HSM can be referenced in 6.2.3, “Hierarchical Space Management” on page 138.

2. Start the IBM Spectrum Archive EE program by running the following command:

```
/opt/ibm/l1fsee/bin/eeadm cluster start
```

Important: If the `eeadm cluster 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 78.

Example 5-7 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 5-7 Start IBM Spectrum Archive EE the first time

```
[root@l1tfsml1 ~]# 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@l1tfsml1 ~]# eeadm cluster start
Library name: libb, library serial: 0000013400190402, control node (l1tfssee_md) IP address:
9.11.244.46.
Starting - sending a startup request to libb.
Starting - waiting for startup completion : libb.
Starting - opening a communication channel : libb.
.
Starting - waiting for getting ready to operate : libb.
.....
Started the IBM Spectrum Archive EE services for library libb with good status.
[root@l1tfsml1 ~]# eeadm node list
```

Node ID	State	Node IP	Drives	Ctrl Node	Library	Node Group	Host Name
1	available	9.11.244.46	4	yes(active)	libb	G0	l1tfsml1

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 6, “Managing daily operations of IBM Spectrum Archive Enterprise Edition” on page 129.

5.3.1 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, “Installing IBM Spectrum Archive Enterprise Edition” on page 73 to install Spectrum Archive EE and then in Chapter 5, “Configuring IBM Spectrum Archive Enterprise Edition” on page 97. If performed in that order, you can skip this section. See 7.10.3, “IBM Spectrum Archive EE migration policy with AFM” on page 246 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 more information about configuring IBM Spectrum Scale AFM, see the AFM documentation at [IBM Documentation](#).

5.3.2 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 8-14 on page 291.

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 5-8.

Example 5-8 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@l1fseeecachesrv ~]# systemctl start hsm
[root@l1fseeecachesrv ~]# 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@l1fseeecachesrv ~]# 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 5-9 shows example content of the exports file for file set IWhome.

Example 5-9 Contents of an exports file

```
[root@l1tfseehomesrv ~]# 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 5-10 shows an example of starting and checking the NFS service.

Example 5-10 Starting and checking the status of NFS service

```
[root@l1tfseehomesrv ~]# systemctl start nfs
[root@l1tfseehomesrv ~]# 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 l1tfseehomesrv.tuc.stglabs.ibm.com systemd[1]: Starting NFS
server and services...
Mar 21 15:58:43 l1tfseehomesrv.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 5-11 shows the execution of the **mmafmconfig** command with the IWhome file set.

Example 5-11 Execution of mmafmconfig enable <path-to-fileset>

```
[root@l1tfseehomesrv ~]# mmafmconfig enable /ibm/glues/IWhome/
[root@l1tfseehomesrv ~]#
```

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 **eeadm cluster 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 5-12 shows the output of running `mmchnode` on one cache node.

Example 5-12 Setting gateway nodes for cache clusters

```
[root@l1fseecachesrv ~]# mmchnode --gateway -N l1fseecachesrv
Tue Mar 21 16:49:16 MST 2017: mmchnode: Processing node
l1fseecachesrv.tuc.stglabs.ibm.com
[root@l1fseecachesrv ~]#
```

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 5-13:
 - a. `mmmstartup -a`
 - b. `mmgetstate -a`
 - c. `mmm mount all -a` (optional, only if the GPFS file system is not mounted automatically)

Example 5-13 Starting and mounting IBM Spectrum Scale and GPFS file system

```
[root@l1fseecachesrv ~]# mmstartup -a
Tue Mar 21 14:37:57 MST 2017: mmstartup: Starting GPFS ...
[root@l1fseecachesrv ~]# mmgetstate -a
```

Node number	Node name	GPFS state
1	l1fseecachesrv	arbitrating

```
[root@l1fseecachesrv ~]# mmgetstate -a
```

Node number	Node name	GPFS state
1	l1fseecachsrv1	active

```
[root@l1fseecachesrv ~]# mmmount all -a
Tue Mar 21 14:40:36 MST 2017: mmmount: Mounting file systems ...
[root@l1fseecachesrv ~]#
```

8. 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 5-14 shows the execution of `mmcrfileset` to create a cache fileset.

Example 5-14 Creating a cache fileset that targets the home fileset

```
[root@l1fseecachesrv ~]# mmcrfileset gpfs iwcache -p afmmode=iw -p
afmtarget=l1fseehomesrv:/ibm/glues/IWhome --inode-space=new
Fileset iwcache created with id 1 root inode 4194307.
```

9. 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 5-15 shows output of running `mmlinkfileset`.

Example 5-15 Linking the GPFS fileset to a directory on the GPFS file system

```
[root@l1fseecachesrv 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.

5.3.3 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 8.11.2, “Asynchronous archive replication” on page 291.

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.

1. 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 5-16.

Example 5-16 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@l1fseehomesrv ~]# systemctl start hsm
[root@l1fseehomesrv ~]# 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@l1fseehomesrv ~]# 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 5-17 shows example content of the exports file for fileset IWhome.

Example 5-17 Contents of an exports file

```
[root@l1tfseehomesrv ~]# 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 5-18 shows an example of starting and checking the NFS service.

Example 5-18 Starting and checking the status of NFS service

```
[root@l1tfseehomesrv ~]# systemctl start nfs
[root@l1tfseehomesrv ~]# 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 l1tfseehomesrv.tuc.stglabs.ibm.com systemd[1]: Starting NFS
  server and services...
Mar 21 15:58:43 l1tfseehomesrv.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 5-19 shows the execution of the `mmafmconfig` command with the IWhome fileset.

Example 5-19 Execution of mmafmconfig enable <path-to-fileset>

```
[root@l1tfseehomesrv ~]# mmafmconfig enable /ibm/glues/IWhome/
[root@l1tfseehomesrv ~]#
```

5. After the fileset has been enabled for AFM, start Spectrum Archive if it has not been started previously by running `eadm cluster 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 file sets.

6. If IBM Spectrum Archive EE is active, properly shut it down by using the commands in Example 5-20.

Example 5-20 Shutting down IBM Spectrum Archive EE

```
[root@l1tfseecachesrv ~]# eadm cluster stop
Library name: libb, library serial: 0000013400190402, control node (l1tfsee_md)
IP address: 9.11.244.46.
Stopping - sending request and waiting for the completion.
.
Stopped the IBM Spectrum Archive EE services for library libb.
[root@l1tfseecachesrv ~]# pidof mmm
[root@l1tfseecachesrv ~]# umount /l1tfs
[root@l1tfseecachesrv ~]# pidof l1tfs
[root@l1tfseecachesrv ~]#
```

7. If IBM Spectrum Scale is active, properly shut it down by using the commands in Example 5-21.

Example 5-21 Shutting down IBM Spectrum Scale

```
[root@l1tfseecachesrv ~]# 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@l1tfseecachesrv ~]# 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@l1tfseecachesrv ~]# systemctl stop hsm
[root@l1tfseecachesrv ~]# mmumount all -a
Wed Mar 22 13:31:44 MST 2017: mmumount: Unmounting file systems ...
[root@l1tfseecachesrv ~]# 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
l1tfseecachesrv.tuc.stglabs.ibm.com: Shutting down!
l1tfseecachesrv.tuc.stglabs.ibm.com: 'shutdown' command about to kill process
24101
l1tfseecachesrv.tuc.stglabs.ibm.com: Unloading modules from
/lib/modules/3.10.0-229.el7.x86_64/extra
l1tfseecachesrv.tuc.stglabs.ibm.com: Unloading module mmfs26
l1tfseecachesrv.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 5-22.

Example 5-22 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 5-23.

Example 5-23 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 5-24.

Example 5-24 Creating an IW fileset

```
[root@ltfseecachesrv ~]# mmcrfileset gpfs iwcache -p afmmode=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 5-25.

Example 5-25 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 5-26.

Example 5-26 Obtaining metadata and HSM file system(s)

```
[root@l1tfseeecachesrv ~]# l1tfsee_config -m INFO
The EE configuration script is starting: /opt/ibm/l1tfsee/bin/l1tfsee_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
l1tfsee_config.filesystem exists.
Metadata Filesystem:
  /ibm/glues
HSM Filesystems:
  /ibm/glues
Library: Name=l1tfsee_lib1, S/N=0000013400190402
Node Group: Name=l1tfseeecachesrv
Node: l1tfseeecachesrv.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 that is shown in Example 5-27 with the same file systems that were recorded from step 7.

Example 5-27 Modify IBM Spectrum Archive EE configuration for IBM Spectrum Scale AFM

```
[root@l1tfseeecachesrv ~]# l1tfsee_config -m UPDATE_FS_INFO
The EE configuration script is starting: /opt/ibm/l1tfsee/bin/l1tfsee_config -m
UPDATE_FS_INFO
UPDATE_FS_INFO mode starts .
## Step-1. Check to see if the cluster is already created ##
The cluster is already created and the configuration file
l1tfsee_config.filesystem exists.
Successfully validated the prerequisites.
** Select file systems to configure for the 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/flash Mount point(/flash)
  2. /dev/gpfs Mount point(/ibm/gpfs)
  a. Select all file systems
  q. Quit
Input number > a
## Step-2. Add selected file systems to Space Management ##
Added the selected file systems to the space management.
## Step-3. 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.
Disabling runtime AFM file state checking.
UPDATE_FS_INFO mode completed.

14. Start IBM Spectrum Archive EE by using the commands in Example 5-28.

Example 5-28 Start IBM Spectrum Archive EE

```
[root@ltfseeecachesrv ~]# eeadm cluster start
Library name: libb, library serial: 0000013400190402, control node (ltfsee_md)
IP address: 9.11.244.46.
Starting - sending a startup request to libb.
Starting - waiting for startup completion : libb.
Starting - opening a communication channel : libb.
.
Starting - waiting for getting ready to operate : libb.
.....
Started the IBM Spectrum Archive EE services for library libb 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 5-29. If it has not been properly set, modify the file so that the AFMSKIPUNCACHEDFILES flag is set to yes.

Example 5-29 Validating AFMSKIPUNCACHEDFILES is set to yes

```
[root@ltfseeecachsrv ~]# 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.



Managing daily operations of IBM Spectrum Archive Enterprise Edition

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:

- ▶ 6.1, “Overview” on page 131
- ▶ 6.2, “Status information” on page 136
- ▶ 6.3, “Upgrading components” on page 142
- ▶ 6.4, “Starting and stopping IBM Spectrum Archive EE” on page 146
- ▶ 6.5, “Task command summaries” on page 147
- ▶ 6.6, “IBM Spectrum Archive EE database backup” on page 152
- ▶ 6.7, “IBM Spectrum Archive EE automatic node failover” on page 153
- ▶ 6.8, “Tape library management” on page 154
- ▶ 6.9, “Tape storage pool management” on page 162
- ▶ 6.10, “Pool capacity monitoring” on page 163
- ▶ 6.11, “Migration” on page 165
- ▶ 6.12, “Premigration” on page 184
- ▶ 6.13, “Preserving file system objects on tape” on page 186
- ▶ 6.14, “Recall” on page 190
- ▶ 6.15, “Recalling files to their resident state” on page 196
- ▶ 6.16, “Reconciliation” on page 197
- ▶ 6.17, “Reclamation” on page 200
- ▶ 6.18, “Checking and repairing tapes” on page 201
- ▶ 6.19, “Importing and exporting” on page 202
- ▶ 6.20, “Drive Role settings for task assignment control” on page 207
- ▶ 6.21, “Tape drive intermix support” on page 208
- ▶ 6.22, “Obtaining the location of files and data” on page 213
- ▶ 6.23, “Obtaining system resources, and tasks information” on page 214
- ▶ 6.24, “Monitoring the system with SNMP” on page 216
- ▶ 6.25, “Configuring Net-SNMP” on page 218

- ▶ 6.26, “IBM Spectrum Archive REST API” on page 219
- ▶ 6.27, “File system migration” on page 231

Note: The steps that you perform are the same as we describe in this Redbooks publication. You might see slightly different output responses in your environment, depending on your version and release of the product.

6.1 Overview

The following terms specific to IBM Spectrum Archive EE operations are used in this chapter:

Migration	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.
Premigration	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.
Recall	The movement of migrated files from tape cartridges back to the originating GPFS file system on disk, which is the reverse of migration.
Reconciliation	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.
Reclamation	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 Spectrum Protect from the IBM Spectrum Archive family. All active files are consolidated onto a second tape cartridge, which improves overall tape usage.
Library rescan	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.
Tape validate	The process of validating the current condition of a tape by loading it to the tape drive and updating the tape state.
Tape replace	The process of moving the contents of a tape that previously suffered some error to another tape in the same pool.
Import	The addition of an LTFS tape cartridge to IBM Spectrum Archive EE.
Export	The removal of an LTFS tape cartridge from IBM Spectrum Archive EE.
Data migration	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">- within a pool- pool to pool

6.1.1 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 following two syntaxes `eeadm <resource_type> <action> [options]`, and `eeadm <subcommand> [options]`.

The following **eadm** command options are available. All options, except **listing resource types**, can be run only with root user permissions:

▶ **eadm cluster start**

Use this command to start the process of the IBM Spectrum Archive EE system on all configured servers, or on a specific library.

Important: If the **eadm cluster 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 78.

▶ **eadm cluster stop**

Use this command to Stop the process of the IBM Spectrum Archive EE system on all configured servers, or on a specific library.

▶ **eadm cluster failover**

Use this command to manually initiate a node failover process.

▶ **eadm cluster set**

Use this command to change the global configuration attributes of IBM Spectrum Archive cluster.

▶ **eadm cluster show**

Use this command to display the global configuration attributes of IBM Spectrum Archive cluster.

▶ **eadm drive assign**

Use this command to assign tape drives to the IBM Spectrum Archive EE server.

▶ **eadm drive unassign**

Use this command to unassign tape drives from the IBM Spectrum Archive EE server.

▶ **eadm drive up**

Use this command to enable a tape drive. The enabled drive can be used as a part of the IBM Spectrum Archive EE system.

▶ **eadm drive down**

Use this command to disable a tape drive. The disabled drive cannot be used as a part of the IBM Spectrum Archive EE system.

▶ **eadm drive list**

Use this command to list all the configured tape drives.

▶ **eadm drive set**

Use this command to change the configuration attributes of a tape drive.

▶ **eadm drive show**

Use this command to display the configuration attributes of a tape drive.

▶ **eadm file state**

Use this command to display the current data placement of files. Each file is in one of the following states:

- **resident:** The data is on disk

- **premigrated:** The data is both on disk and tapes
- **migrated:** The data is on tapes while its stub file remains on disk.
- ▶ **eeadm library list**
Use this command to list all the managed tape libraries.
- ▶ **eeadm library rescan**
Use this command to force the tape library to check and report its physical resources, and update the resource information kept in IBM Spectrum Archive EE.
- ▶ **eeadm library show**
Use this command to display the configuration attributes of a tape library
- ▶ **eeadm node down**
Use this command to disable one of the IBM Spectrum Archive EE servers temporarily for maintenance. The disabled node does not participate in the system.
- ▶ **eeadm node list**
Use this command to list the configuration and status of all the configured nodes.
- ▶ **eeadm node show**
Use this command to display the configuration attributes of the node.
- ▶ **eeadm node up**
Use this command to enable one of the IBM Spectrum Archive EE servers. The enabled node can be used as a part of the IBM Spectrum Archive EE system.
- ▶ **eeadm nodegroup list**
Use this command to list all the configured node groups.
- ▶ **eeadm pool create**
Use this command to create a tape pool.
- ▶ **eeadm pool delete**
Use this command to delete a tape pool to which no tapes are assigned.
- ▶ **eeadm pool list**
Use this command to list all the configured tape pools
- ▶ **eeadm pool set**
Use this command to change the configuration attributes of a pool.
- ▶ **eeadm pool show**
Use this command to display the configuration attributes of the tape pool.
- ▶ **eeadm tape list**
Use this command to list the configuration and status of all the tapes.
- ▶ **eeadm tape set**
Use this command to change the configuration attribute of a tape.
- ▶ **eeadm tape show**
Use this command to display the configuration attributes of a tape that is already assigned to the tape pool.

► **eeadm tape assign**

Use this command to format the tapes and assign it to the tape pool. The command fails if it contains any file object unless the `-f` option is used. Use the **eeadm tape import** command to make the tape that contains files a new member of the pool.

► **eeadm tape export**

Use this command to export the tape permanently from the IBM Spectrum Archive EE system and purges the GPFS files referring to the tape. The command internally runs the reconciliation process and identifies the active GPFS files in migrated or premigrated states. If a file refers to the tape to be exported and if the tape contains the last replica of the file, it deletes the GPFS file. After the successful completion of the command, the tape is unassigned from the tape pool, and the tape state becomes **exported**. Prior to running this command, all GPFS file systems must be mounted so that the command is able to check for the existence of the files on disk.

► **eeadm tape import**

Use this command to import tapes created and managed by other system to the IBM Spectrum Archive EE system, and makes the files on tape accessible from the GPFS namespace. The command will create the stub files and leave the files in the migrated state without transferring the file data back to disk immediately.

► **eeadm tape move**

Use this command to move the tape physically within the tape library. The command can move the tape to its home slot either from the tape drive or from the I/E slot (or I/O station) of the tape library. It can move the tape to I/E slot if the tape is:

- in the offline state
- the tape belongs to a tape pool that is undergoing pool relocation
- tape is currently not assigned to a pool

► **eeadm tape offline**

Use this command to set the tape to the offline state to prepare for moving the tape temporarily out of the tape library, until access to the data is required. The tape needs to come back to the original IBM Spectrum Archive EE system by using the **eeadm tape online** command.

► **eeadm tape online**

Use this command to make the offline tape accessible from the IBM Spectrum Archive EE system.

► **eeadm tape unassign**

Use this command to unassign the member tapes from the tape pool.

► **eeadm tape datamigrate**

Use this command to move the active contents of specified tapes to a different tape pool and updates the stub files on disk to point to the new data location. After the successful completion of the command, the tape is automatically unassigned from the source tape pool. The **datamigrate** command can be used to move the data on older technology tapes in the source tape pool to newer technology tapes in the destination tape pool.

► **eeadm tape reclaim**

Use this command to reclaim the unreferenced space of the specified tapes. It moves the active contents on the tapes to different tapes, then recycles the specified tapes. If the `--unassign` option is specified, the tape is automatically unassigned from the tape pool after the successful completion of the **eeadm tape reconcile** command.

Use this command to compare the contents of the tape with the files on the GPFS file systems, and reconciles the differences between them.

► **eeadm tape replace**

Use this command to move the contents of a tape that previously suffered some error to another tape in the same pool. This command is used on tapes in the **require_replace** or **need_replace** state. After the successful completion of the replacement process, the tape that had the error is automatically unassigned from the tape pool.

► **eeadm tape validate**

Use this command to validate the current condition of a tape by loading it to the tape drive, and update the tape state. The tape must be a member of the tape pool and online.

► **eeadm task cancel**

Use this command to cancel the active task. The command supports the cancellation of reclaim and datamigrate tasks only.

► **eeadm task clearhistory**

Use this command to delete the records of completed tasks to free up disk space.

► **eeadm task list**

Use this command to list the active or completed tasks.

► **eeadm task show**

Use this command to show the detailed information of the specified task.

► **eeadm migrate**

Use this command to move the file data to the tape pools to free up disk space, and sets the file state to **migrated**.

► **eeadm premigrate**

Use this command to copy the file data to the tape pools, and sets the file state to **premigrated**.

► **eeadm recall**

Use this command to recall the file data back from the tape and places the file in the **premigrated** state, or optionally in the **resident** state.

► **eeadm save**

Use this command to save the name of empty files, empty directories, and symbolic links on the tape pools.

6.1.2 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 which allows administrators to perform operations, see “IBM Spectrum Archive EE dashboard” on page 34 for more info.

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 **eeadm** command uses the following two syntax:

`eeadm <resource_type> <action> [options]`, and `eeadm <subcommand> [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 10.1, “Command-line reference” on page 316.

6.2 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 9, “Troubleshooting IBM Spectrum Archive Enterprise Edition” on page 293.

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.

6.2.1 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 following publications:

- ▶ *General Parallel File System Version 4 Release 1.0.4 Advanced Administration Guide*, SC23-7032
- ▶ *IBM Spectrum Scale: Administration Guide*, which is available at [IBM Documentation](#).

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 6-1 shows how to get status about the GPFS/IBM Spectrum Scale daemon on one or more nodes.

Example 6-1 Check the GPFS/IBM Spectrum Scale status

```
[root@l1tfs97 ~]# 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 6-2.

Example 6-2 Start GPFS/IBM Spectrum Scale

```
[root@ltfs97 ~]# mmstartup -a
Tue Apr 2 14:41:13 JST 2013: mmstartup: Starting GPFS ...
[root@ltfs97 ~]# mmgetstate -a
Node number  Node name      GPFS state
-----
           1      htohru9      active
```

If the status is active, also check the GPFS/IBM Spectrum Scale mount status by running the command that is shown in Example 6-3.

Example 6-3 Check the GPFS/IBM Spectrum Scale mount status

```
[root@ltfs97 ~]# mmlsmount all
File system gpfs is mounted on 1 nodes.
```

The message confirms that the GPFS file system is mounted.

6.2.2 IBM Spectrum Archive Library Edition component

IBM Spectrum Archive EE constantly checks to see whether the IBM Spectrum Archive Library Edition (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 6-4. The IBM Spectrum Archive LE component must be running on all EE nodes.

Example 6-4 Check the IBM Spectrum Archive LE component status (running)

```
[root@ltfs97 ~]# df -m
Filesystem          1M-blocks      Used Available Use% Mounted on
/dev/mapper/VolGroup-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/sg2       2147483648     0 2147483648    0% /ltfs
```

To start IBM Spectrum Archive LE run the `eeadm cluster start` command. If errors occur during the start of the IBM Spectrum Archive EE system, run the `eeadm node list` command to display which component failed to start. For more information about the updated `eeadm node list` command, see 6.7, “IBM Spectrum Archive EE automatic node failover” on page 153.

6.2.3 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 6-5.

Example 6-5 Check the HSM status by running ps

```
[root@l1tfs97 /]# 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 6-6.

Example 6-6 Check the HSM status by using dsmmigfs

```
[root@yurakucho ~]# dsmmigfs query -detail
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 11.0
  Client date/time: 03/06/21  05:25:59
(c) Copyright by IBM Corporation and other(s) 1990, 2020. All Rights Reserved.
```

The local node has Node ID: 1
The failover environment is active on the local node.
The recall distribution is enabled.
The monitoring of local space management daemons is active.

```
File System Name:/ibm/gpfs_data
High Threshold:90
Low Threshold:80
Premig Percentage:10
Quota: 21273588
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
Min Mig File Size: 0
Inline Copy Mode: MIG
Preferred Node: yurakucho Node ID: 1
Owner Node: yurakucho Node ID: 1

File System Name:/ibm/gpfs_meta
High Threshold:90
Low Threshold:80
Premig Percentage:10
Quota: 1100800
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
Min Mig File Size: 0
Inline Copy Mode: MIG
Preferred Node: yurakucho Node ID: 1
Owner Node: yurakucho 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 6-7.

Example 6-7 Check GPFS file system

```
[root@l1fs97 /] 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 5.2, “Configuring IBM Spectrum Archive EE” on page 105.

Permissions: Starting HSM requires root user permissions.

If the HSM watch daemon (`dsmwatchd`) is not running, Example 6-8 shows you how to start it.

Example 6-8 Start the HSM watch daemon

```
[root@l1tfsr11 ~]# systemctl start hsm.service
[root@l1tfsr11 ~]# 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 6-9 shows you how to start them.

Example 6-9 Start HSM

```
[root@yurakucho ~]# dsmmigfs start
IBM Spectrum Protect
Command Line Space Management Client Interface
  Client Version 8, Release 1, Level 11.0
  Client date/time: 03/06/21  05:29:22
(c) Copyright by IBM Corporation and other(s) 1990, 2020. 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.

6.2.4 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 `eeadm node list` command. You can also use this command to determine whether any component required for IBM Spectrum Archive EE failed to start. 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 does not require root user permissions.

If the MMM service is running correctly, you see a message similar to the message shown in Example 6-10.

Example 6-10 Check the IBM Spectrum Archive EE status

```
[root@l1tfsml1 ~]# eeadm node list
Node ID  State      Node IP      Drives  Ctrl Node    Library  Node Group  Host
Name
1        available  9.11.244.46  3       yes(active) libb     G0
lib_l1tfsml1
```

If the MMM service is not running correctly, you may see a message that is similar to the message shown in Example 6-11.

Example 6-11 Check the IBM Spectrum Archive EE status

```
[root@l1tfsml1 ~]# eeadm node list
```

```
Spectrum Archive EE service (MMM) for library libb fails to start or is not
running on lib_l1tfsml1 Node ID:1
```

```
Problem Detected:
Node ID  Error Modules
      1  LE; MMM;
```

In Example 6-11, IBM Spectrum Archive EE failed to start MMM because it was unable to mount LE. In this example, MMM failure to mount LE was caused by the server having no control path drives connected; therefore, the work around is to assign a control path drive to the server and have the monitor daemon automatically mount LE.

To monitor the process of IBM Spectrum Archive EEs start, run the **eeadm node list** command. If IBM Spectrum Archive EE is taking too long to recover, stop and start the process by using the **eeadm cluster stop/start** command. Example 6-12 shows the process of IBM Spectrum Archive EE recovering after a control path drive was connected.

Example 6-12 Start IBM Spectrum Archive EE

```
[root@l1tfsml1 ~]# eeadm node list
```

```
Spectrum Archive EE service (MMM) for library libb fails to start or is not
running on lib_l1tfsml1 Node ID:1
```

```
Problem Detected:
Node ID  Error Modules
      1  LE(Starting); MMM;
```

```
[root@kyoto ~]# eeadm node list
```

```
Node ID  State      Node IP      Drives  Ctrl Node   Library  Node Group  Host
Name
1        available  9.11.244.46   3      yes(active) libb     G0
lib_l1tfsml1
```

Important: If the **eeadm cluster 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 78.

6.3 Upgrading components

The following sections describe the process that is used to upgrade IBM Spectrum Scale and other components of IBM Spectrum Archive EE.

6.3.1 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 [IBM Documentation](#) and select the IBM Spectrum Scale release under the Cluster product libraries topic in the navigation pane that applies to your installation.

To update IBM Spectrum Scale, complete the following steps:

1. Stop IBM Spectrum Archive EE by running the command that is shown in Example 6-13.

Example 6-13 Stop IBM Spectrum Archive EE

```
[root@l1tfsml1 ~]# eeadm cluster stop
```

```
Library name: libb, library serial: 0000013400190402, control node (l1tfssee_md)  
IP address: 9.11.244.46.
```

```
Stopping - sending request and waiting for the completion.
```

```
..
```

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

2. Run the `pidof mmm` command on all EE Control Nodes until all MMM processes have been terminated.
3. Run the `pidof l1tfs` command on all EE nodes until all l1tfs processes are stopped.
4. Disable DSM failover by running the command that is shown in Example 6-14.

Example 6-14 Disable failover

```
[root@l1tfsml1 ~]# 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 IBM Spectrum Protect for Space Management HSM by running the command that is shown in Example 6-15.

Example 6-15 Stop HSM

```
[root@l1tfsml1~]# dsmmigfs stop
```

```
IBM Spectrum Protect
```

Command Line Space Management Client Interface
Client Version 8, Release 1, Level 11.0
Client date/time: 03/06/21 05:35:38
(c) Copyright by IBM Corporation and other(s) 1990, 2020. 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 6-16.

Example 6-16 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 6-17.

Example 6-17 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 6-18), you must wait for them to finish and then run the **mmumount all** command again.

Example 6-18 Processes running that prevent the unmounting of the GPFS file system

```
[root@ltfs97 ~]# 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 6-19.

Example 6-19 Shut down GPFS

```
[root@ltfs97 ~]# 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 6-20.

Example 6-20 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 6-21.

Example 6-21 Rebuild GPFS

```
mmbuildgpl
```

11. Start GPFS by running the command that is shown in Example 6-22.

Example 6-22 Start GPFS

```
[root@ltfs97 ~]# 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 6-23.

Example 6-23 Mount GPFS file systems

```
[root@ltfs97 ~]# mmmount all  
Tue Apr 16 23:48:09 JST 2013: mmmount: Mounting file systems ...
```

13. Start the watch daemon by running the command that is shown in Example 6-24.

Example 6-24 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 6-25.

Example 6-25 Start HSM

```
[root@yurakucho ~]# dsmmigfs start  
IBM Spectrum Protect  
Command Line Space Management Client Interface  
Client Version 8, Release 1, Level 11.0  
Client date/time: 03/06/21 05:42:15  
(c) Copyright by IBM Corporation and other(s) 1990, 2020. 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 6-26.

Example 6-26 Enable failover

```
[root@l1tfsml1 ~]# 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 6-27.

Example 6-27 Start IBM Spectrum Archive EE

```
[root@l1tfsml1 ~]# eadm cluster start
Library name: libb, library serial: 0000013400190402, control node (l1tfsml1)
IP address: 9.11.244.46.
Starting - sending a startup request to libb.
Starting - waiting for startup completion : libb.
Starting - opening a communication channel : libb.
.
Starting - waiting for getting ready to operate : libb.
.....
Started the IBM Spectrum Archive EE services for library libb with good status.
```

Optionally, you can check the status of each component when it is started, as described in 6.2, “Status information” on page 136.

6.3.2 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 78. 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.

6.3.3 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 78. Because HSM is a component of IBM Spectrum Archive EE, it is upgraded as part of the IBM Spectrum Archive EE upgrade.

6.3.4 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 78.

6.4 Starting and stopping IBM Spectrum Archive EE

This section describes how to start and stop IBM Spectrum Archive EE.

6.4.1 Starting IBM Spectrum Archive EE

Run the `eeadm cluster start` command to start the IBM Spectrum Archive the IBM Spectrum Archive EE system. The HSM components must be running before you can use this command. You can run the `eeadm cluster 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 6-28.

Example 6-28 Start IBM Spectrum Archive EE

```
[root@ltfsm11 ~]# eeadm cluster start
Library name: libb, library serial: 0000013400190402, control node (ltfsee_md) IP
address: 9.11.244.46.
Starting - sending a startup request to libb.
Starting - waiting for startup completion : libb.
Starting - opening a communication channel : libb.
.
Starting - waiting for getting ready to operate : libb.
.....
Started the IBM Spectrum Archive EE services for library libb with good status.
```

Important: If the `eeadm cluster 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 78.

You can confirm that IBM Spectrum Archive EE is running by referring to the steps in Example 6-10 on page 140 or by running the command in Example 6-29.

Example 6-29 Check the status of all available IBM Spectrum Archive EE nodes

```
[root@ltfsm11 ~]# eeadm node list
Node ID  State      Node IP      Drives  Ctrl Node  Library  Node Group  Host
Name
1        available  9.11.244.46  3      yes(active)  libb    G0
lib_ltfsm11
```

6.4.2 Stopping IBM Spectrum Archive EE

The **eeadm cluster stop** command stops the IBM Spectrum Archive EE system on all EE Control Nodes.

For example, to start IBM Spectrum Archive EE, run the command that is shown in Example 6-30.

Example 6-30 Stop IBM Spectrum Archive EE

```
[root@l1tfsm11 ~]# eeadm cluster stop
Library name: libb, library serial: 0000013400190402, control node (l1tfsee_md) IP
address: 9.11.244.46.
Stopping - sending request and waiting for the completion.
..
Stopped the IBM Spectrum Archive EE services for library libb.
```

In some cases, you might see the GLESM658I informational message if there are active tasks on the task queue in IBM Spectrum Archive EE:

There are still tasks in progress.
To terminate IBM Spectrum Archive EE for this library,
run the "eeadm cluster stop" command with the "-f" or "--force" option.

If you are sure that you want to stop IBM Spectrum Archive EE, run the **eeadm cluster stop** command with the **-f** option, which forcefully stops any running IBM Spectrum Archive EE tasks abruptly. Note that this option will stop all tasks that are currently running on EE, and will need to be manually re-run if necessary once the cluster has restarted.

6.5 Task command summaries

For any IBM Spectrum Archive EE commands that generates a task, the **eeadm task list** and **eeadm task show** commands will be used to display the task information. For the list of task generating commands, refer to “User Task Reporting” on page 7. The **eeadm task list** and **eeadm task show** commands has replaced the **l1tfsee info scans** and **l1tfsee info jobs** commands.

6.5.1 eeadm task list

The most common use will be to get the list of active tasks which will either be in the running, waiting, or interrupted status. The output will be sorted by status in the following order:

1. Running
2. Waiting or interrupted

Within each status group, the entries will be sorted by Priority (H, M, L order).

- H will be any transparent or selective recall commands
- M will be any premigration, migration, or save commands
- L will be everything else (general commands)

Within Running status, within the same Priority, the entries are sorted by Started Time (oldest at the top).

Example 6-31 shows active tasks.

Example 6-31 Viewing active tasks

```
[root@kyoto ~]# eadm task list
TaskID Type Priority Status #DRV CreatedTime(-0700)
StartedTime(-0700)
 18822 selective_recall H running 1 2019-01-07_11:43:04
2019-01-07_11:43:04
 18823 selective_recall H waiting 0 2019-01-07_11:43:09
2019-01-07_11:43:09
 18824 migrate M waiting 0 2019-01-07_11:43:11
2019-01-07_11:43:11
```

The other common use will be to get the list of completed tasks which returns the prior task IDs, result, and date/time information, sorted by completed time (oldest at the top). For administrators, this allows a quick view into task IDs, the history of what task has been executed recently and their results. For more information about the specified task, see “eadm task show” on page 148.

Example 6-32 shows five completed tasks.

Example 6-32 View previous 5 completed tasks

```
[root@kyoto ~]# eadm task list -c -n 5
TaskID Type Result CreatedTime(-0700) StartedTime(-0700)
CompletedTime(-0700)
 18819 selective_recall succeeded 2019-01-07_10:56:23 2019-01-07_10:56:23
2019-01-07_10:57:51
 18820 selective_recall succeeded 2019-01-07_11:39:03 2019-01-07_11:39:03
2019-01-07_11:39:58
 18822 selective_recall succeeded 2019-01-07_11:43:04 2019-01-07_11:43:04
2019-01-07_11:44:31
 18823 selective_recall succeeded 2019-01-07_11:43:09 2019-01-07_11:43:09
2019-01-07_11:45:49
 18824 migrate succeeded 2019-01-07_11:43:11 2019-01-07_11:43:11
2019-01-07_11:45:58
```

For more information about the task status, see “User Task Reporting” on page 7.

6.5.2 eadm task show

The most common use will be to get the detailed information of the specified task, including a verbose option which shows the output messages and subtask information. For any failed tasks, the administrator can preform resubmission of those tasks or next step recovery procedures. Example 6-33 shows the output verbosely from an active migration task.

Example 6-33 Verbose output of an active migration task

```
[root@kyoto ~]# eadm task show 18830 -v
=== Task Information ===
Task ID: 18830
Task Type: migrate
Command Parameters: eadm migrate mig3 -p pool1
Status: running
Result: -
```

```

Accepted Time:      Mon Jan  7 11:58:34 2019 (-0700)
Started Time:      Mon Jan  7 11:58:34 2019 (-0700)
Completed Time:    -
In-use Resources:  1068045923(D00369L5):pool1:G0:libb
Workload:         100 files. 1 replicas.
                  7545870371 bytes to copy. 1 copy tasklets on pool1@libb.
Progress:         -
                  0/1 copy tasklets completed on pool1@libb.
Result Summary:   -
Messages:
  2019-01-07 11:58:34.231005 GLESM896I: Starting the stage 1 of 3 for migration
  task 18830 (qualifying the state of migration candidate files).
  2019-01-07 11:58:37.133670 GLESM897I: Starting the stage 2 of 3 for migration
  task 18830 (copying the files to 1 pools).

```

```

--- Subtask(level 1) Info ---
Task ID:          18831
Task Type:        copy_replica
Status:           running
Result:           -
Accepted Time:    Mon Jan  7 11:58:37 2019 (-0700)
Started Time:    Mon Jan  7 11:58:37 2019 (-0700)
Completed Time:  -
In-use Libraries: libb
In-use Node Groups: G0
In-use Pools:    pool1
In-use Tape Drives: 1068045923
In-use Tapes:    D00369L5
Workload:        7545870371 bytes to copy. 1 copy tasklets on pool1@libb.
Progress:        0/1 copy tasklets completed on pool1@libb.
Result Summary:  -
Messages:
  2019-01-07 11:58:37.196346 GLESM825I: The Copy tasklet (0x3527880) will be
  dispatched on drive 1068045923 from the write queue (tape=D00369L5).
  2019-01-07 11:58:37.202930 GLESM031I: A list of 100 file(s) has been added to
  the migration and recall queue.

```

Example 6-34 shows the output verbosely of a completed migration task.

Example 6-34 Verbose output of a completed migration task

```

[root@kyoto ~]# eadm task show 18830 -v
=== Task Information ===
Task ID:          18830
Task Type:        migrate
Command Parameters: eadm migrate mig3 -p pool1
Status:           completed
Result:           succeeded
Accepted Time:    Mon Jan  7 11:58:34 2019 (-0700)
Started Time:    Mon Jan  7 11:58:34 2019 (-0700)
Completed Time:  Mon Jan  7 11:59:34 2019 (-0700)
Workload:        100 files. 1 replicas.
                  7545870371 bytes to copy. 1 copy tasklets on pool1@libb.
Progress:        -
                  1/1 copy tasklets completed on pool1@libb.

```

Result Summary: 100 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong pool, 0 not found, 0 too small, 0 too early.

(GLESM899I) All files have been successfully copied on pool1/libb.

Messages:

2019-01-07 11:58:34.231005 GLESM896I: Starting the stage 1 of 3 for migration task 18830 (qualifying the state of migration candidate files).

2019-01-07 11:58:37.133670 GLESM897I: Starting the stage 2 of 3 for migration task 18830 (copying the files to 1 pools).

2019-01-07 11:59:30.564568 GLESM898I: Starting the stage 3 of 3 for migration task 18830 (changing the state of files on disk).

2019-01-07 11:59:34.123713 GLESL038I: Migration result: 100 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too early for migration.

--- Subtask(level 1) Info ---

Task ID: 18831

Task Type: copy_replica

Status: completed

Result: succeeded

Accepted Time: Mon Jan 7 11:58:37 2019 (-0700)

Started Time: Mon Jan 7 11:58:37 2019 (-0700)

Completed Time: Mon Jan 7 11:59:30 2019 (-0700)

Workload: 7545870371 bytes to copy. 1 copy tasklets on pool1@libb.

Progress: 1/1 copy tasklets completed on pool1@libb.

Result Summary: (GLESM899I) All files have been successfully copied on pool1/libb.

Messages:

2019-01-07 11:58:37.196346 GLESM825I: The Copy tasklet (0x3527880) will be dispatched on drive 1068045923 from the write queue (tape=D00369L5).

2019-01-07 11:58:37.202930 GLESM031I: A list of 100 file(s) has been added to the migration and recall queue.

2019-01-07 11:59:30.556694 GLESM134I: Copy result: 100 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too early for migration.

2019-01-07 11:59:30.556838 GLESM899I: All files have been successfully copied on pool1/libb.

The other common use will be to show the file results of each individual file from any premigration, migration or recall tasks. With this information, the administrator can determine which files were successful and which files failed including the error code (reason) and date/time the error occurred. The administrator can quickly determine which files failed, and take corrective actions including resubmission for those failed files. Example 6-35 shows the completed task results.

Example 6-35 Completed task results

```
[root@kyoto prod]# eeadm task show 18835 -r
Result  Failure Code  Failed time  Node File name
Success - - - -
/ibm/gpfs/prod/LTFS_EE_FILE_1iv1y5z3SkWhzD_48fp5.bin
Success - - - -
/ibm/gpfs/prod/LTFS_EE_FILE_eMBse7fTbESNileaQnhHvOK6V621WuTxs_zQ1.bin
Success - - - -
/ibm/gpfs/prod/LTFS_EE_FILE_ZYyoDMD3WnwRyN50j59wJxjARox66YKq1OMw_NsE.bin
```

```

Success - - -
/ibm/gpfs/prod/LTFS_EE_FILE_h7SaXit10f9vrUo_3yYT.bin
Success - - -
/ibm/gpfs/prod/LTFS_EE_FILE_RPrsQ0xxKAu3nJ9_xResu.bin
Success - - -
/ibm/gpfs/prod/LTFS_EE_FILE_CLS7aXD9YBwUNHhhfL1FSaVf4q7eMBtwHYVnMpcAWR6XwnPYL_rsQ.
bin
Success - - -
/ibm/gpfs/prod/LTFS_EE_FILE_XEvHOLABXWx4CZY7cmwnvyT9W5i5uu_bUvNC.bin
Success - - -
/ibm/gpfs/prod/LTFS_EE_FILE_3_wjhe.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_7QVSfmURbF1kQZJAYNv1Px82frnUe1fyKSH0c7ZqJNs1_swA.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_8hB1_B.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_n19T7Y4Z_1.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_W2E77x4f3CICypMbLewnUzQq91hDojdQVJHymiXZuHMJKPY_X.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_s0yPUWwKaMu3Y_VzS.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_tc5xwE1J1SM_x.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_y1_73YEI.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_UR65_nyJ.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_DXhfSFK8N2TrN7bhr0tNfNARwT3K1tZbp5SmBb8RbK_d.bin
Fail GLESC012E 2019/01/07T12:06:49 1
/ibm/gpfs/prod/LTFS_EE_FILE_ZRthniqdYS70yoblCUKc9uz9NECTtLnC81NODxsQhj_QEWas.bin

```

For more information about the task status, see “User Task Reporting” on page 7.

6.6 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:

- ▶ `eeadm pool create`
- ▶ `eeadm pool delete`
- ▶ `eeadm tape assign`
- ▶ `eeadm tape unassign`
- ▶ `eeadm drive assign`
- ▶ `eeadm drive unassign`
- ▶ `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 **eeadm tape assign** 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.

6.7 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.

6.7.1 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 daemons
- ▶ IBM Spectrum Scale (GPFS) daemon called `mmfsd`
- ▶ IBM Spectrum Protect for Space Management (HSM) recall daemon called `dsmrecalld`
- ▶ `Rpcbind`
- ▶ `Rsyslog`
- ▶ SSH

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 attempt to restart the MMM service on the current node is made and if it fails a 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. If GPFS, HSM, `rpcbind`, or `rsyslog` are having problems, the issues can be viewed by using the `eeadm node list` command.

Example 6-36 shows output from running `eeadm node list` 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 6-36 The eeadm info nodes output

```
[root@daito ~]# eeadm node list
```

```
Spectrum Archive EE service (MMM) for library libb fails to start or is not running on daito Node ID:1
```

```
Problem Detected:  
Node ID Error Modules
```

```
1 MMM; rpcbind;
```

Node ID	State	Node IP	Drives	Ctrl Node	Library	Node Group	Host Name
2	available	9.11.244.62	2	yes(active)	liba	G0	nara

In addition to the automatic failover, there is also an option to perform a manual failover by running the **eeadm cluster 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 6-37 shows output from running the **eeadm cluster failover** command.

Example 6-37 The eeadm failover command

```
[root@kyoto ~]# eeadm cluster failover
2019-01-07 13:22:12 GLESL659I: Failover on library lto is requested to start.
          Use the "eeadm node list" command to see if the control node is
switched over.
```

6.8 Tape library management

This section describes how to use **eeadm** commands to add and remove tape drives and tape cartridges from your LTFS library.

6.8.1 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 LTFS creates the required LTFS 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 **eeadm library rescan** command. First, run the **eeadm tape list** command. If the tape cartridges are missing, run the **eeadm library rescan** 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 **eeadm tape list** 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 **eeadm tape list** command to see whether your tape appears in the list, as shown in Example 6-38 on page 155. In this example, the **-l** option is used to limit the tapes to one tape library.

Example 6-38 Run the eeadm tape list command to check whether a tape cartridge must be synchronized

```
[root@mikasa1 ~]# eeadm tape list -l lib_saitama
Tape ID  Status  State      Usable(GiB)  Used(GiB)  Available(GiB)  Reclaimable%  Pool  Library  Location  Task ID
JCA224JC ok      appendable  6292         0           6292           0%    pool1  lib_saitama  homeslot  -
JCC093JC ok      appendable  6292         496         5796           0%    pool1  lib_saitama  homeslot  -
JCB745JC ok      appendable  6292         0           6292           0%    pool2  lib_saitama  homeslot  -
```

Tape cartridge JCA561JC is not in the list.

3. Because tape cartridge JCA561JC is not in the list, synchronize the data in the IBM Spectrum Archive EE inventory with the tape library by running the **eeadm library rescan** command, as shown in Example 6-39.

Example 6-39 Synchronize the tape

```
[[root@mikasa1 ~]# eeadm library rescan
2019-01-07 13:41:34 GLESL036I: library rescan for lib_saitama completed.
(id=ebc1b34a-1bd8-4c86-b4fb-bee7b60c24c7, ip_addr=9.11.244.44)
2019-01-07 13:41:47 GLESL036I: library rescan for lib_mikasa completed.
(id=8a59cc8b-bd15-4910-88ae-68306006c6da, ip_addr=9.11.244.42)
```

4. Repeating the **eeadm tape list** command shows that the inventory was corrected, as shown in Example 6-40.

Example 6-40 Tape cartridge JCA561JC is synchronized

```
[root@mikasa1 ~]# eeadm tape list -l lib_saitama
Tape ID  Status  State      Usable(GiB)  Used(GiB)  Available(GiB)  Reclaimable%  Pool  Library  Location  Task ID
JCA224JC ok      appendable  6292         0           6292           0%    pool1  lib_saitama  homeslot  -
JCC093JC ok      appendable  6292         496         5796           0%    pool1  lib_saitama  homeslot  -
JCB745JC ok      appendable  6292         0           6292           0%    pool2  lib_saitama  homeslot  -
JCA561JC ok      unassigned  0            0           0              0%    -      lib_saitama  ieslot    -
```

5. If necessary, move the tape cartridge from the I/O station to a storage slot by running the **eeadm tape move** command (see Example 6-41) with the **-L homeslot** option. The example also requires a **-l** option because of multiple tape libraries.

Example 6-41 Move tape to homeslot

```
[root@mikasa1 ~]# eeadm tape move JCA561JC -L homeslot -l lib_saitama
2019-01-07 14:02:36 GLESL700I: Task tape_move was created successfully, task id is 6967.
2019-01-07 14:02:50 GLESL103I: Tape JCA561JC 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 **eeadm tape 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 **eeadm tape assign** 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 **eeadm library rescan** command, as shown in Example 6-39 on page 155.
3. If necessary, move the tape cartridge from the I/O station to a storage slot by running the **eeadm tape move** command with the **-L homeslot** option, as shown in Example 6-41 on page 155.
4. The **eeadm tape assign** command automatically formats tapes when assigning them to a pool. Use the **-f** or **--format** option only when the user is aware that the tapes still contain files and are no longer needed. For example, Example 6-42 shows the output of the **eeadm tape assign** command.

Example 6-42 Format a scratch tape

```
[root@mikasa1 ~]# eeadm tape assign JCA561JC -p pool1 -l lib_saitama
2019-01-07 14:06:23 GLESL700I: Task tape_assign was created successfully, task
id is 6968.
2019-01-07 14:11:14 GLESL087I: Tape JCA561JC successfully formatted.
2019-01-07 14:11:14 GLESL360I: Assigned tape JCA561JC to pool pool1
successfully.
```

Note: Media optimization will run on the each initial assignment starting from LTO 9, which may take some more time. Refer to 7.29, “LTO 9 Media Optimization” on page 272 for more information.

For more information about other formatting options, see 6.8.3, “Formatting tape cartridges” on page 158.

6.8.2 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 6-43.

Example 6-43 Error message when removing a tape cartridge from a pool with migrated/saved files

```
[root@mikasa1 ~]# eeadm tape unassign JCC093JC -p pool1 -l lib_saitama
2019-01-07 15:10:36 GLESL700I: Task tape_unassign was created successfully, task
id is 6970.
2019-01-07 15:10:36 GLESL357E: Tape JCC093JC has migrated files or saved files. It
has not been unassigned 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 6-44.

Example 6-44 Remove an empty tape cartridge from one tape cartridge pool and add it to another

```
[root@mikasa1 ~]# eeadm tape unassign JCA561JC -p pool1 -l lib_saitama
2019-01-07 15:11:26 GLESL700I: Task tape_unassign was created successfully, task id is
6972.
2019-01-07 15:11:26 GLESM399I: Removing tape JCA561JC from pool pool1 (Normal).
2019-01-07 15:11:26 GLESL359I: Unassigned tape JCA561JC from pool pool1 successfully.

[root@mikasa1 ~]# eeadm tape assign JCA561JC -p pool2 -l lib_saitama
2019-01-07 15:12:10 GLESL700I: Task tape_assign was created successfully, task id is 6974.
2019-01-07 15:16:56 GLESL087I: Tape JCA561JC successfully formatted.
2019-01-07 15:16:56 GLESL360I: Assigned tape JCA561JC to pool pool2 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 6.17, “Reclamation” on page 200.

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 6-45. 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 6-45 Move a tape cartridge from a tape drive to its homeslot

```
[root@kyoto prod]# eeadm tape move D00369L5 -p pool1 -L homeslot
2019-01-07 15:47:26 GLESL700I: Task tape_move was created successfully, task id is
18843.
2019-01-07 15:49:14 GLESL103I: Tape D00369L5 is moved successfully.
```

Moving to the I/O station

The command that is shown in Example 6-46 moves a tape cartridge to the `ieslot` (I/O station). This might be required when tape cartridges are exported, offline, or unassigned.

Example 6-46 Move a tape cartridge to the ieslot after an offline operation

```
[root@mikasa1 ~]# eeadm tape offline JCA561JC -p pool2 -l lib_saitama
2019-01-07 15:50:17 GLESL700I: Task tape_offline was created successfully, task id
is 6976.
2019-01-07 15:50:17 GLESL073I: Offline export of tape JCA561JC has been requested.
2019-01-07 15:51:51 GLESL335I: Updated offline state of tape JCA561JC to offline.

[root@mikasa1 ~]# eeadm tape move JCA561JC -p pool2 -l lib_saitama -L ieslot
2019-01-07 15:53:45 GLESL700I: Task tape_move was created successfully, task id is
6978.
2019-01-07 15:53:45 GLESL103I: Tape JCA561JC 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 or offlining 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.

6.8.3 Formatting tape cartridges

This section describes how to format a medium in the library for the IBM Spectrum Archive EE. To format a scratch tape use the **eeadm tape assign** command, and only use the **-f/--format** option when the user no longer requires access to the data on the tape.

If the tape cartridge is already formatted for IBM Spectrum Archive EE and contains file objects, the format fails, as shown in Example 6-47.

Example 6-47 Format failure

```
[root@kyoto prod]# eeadm tape assign 1FB922L5 -p pool2
2019-01-08 08:29:08 GLESL700I: Task tape_assign was created successfully, task id
is 18850.
2019-01-08 08:30:21 GLESL138E: Failed to format the tape 1FB922L5, because it is
not empty.
```

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 was previously formatted for LTFS and has no data written.
- ▶ The medium has an invalid label.
- ▶ Labels in both partitions do not have the same value.

If none of these conditions are found, the format fails. If the format fails because files exist on the tape, the user should add the tape to their designated pool by using the **eeadm tape import** command. If the user no longer requires what is on the tape, the **-f/--format** option can be added to the **eeadm tape assign** command to force a format.

Example 6-48 shows a tape cartridge being formatted by using the **-f** option.

Example 6-48 Forced format

```
[root@kyoto prod]# eeadm tape assign 1FB922L5 -p pool2 -f
2019-01-08 08:32:42 GLESL700I: Task tape_assign was created successfully, task id
is 18852.
2019-01-08 08:35:08 GLESL087I: Tape 1FB922L5 successfully formatted.
2019-01-08 08:35:08 GLESL360I: Assigned tape 1FB922L5 to pool pool2 successfully.
```

Multiple tape cartridges can be formatted by specifying multiple tape VOLSERS. Example 6-49 shows three tape cartridges that are formatted sequentially or simultaneously.

Example 6-49 Format multiple tape cartridges

```
[root@mikasa1 ~]# eeadm tape assign JCC075JC JCB610JC JCC130JC -p pool1 -l
lib_saitama
2019-01-08 09:25:32 GLESL700I: Task tape_assign was created successfully, task id
is 6985.
2019-01-08 09:30:26 GLESL087I: Tape JCB610JC successfully formatted.
2019-01-08 09:30:26 GLESL360I: Assigned tape JCB610JC to pool pool1 successfully.
2019-01-08 09:30:33 GLESL087I: Tape JCC130JC successfully formatted.
```

```
2019-01-08 09:30:33 GLESL360I: Assigned tape JCC130JC to pool pool1 successfully.
2019-01-08 09:30:52 GLESL087I: Tape JCC075JC successfully formatted.
2019-01-08 09:30:52 GLESL360I: Assigned tape JCC075JC to pool pool1 successfully.
```

When multiple format tasks are submitted, IBM Spectrum Archive EE uses all available drives with the 'g' drive attribute for the format tasks, 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 uses the **-E** option, to the **eeadm tape unassign** command. When ran, this command checks whether the tape cartridge contains any active data. For example, if all files that were migrated on the tape has been made resident, there will be no active data on tape. If there is no active data, or if all 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 ran, 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.
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 6-50 shows the output of the **eeadm tape unassign -E** command with a tape which contains files but all files on the gpfs file system deleted.

Example 6-50 Removing tape cartridge from pool with active file check

```
[root@mikasa1 prod]# eeadm tape unassign JCB350JC -p test2 -l lib_saitama -E
2019-01-08 10:20:07 GLESL700I: Task tape_unassign was created successfully,
task id is 7002.
2019-01-08 10:20:09 GLESM399I: Removing tape JCB350JC from pool test2
(Empty_Force).
2019-01-08 10:20:09 GLESL572I: Unassign tape JCB350JC from pool test2
successfully. Format the tape when assigning it back to a pool.
```

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.

Example 6-51 shows the output of the **eeadm tape unassign -E** command with a tape which contains active files.

Example 6-51 Tape cartridges containing inactive data are removed from the pool

```
[root@mikasa1 prod]# eeadm tape unassign JCB350JC -p test2 -l lib_saitama -E
2019-01-08 10:17:15 GLESL700I: Task tape_unassign was created successfully,
task id is 6998.
2019-01-08 10:17:16 GLESL357E: Tape JCB350JC has migrated files or saved files.
It has not been unassigned from the pool.
```

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.

6.8.4 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 task is in-progress, the tape is unloaded automatically when the task completes.

After mounting the library, the user can run **eeadm** 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 **eeadm drive unassign** command, as shown in Example 6-52. A medium in the tape drive is automatically moved to the home slot (if one exists).

Example 6-52 Remove a tape drive

```
[root@yurakucho ~]# eeadm drive list -l lib1
Drive S/N  Status  State      Type Role  Library  Node ID  Tape  Node Group
Task ID
00078DF0CF ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0D3 ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0DC ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0E2 ok      not_mounted LT08  mrg  lib1    1        -    GO -
1013000508 ok      not_mounted LT08  mrg  lib1    1        -    GO -
```

```
[root@yurakucho ~]# eeadm drive unassign 00078DF0E2 -l lib1
2021-03-06 06:48:13 GLESL700I: Task drive_unassign was created successfully,
task ID is 2916.
2021-03-06 06:48:13 GLESL817I: Disabling drive 00078DF0E2.
2021-03-06 06:48:14 GLESL813I: Drive 00078DF0E2 is disabled successfully.
2021-03-06 06:48:14 GLESL819I: Unassigning drive 00078DF0E2.
2021-03-06 06:48:14 GLESL121I: Drive 00078DF0E2 is unassigned successfully.
```

```
[root@yurakucho ~]# eeadm drive list -l lib1
Drive S/N  Status  State      Type Role  Library  Node ID  Tape  Node Group
Task ID
00078DF0CF ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0D3 ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0DC ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0E2 info    unassigned  NONE  ---  lib1    -        -    - -
1013000508 ok      not_mounted LT08  mrg  lib1    1        -    GO -
```

2. Physically remove the tape drive from the tape library.

For more information about how to remove tape drives, see the IBM Documentation website for your IBM tape library.

6.8.5 Adding tape drives

Add the tape drive to the LTFS inventory by running the **eeadm drive assign** command, as shown in Example 6-53.

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 tasks on the task queue can be run on that drive. The drive attributes can be specified using the `-r` option and must be a decimal number or the combination of the letters `m`, `r`, and `g`.

In Example 6-53, 6 is the logical OR of `migrate(4)` and `recall(2)`, so migration tasks and recall tasks can be performed on this drive. For more information, see 6.20, “Drive Role settings for task assignment control” on page 207.

The node ID is required for the **eeadm drive assign** command.

Example 6-53 Add a tape drive

```
[root@yurakucho ~]# eeadm drive assign 00078DF0E2 -r 6 -n 1 -l lib1
2021-03-06 06:53:06 GLESL700I: Task drive_assign was created successfully, task ID
is 2918.
2021-03-06 06:53:06 GLESL818I: Assigning drive 00078DF0E2 to node 1.
2021-03-06 06:53:06 GLESL119I: Drive 00078DF0E2 assigned successfully.
2021-03-06 06:53:06 GLESL816I: Enabling drive 00078DF0E2.
2021-03-06 06:53:06 GLESL812I: Drive 00078DF0E2 is enabled successfully.
[root@yurakucho ~]# eeadm drive list -l lib1
Drive S/N  Status  State      Type Role  Library  Node ID  Tape  Node Group
Task ID
00078DF0CF ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0D3 ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0DC ok      not_mounted LT08  mrg  lib1    1        -    GO -
00078DF0E2 ok      not_mounted LT08  mr-  lib1    1        -    GO -
1013000508 ok      not_mounted LT08  mrg  lib1    1        -    GO -
```

6.9 Tape storage pool management

This section describes how to use the `eeadm pool` command to manage tape cartridge pools with IBM Spectrum Archive EE.

Permissions: Managing tape cartridge pools by running the `eeadm 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 6.8.2, “Moving tape cartridges” on page 156.
- ▶ Multiple tasks 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 7.2, “Maximizing migration performance with redundant copies” on page 235.
- ▶ 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 `eeadm tape export` command, they are removed from their tape cartridge pool and the files are not accessible for recall.

6.9.1 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 `eeadm pool` command. For example, the command that is shown in Example 6-54 creates the tape cartridge pool named `MPEGpool`.

Example 6-54 Create a tape cartridge pool

```
[root@kyoto prod]# eeadm pool create 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 `eeadm pool list` command.

6.9.2 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 `eeadm pool delete` command. For example, the command in Example 6-55 on page 163 deletes the tape cartridge pool that is named MPEGpool.

Example 6-55 Delete a tape cartridge pool

```
[root@kyoto prod]# eeadm pool delete 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.

To allow the deletion of the tape cartridge pool, you must remove all tape cartridges from it by running the `eeadm tape unassign` command, as described in 6.8.2, “Moving tape cartridges” on page 156.

6.10 Pool capacity monitoring

IBM Spectrum Archive EE allows users the ability to automatically monitor the capacity of designated pools for low space threshold or if a pool has ran out of space due to a migration failure. This feature uses SNMP traps to inform administrators of such occurrences. 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, use the `eeadm pool set` command with the `lowspacewarningenable yes` and the `nospacewarningenable yes` option, then set a threshold limit for the pool with the `lowspacewarningthreshold` option. The value for the `lowspacewarningthreshold` must be an integer and is in TiB. The pool capacity monitor thread checks each set pool for a low capacity every 30 minutes, and sends a trap every 24 hours.

Run the **eeadm pool show <pool_name>** command to view the current attributes for your pools. Example 6-56 shows the output of the attributes of a pool from the **eeadm pool show** command.

Example 6-56 The eeadm pool show command

```
[root@l1fseesrv1 ~]# eeadm pool show pool1
Attribute          Value
poolname           pool1
poolid             aefeea24-661e-48ba-8abd-d4948b020d74
devtype            LT0
mediarestriction  none
format             Not Applicable (0xFFFFFFFF)
worm               no (0)
nodegroup          G0
fillpolicy         Default
owner              System
mountlimit         0
lowspacewarningenable yes
lowspacewarningthreshold 0
nospacewarningenable yes
mode               normal
```

By default, the **lowspacewarningenable** and **nospacewarningenable** attributes are set to **yes**, and **lowspacewarningthreshold** is set to 0, which indicates no traps will be sent for a pool with low space. 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 **eeadm pool set <pool_name> -a <attribute> -v <value>** command.

Example 6-57 shows the output of modifying pool1's **lowspacewarningthreshold** attribute to 2 TiB.

Example 6-57 Output of modifying pool1's lowspacewarningthreshold attribute

```
[root@kyoto prod]# eeadm pool set pool1 -a lowspacewarningthreshold -v 30

[root@kyoto prod]# eeadm pool show pool1
Attribute          Value
poolname           pool1
poolid             93499f33-c67f-4aae-a07e-13a56629b057
devtype            LT0
mediarestriction  none
format             Not Applicable (0xFFFFFFFF)
worm               no (0)
nodegroup          G0
fillpolicy         Default
owner              System
mountlimit         0
lowspacewarningenable yes
lowspacewarningthreshold 30
nospacewarningenable yes
mode               normal
```

With `lowspacewarningthreshold` set to 30 TiB, when `pool1`'s capacity drops below 30 TiB, a trap will be sent to the user when the next check cycle occurs. Example 6-58 shows the traps generated when `pool1`'s capacity drops below 30 TiB.

Example 6-58 Traps sent when pool capacity is below the set threshold

```
2018-11-26 09:08:42 tora.tuc.stglabs.ibm.com [UDP:
[9.11.244.63]:60811->[9.11.244.63]:162]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (147206568) 17 days, 0:54:25.68
SNMPv2-MIB::snmpTrapOID.0 = OID: IBMSA-MIB::ibmsaWarnM609PoolLowSpace
IBMSA-MIB::ibmsaMessageSeverity.0 = INTEGER: warning(40)   IBMSA-MIB::ibmsaJob.0
= INTEGER: other(7)   IBMSA-MIB::ibmsaEventNode.0 = STRING:
"tora.tuc.stglabs.ibm.com"   IBMSA-MIB::ibmsaMessageText.0 = STRING: "GLES609W:
Pool space is going to be small, library: lib_tora, pool: pool1, available
capacity: 23.6(TiB), threshold: 30(TiB)"
```

6.11 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 LTFS 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 file system space 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 LTFS tape tier. For each specified GPFS file, the file content, GPFS path, and user-defined extended attributes are stored in LTFS so that they can be re-created at import. Empty GPFS directories or files are not migrated.

In addition, it is possible to migrate an arbitrary list of files directly by running the `eeadm 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 IBM Spectrum Protect 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 LTFS node. For more information, see 6.2.4, “IBM Spectrum Archive EE” on page 140.
- ▶ Ensure that one or more storage pools are created and each has one or more assigned tapes. For more information, see 6.9.1, “Creating tape cartridge pools” on page 162.
- ▶ Ensure that space management is turned on. For more information, see 6.2.3, “Hierarchical Space Management” on page 138.
- ▶ 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 `eeadm 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. Tasks are put on the queue and the scheduler looks at the queue to decide which tasks 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.

IBM Spectrum Archive EE internally groups files into file lists and schedules these lists on the task 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 task queue for available drives.

For more information about performance references, see 3.7.4, “Performance” on page 65.

Note: Files that are recently created need to wait two minutes before being migrated. Otherwise, the migrations will fail.

Example 6-59 shows the output of running the `mmapplypolicy` command that uses a policy file called `sample_policy.txt`.

Example 6-59 Output of the mmapplypolicy command

```
[root@kyoto ~]# mmapplypolicy /ibm/gpfs/prod -P sample_policy.txt
[I] GPFS Current Data Pool Utilization in KB and %
Pool_Name          KB_Occupied      KB_Total  Percent_Occupied
system             1489202176      15435038720  9.648192033%
[I] 682998 of 16877312 inodes used: 4.046841%.
[I] Loaded policy rules from cmt_policy.txt.
Evaluating policy rules with CURRENT_TIMESTAMP = 2019-01-08@22:26:33 UTC
Parsed 3 policy rules.
```

```
RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'
```

```
RULE EXTERNAL POOL 'md1'
EXEC '/opt/ibm/ltfsee/bin/eedm'
OPTS '-p pool3@libb'
```

```
RULE 'LTFS_EE_FILES' MIGRATE FROM POOL 'system'
TO POOL 'md1'
WHERE FILE_SIZE > 0
AND NAME LIKE '%.bin'
AND PATH_NAME LIKE '/ibm/gpfs/%'
/*AND (is_cached)
AND NOT (is_dirty)*/
AND ((NOT MISC_ATTRIBUTES LIKE '%M%')
OR (MISC_ATTRIBUTES LIKE '%M%' AND MISC_ATTRIBUTES NOT LIKE '%V%'))
)
AND NOT ((NAME = 'dsmerror.log' OR NAME LIKE '%DS_Store%'))
)
AND (NOT ((
```

```

        FALSE
        OR PATH_NAME LIKE '/ibm/gpfs/.ltfsee/%'
        OR PATH_NAME LIKE '%/.SpaceMan/%'
    )
) OR ((
    FALSE
    )
))

```

[I] 2019-01-08@22:26:33.153 Directory entries scanned: 22.
[I] Directories scan: 21 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.

[I] 2019-01-08@22:26:33.157 Sorting 22 file list records.
[I] Inodes scan: 21 files, 1 directories, 0 other objects, 0 'skipped' files and/or errors.

[I] 2019-01-08@22:26:33.193 Policy evaluation. 22 files scanned.
[I] 2019-01-08@22:26:33.197 Sorting 10 candidate file list records.
[I] 2019-01-08@22:26:33.197 Choosing candidate files. 10 records scanned.

[I] Summary of Rule Applicability and File Choices:

Rule#	Hit_Cnt	KB_Hit	Chosen	KB_Chosen
KB_I11	Rule			
0	10	953344	10	953344
0	RULE 'LTFS_EE_FILES' MIGRATE FROM POOL 'system' TO POOL 'md1' WHERE(.)			

[I] Filesystem objects with no applicable rules: 12.

[I] GPFS Policy Decisions and File Choice Totals:
Chose to migrate 953344KB: 10 of 10 candidates;
Predicted Data Pool Utilization in KB and %:

Pool_Name	KB_Occupied	KB_Total	Percent_Occupied
system	1488248832	15435038720	9.642015540%

2019-01-08 15:26:33 GLESL700I: Task migrate was created successfully, task id is 18860.

2019-01-08 15:26:33 GLESM896I: Starting the stage 1 of 3 for migration task 18860 (qualifying the state of migration candidate files).

2019-01-08 15:26:33 GLESM897I: Starting the stage 2 of 3 for migration task 18860 (copying the files to 1 pools).

2019-01-08 15:27:18 GLESM898I: Starting the stage 3 of 3 for migration task 18860 (changing the state of files on disk).

2019-01-08 15:27:19 GLESL038I: Migration result: 10 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for migration, 0 too early for migration.

[I] 2019-01-08@22:27:20.722 Policy execution. 10 files dispatched.

[I] A total of 10 files have been migrated, deleted or processed by an EXTERNAL EXEC/script;

0 'skipped' files and/or errors.

6.11.1 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 `eeadm recall` command with the `--resident` option. For more information about creating replicas, see 6.11.4, “Replicas and redundant copies” on page 177. Recalling the file into resident state invalidates the LTFS copy from the reconcile and export perspective.

6.11.2 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 LIST, 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 6-60, 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 7.10, “Use cases for `mmapplypolicy`” on page 244.

Example 6-60 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*, which is available at [IBM Documentation](#)

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 6.11.3, “Manual migration” on page 173.

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/eadm`) 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 6-61.

Example 6-61 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 6-62. 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 6-62 List an IBM Spectrum Scale policy

```
[root@ltfs97]# 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/eeadm'
OPTS -p 'pool1@libb'
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 file system 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. This occurs in cases where migrations are continuously called while the past migrations has not finished, because this state may appear for the values to be still over the threshold.

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 `<GPFS mount point>/.ltfsee/`.

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 6-63 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/eeadm`), 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 6-63 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/eeadm'
OPTS -p 'Tapepool1@liba' /* 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 6-63, 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 [IBM Documentation](#).

Also, see 7.4, “Setting `mmapplypolicy` options for increased performance” on page 237.

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 [IBM Documentation](#).

6.11.3 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 `eeadm` 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*, which is available at [IBM Documentation](#).

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 making any data movements. The `-f /tmp` option specifies a directory or file to which each migration rule is output. This option is helpful when dealing with many 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 -P test_policy.txt -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 6-64 shows a `mmapplypolicy` command that tests, but does not apply, an IBM Spectrum Scale policy by using the `testpolicy` policy file.

Example 6-64 Test an IBM Spectrum Scale policy

```
[root@kyoto ~]# mmapplypolicy /ibm/gpfs/prod -P sample_policy.txt -I test
[I] GPFS Current Data Pool Utilization in KB and %
Pool_Name          KB_Occupied      KB_Total  Percent_Occupied
system            1489192960      15435038720  9.648132324%
[I] 683018 of 16877312 inodes used: 4.046960%.
[I] Loaded policy rules from cmt_policy.txt.
Evaluating policy rules with CURRENT_TIMESTAMP = 2019-01-08@22:34:42 UTC
Parsed 3 policy rules.

RULE 'SYSTEM_POOL_PLACEMENT_RULE' SET POOL 'system'

RULE EXTERNAL POOL 'md1'
EXEC '/opt/ibm/ltfsee/bin/eedm'
OPTS '-p copy_ltfsm11@lib_ltfsm11'

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/.ltfsee/%' OR PATH_NAME LIKE '%/.SpaceMan/%'))

[I] 2019-01-08@22:34:42.327 Directory entries scanned: 22.
[I] Directories scan: 21 files, 1 directories, 0 other objects, 0 'skipped' files
and/or errors.
[I] 2019-01-08@22:34:42.330 Sorting 22 file list records.
[I] Inodes scan: 21 files, 1 directories, 0 other objects, 0 'skipped' files
and/or errors.
[I] 2019-01-08@22:34:42.365 Policy evaluation. 22 files scanned.
[I] 2019-01-08@22:34:42.368 Sorting 10 candidate file list records.
[I] 2019-01-08@22:34:42.369 Choosing candidate files. 10 records scanned.
[I] Summary of Rule Applicability and File Choices:
  Rule#      Hit_Cnt      KB_Hit      Chosen      KB_Chosen
KB_I11      Rule
  0          10          947088      10          947088
0    RULE 'LTFS_EE_FILES' MIGRATE FROM POOL 'system' TO POOL 'md1' WHERE(.)

[I] Filesystem objects with no applicable rules: 12.

[I] GPFS Policy Decisions and File Choice Totals:
  Chose to migrate 947088KB: 10 of 10 candidates;
  Predicted Data Pool Utilization in KB and %:
Pool_Name          KB_Occupied      KB_Total  Percent_Occupied
system            1488245872      15435038720  9.641996363%
```

The policy in Example 6-64 on page 175 is configured to select files that have the file extension `.IMG` for migration to the IBM Spectrum Archive EE tape cartridge pool named `copy_1tfsml1` in library name `lib_1tfsml1` if the usage of the `/ibm/gpfs` file system exceeds 50% for any `.IMG` file that exceeds 5 MB.

Using eeadm

The `eeadm 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`

Example 6-65 shows the output of running such a `migrate` command.

Example 6-65 Manual migration by using a scan result file

```
[root@kyoto prod]# eeadm migrate gpfs-scan.txt -p MPEGpool
2019-01-08 15:40:44 GLESL700I: Task migrate was created successfully, task id is
18864.
2019-01-08 15:40:44 GLESM896I: Starting the stage 1 of 3 for migration task 18864
(qualifying the state of migration candidate files).
2019-01-08 15:40:44 GLESM897I: Starting the stage 2 of 3 for migration task 18864
(copying the files to 1 pools).
2019-01-08 15:40:56 GLESM898I: Starting the stage 3 of 3 for migration task 18864
(changing the state of files on disk).
2019-01-08 15:40:57 GLESL038I: Migration result: 10 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 `eeadm` 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 `eeadm` as an external program. In this case, the full path to `eeadm` 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 6-66 shows how to start the **crond** process and create a single cron job that performs migrations every six hours.

Example 6-66 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
```

6.11.4 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 copy 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 copy 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 copy 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 task 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 **eeadm recall** command, an available copy is selected from the available replicas and the recall task 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 **-1** 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 where a migrated file is truncated to 0, 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 **eeadm migrate** (or **eeadm 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 6-67.

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 6-67 shows the OPTS line in a policy file, which makes replicas of files in two tape cartridge pools in a single tape library.

Example 6-67 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 6.11.2, “Threshold-based migration” on page 168.

If you are running the **eeadm migrate** (or **eeadm premigrate**) command, a scan list file must be passed along with the designated pools to which the user wants to migrate the files. This process can be done by calling **eeadm migrate <inputfile> -p <list_of_pools> [OPTIONS]** (or **eeadm premigrate**).

Example 6-68 shows what the scan list looks like when selecting files to migrate.

Example 6-68 Example scan list file

```
[root@lufs97 /]# cat migrate.txt
-- /ibm/glues/document10.txt
-- /ibm/glues/document20.txt
```

Example 6-69 shows how one would run a manual migration using the **eeadm migrate** command on the scan list from Example 6-68 to two tapes.

Example 6-69 Creation of replicas during migration

```
[root@kyoto prod]# eeadm migrate migrate.txt -p pool1 pool2
2019-01-08 15:52:39 GLESL700I: Task migrate was created successfully, task id is
18866.
2019-01-08 15:52:39 GLESM896I: Starting the stage 1 of 3 for migration task 18866
(qualifying the state of migration candidate files).
```

```

2019-01-08 15:52:39 GLESM897I: Starting the stage 2 of 3 for migration task 18866
(copying the files to 2 pools).
2019-01-08 15:53:26 GLESM898I: Starting the stage 3 of 3 for migration task 18866
(changing the state of files on disk).
2019-01-08 15:53:26 GLESL038I: 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 3.7, “Sizing and settings” on page 61.

By monitoring the **eadm task list** and the **eadm task show** command as the migration is running, you can observe the status of the migration task, as shown in Example 6-70.

Example 6-70 Migration task status

```

[root@mikasa1 ~]# eadm task list
TaskID Type Priority Status #DRV CreatedTime(-0700) StartedTime(-0700)
 7014 migrate M waiting 0 2019-01-08_16:11:28 2019-01-08_16:11:28

```

```

[root@mikasa1 ~]# eadm task list
TaskID Type Priority Status #DRV CreatedTime(-0700) StartedTime(-0700)
 7014 migrate M running 0 2019-01-08_16:11:28 2019-01-08_16:11:28

```

```

[root@mikasa1 ~]# eadm task show 7014
=== Task Information ===
Task ID: 7014
Task Type: migrate
Command Parameters: eadm migrate mig -p pool2@lib_saitama test2@lib_saitama
Status: running
Result: -
Accepted Time: Tue Jan 8 16:11:28 2019 (-0700)
Started Time: Tue Jan 8 16:11:28 2019 (-0700)
Completed Time: -
In-use Resources: 0000078PG24E(JCB745JC):pool2:G0:lib_saitama
Workload: 2 files. 2 replicas.
4566941 bytes to copy. 1 copy tasklets on pool2@lib_saitama.
4566941 bytes to copy. 1 copy tasklets on test2@lib_saitama.
Progress: -
0/1 copy tasklets completed on pool2@lib_saitama.
0/1 copy tasklets completed on test2@lib_saitama.
Result Summary: -

```

```

[root@mikasa1 ~]# eadm task show 7014
=== Task Information ===
Task ID: 7014
Task Type: migrate
Command Parameters: eadm migrate mig -p pool2@lib_saitama test2@lib_saitama
Status: completed
Result: succeeded
Accepted Time: Tue Jan 8 16:11:28 2019 (-0700)
Started Time: Tue Jan 8 16:11:28 2019 (-0700)

```

```

Completed Time:      Tue Jan  8 16:12:14 2019 (-0700)
Workload:           2 files. 2 replicas.
                   4566941 bytes to copy. 1 copy tasklets on pool2@lib_saitama.
                   4566941 bytes to copy. 1 copy tasklets on test2@lib_saitama.

Progress:          -
                  1/1 copy tasklets completed on pool2@lib_saitama.
                  1/1 copy tasklets completed on test2@lib_saitama.

Result Summary:    2 succeeded, 0 failed, 0 duplicate, 0 duplicate wrong pool,
0 not found, 0 too small, 0 too early.
                  (GLESM899I) All files have been successfully copied on
pool2/lib_saitama.
                  (GLESM899I) All files have been successfully copied on
test2/lib_saitama.

```

For more information and command syntax, see the **eeadm migrate** command in 6.11, “Migration” on page 165.

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 **eeadm tape export** or **eeadm tape offline**, depending on how the data should be kept. A second option would be to create the redundant copy in a different tape library.

6.11.5 Data Migration

Because of the need to upgrade tape generations or reuse tape cartridges, IBM Spectrum Archive EE allows users to specify to which tape cartridges they want their data to be migrated. Use the **eeadm tape datamigrate** command to perform pool-to-pool file migrations. 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. Example 6-71 shows how users can use the **eeadm tape datamigrate** command to migrate their data from an older generation pool to a newer one.

Example 6-71 Migrating data to newer generation tape pool

```

[root@mikasa1 ~]# eeadm tape datamigrate -p test3 -d test4 -l lib2
2021-03-24 04:23:29 GLESL700I: Task datamigrate was created successfully, task ID
is 1155.
2021-03-24 04:23:29 GLESR216I: Multiple processes started in parallel. The maximum
number of processes is unlimited.
2021-03-24 04:23:29 GLESL385I: Starting the "eeadm tape datamigrate" command by
the reclaim operation.
Processing 1 tapes in the following list of tapes from source pool test3. The
files on the tapes are moved to the tapes in target pool test4:
2021-03-24 04:23:29 GLESR212I: Source candidates: JCA561JC .

```

2021-03-24 04:23:29 GLESL677I: Files on tape JCA561JC will be copied to tape JCB370JC.
 2021-03-24 04:28:55 GLESL081I: Tape JCA561JC successfully reclaimed as datamigrate process, formatted, and unassigned from tape pool test3.
 2021-03-24 04:28:55 GLESL080I: Reclamation complete as the datamigrate command. 1 tapes reclaimed, 1 tapes unassigned from the tape pool.

For the **eedm tape datamigrate** command syntax, see “The eeadm <resource type> <action> --help command” on page 317.

Note: If the tapes are not specified for **eedm tape reclaim**, the tapes being reclaimed remain in the source pool. If the tapes are specified with the **eedm tape datamigrate** command, 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 and modifying the **media_restriction** or **format** type to make older generation media become **append_fenced**. By doing so, in addition to securing future migration tasks go to newer generation media this also 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 **media_restriction** attribute of the pool, the format type is also automatically updated to the highest generation drive available to the cluster. The **format** attribute is automatically updated only after each time the **media_restriction** is modified. If new drive and media generations are added to a cluster with **media_restriction** already set, users are expected to update **format** or **media_restriction** manually to support new media generations.

Example 6-72 on page 181 shows the automatic format update when changing the **media_restriction** attribute from JC to JE when there are TS1160, TS1155, and TS1150 drives.

Example 6-72 Updating pool media_restriction

```
[root@mikasa1 prod]# eeadm pool show test3 -l lib_saitama
Attribute                Value
poolname                 test3
poolid                   99096daa-0beb-4791-b24f-672804a56440
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@mikasa1 prod]# eeadm pool set test3 -l lib_saitama -a mediarestriction -v JE
```

```
[root@mikasa1 prod]# eeadm pool show test3 -l lib_saitama
Attribute          Value
poolname           test3
poolid             99096daa-0beb-4791-b24f-672804a56440
devtype            3592
mediarestriction   JE
format             60F (0x57)
worm               no (0)
nodegroup          G0
fillpolicy         Default
owner              System
mountlimit         0
lowspacewarningenable yes
lowspacewarningthreshold 0
nospacewarningenable yes
mode               normal
```

All tapes within test3 that fall under the restriction of JE media and 60F format type are all appendable tapes, and everything else is now **append_fenced**.

6.11.6 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 in “resident” state).

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 pool2. Before the migration attempt, tape JCB610JC, which is defined in a different tape cartridge pool (pool1), already contains three of the four files. The state of the files on these tape cartridges before the migration attempt is displayed by the **eeadm file state** command as shown in Example 6-73.

Example 6-73 Before migration

```
[root@mikasa1 prod]# eeadm file state *.bin
Name: /ibm/gpfs/prod/fileA.ppt
State: migrated
ID: 11151648183451819981-3451383879228984073-1939041988-3068866-0
Replicas: 1
Tape 1: JCB610JC@pool1@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/fileB.ppt
State: migrated
ID: 11151648183451819981-3451383879228984073-1844785692-3068794-0
Replicas: 1
Tape 1: JCB610JC@pool1@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/fileC.ppt
State: migrated
```

ID: 11151648183451819981-3451383879228984073-373707969-3068783-0
Replicas: 1
Tape 1: JCB610JC@pool1@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/fileD.ppt
State: resident

The attempt to migrate the files to a different tape cartridge pool produces the results that are shown in Example 6-74.

Example 6-74 Attempted migration of already migrated files

```
[root@mikasa1 prod]# eeadm migrate mig -p pool2@lib_saitama
2019-01-09 15:07:18 GLESL700I: Task migrate was created successfully, task id is
7110.
2019-01-09 15:07:18 GLESM896I: Starting the stage 1 of 3 for migration task 7110
(qualifying the state of migration candidate files).
2019-01-09 15:07:18 GLESM897I: Starting the stage 2 of 3 for migration task 7110
(copying the files to 1 pools).
2019-01-09 15:07:30 GLESM898I: Starting the stage 3 of 3 for migration task 7110
(changing the state of files on disk).
2019-01-09 15:07:31 GLESL159E: Not all migration has been successful.
2019-01-09 15:07:31 GLESL038I: 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 6-75 explain the reason for the failures.

Example 6-75 Migration errors reported in the IBM Spectrum Archive EE log file

```
2019-01-09T15:07:18.713348-07:00 saitama2 mmm[22592]: GLESM148E(00710): File
/ibm/gpfs/prod/fileA.ppt is already migrated and will be skipped.
2019-01-09T15:07:18.714413-07:00 saitama2 mmm[22592]: GLESM148E(00710): File
/ibm/gpfs/prod/fileB.ppt is already migrated and will be skipped.
2019-01-09T15:07:18.715196-07:00 saitama2 mmm[22592]: GLESM148E(00710): File
/ibm/gpfs/prod/fileC.ppt is already migrated and will be skipped.
```

The files on tape JCB610JC (fileA.ppt, fileB.ppt, and fileC.ppt) are already in storage pool pool1. Therefore, the attempt to migrate them to storage pool pool2 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 pool2 tape cartridge pool, the method that is described in 6.11.4, “Replicas and redundant copies” on page 177 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 6-63 on page 171.

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 10.2, “Formats for IBM Spectrum Scale to IBM Spectrum Archive EE migration” on page 325.

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 7.10, “Use cases for `mmapplypolicy`” on page 244.

6.12 Premigration

A premigrated file is a file that the content 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 `eadm 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 content of 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 `eadm 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 `eadm premigrate` or `mmapplypolicy`.
- ▶ Automatic premigration is available by running `eadm premigrate` through the `mmapplypolicy/mmaddcallback` command or a cron job.
- ▶ Migration hints and tips are applicable to premigration.

For the **eadm migrate** command, each migrate task 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.

6.12.1 Premigration with the **eadm premigrate** command

The **eadm premigrate** command is used to premigrate non-empty regular files to tape. The command syntax is the same as for the **eadm migrate** command. The following is an example of the syntax:

```
eadm premigrate <inputfile> -p <list_of_pools> [OPTIONS]
```

The **<input_file>** file includes the list of non-empty regular files to be premigrated. Each line of this file must end with either of the following format.

- ▶ Each line ends with **-- <filename>** with a space before and after the double dash (the file list record format of the **mmapplypolicy** command).
- ▶ Each line contains a file name with an absolute path or a relative path that is based on the working directory. This format is unavailable when you run the command with the **--mmbackup** option.

Specified files 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 commas.

6.12.2 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 **preigrate** 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 '1tfs'
```

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 '1tfs'
```

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 '1tfs' 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 7.10.2, “Creating active archive system policies” on page 245.

```
RULE 'premig4' MIGRATE FROM POOL 'system' THRESHOLD (0,100,0) TO POOL '1tfs' WHERE
(FILE_SIZE > 5242880)
```

6.13 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 **eeadm save** command directly with a scan list file
- ▶ An IBM Spectrum Scale policy with the **mmapp1ypolicy** command

6.13.1 Saving file system objects with the **eeadm save** command

The **eeadm save** command is used to save symbolic links, empty regular files, and empty directories to tape. The command syntax is the same as the **eeadm migrate** or **eeadm premigrate** commands. The following is the syntax of the **eeadm save** command:

```
eeadm save <inputfile> -p <list_of_pools> [OPTIONS]
```

The **<inputfile>** 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 follow either of the following format.

- ▶ Each line ends with -- <filename> with a space before and after the double dash (the file list record format of the mmapplypolicy command).
- ▶ Each line contains a file name with an absolute path or a relative path that is based on the working directory. This format is unavailable when you run the command with the --mmbackup option.

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 commas.

Note: This command is not applicable for non-empty regular files.

6.13.2 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/share/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 IBM Spectrum Scale policies. Example 6-76 shows the sample policy.

Example 6-76 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@sample_library'

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
    )
  )
)
)
)
)
)

```

6.14 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 task 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 task queue. A sixtyish 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 task 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.

6.14.1 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 reading a file.

6.14.2 Selective recall using the `eeadm recall` command

The `eeadm 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 `eeadm 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 tasks 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 like 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 task 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 eeadm recall command

The `eeadm recall` command is used to recall non-empty regular files from tape. The `eeadm recall` command uses the following syntax:

► `eeadm recall <inputfile> [OPTIONS]`

The `<inputfile>` may contain one of two formats. Each line ends with “`-- <filename>`” with a space before and after the double dash. Or each line contains a file name with an absolute path or a relative path that is based on the working directory. The `eeadm recall` command with the output of another command.

The `eeadm recall` command can take as input the output of other commands through a pipe. In Example 6-77, all files with names ending with `.bin` are recalled under the `/ibm/gpfs/prod` directory, including subdirectories. Therefore, it is convenient to recall whole directories with a simple command.

Example 6-77 The eeadm recall command with the output of another command

```
[root@saitama2 prod]# find /ibm/gpfs/prod -name "*.bin" -print | eeadm recall -l lib_saitama
2019-01-09 15:55:02 GLESL277I: The "eeadm recall command" is called without specifying an input file waiting for
standard input.
    If necessary press ^D to exit.
2019-01-09 15:55:02 GLESL268I: 4 file name(s) have been provided to recall.
2019-01-09 15:55:03 GLESL700I: Task selective_recall was created successfully, task id is 7112.
2019-01-09 15:55:09 GLESL263I: Recall result: 4 succeeded, 0 failed, 0 duplicate, 0 not migrated, 0 not found, 0
unknown.
```

6.14.3 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.

Use the following **dsmmigfs** command options to set both the stub size and preview size of the file system being managed by IBM Spectrum Archive EE:

```
dsmmigfs Update -STUBsize  
dsmmigfs Update -PREViewsize
```

The value for the **STUBsize** is a multiple of the IBM Spectrum Scale file system's block size. this value can be obtained by running the **mm1sfs <filesystem>**. The **PREViewsize** parameter must be equal to or less than the **STUBsize** value. Both parameters take a positive integer in bytes.

Example 6-78 shows how to set both the **STUBsize** and **PREViewsize** on the IBM Spectrum Scale file system.

Example 6-78 Updating STUBsize and PREViewsize

```
[root@kyoto prod]# dsmmigfs Update -STUBsize=3145728 -PREViewsize=1048576  
/ibm/gpfs  
IBM Spectrum Protect  
Command Line Space Management Client Interface  
  Client Version 8, Release 1, Level 6.0  
  Client date/time: 01/11/2019 12:42:20  
(c) Copyright by IBM Corporation and other(s) 1990, 2018. All Rights Reserved.  
  
ANS9068I dsmmigfs: dsmmigfstab file updated for file system /ibm/gpfs.
```

For more information about the **dsmmigfs update** command, see [IBM Documentation](#).

6.14.4 Recommend Access Order

The Recommended Access Order (RAO) is a feature provided by drives¹ that serves a method to enable efficient recalls of files from a single cartridge. Drives that support RAO will give the most efficient access order when reading multiple records in a tape cartridge.

In a recall, IBM Spectrum Archive EE will queue the tasks to send requests in an order that is based on tape and starting blocks (refer to 6.14, "Recall" on page 190). This reading optimization is efficient when significantly large numbers of files that are located in linear positions of the tape are requested. However, in a case where the requested files are distributed across multiple bands and wraps, this optimization method was not as effective since the data could not be read efficiently with a simple linear read.

RAO was introduced to address such cases. Though additional calculation time will be needed to create the recommended reading order of files, effective use of RAO will decrease read time significantly in an ideal case.

¹ Refer to 3.3, "System requirements" on page 50 for RAO supported hardware.

Figure 6-1 is a conceptual diagram displaying the seeking path of five files based on block order. The dotted arrows represents the servo head seek path, in this case, starting from File 1 and ending with File 5.

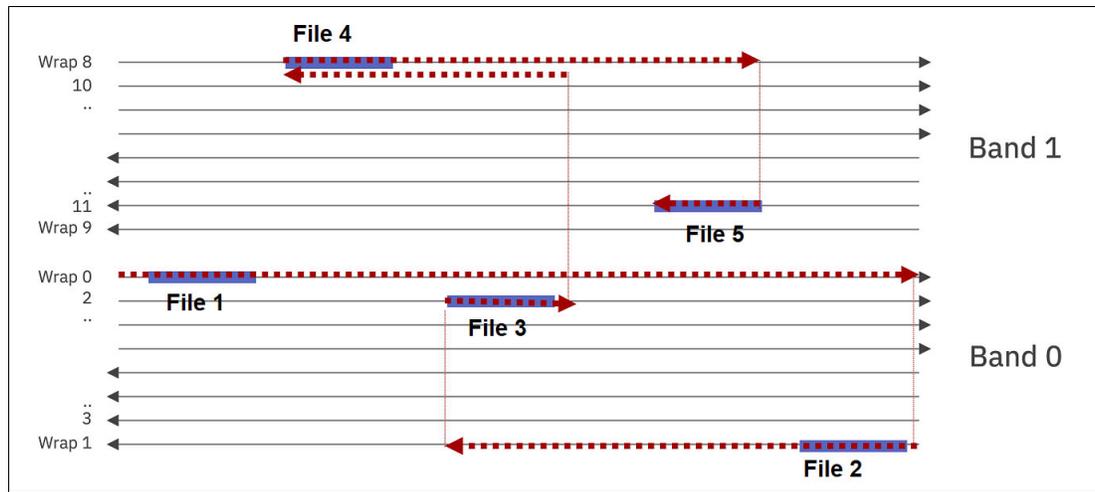


Figure 6-1 A conceptual diagram displaying the seek path of files based on block order

In contrast, Figure 6-2 shows the servo head movement when same files are read based on RAO. When compared with the previous figure, the total seek path length represented with the dotted arrow will be significantly shorter in RAO, therefore indicating that the total recall time will be reduced.

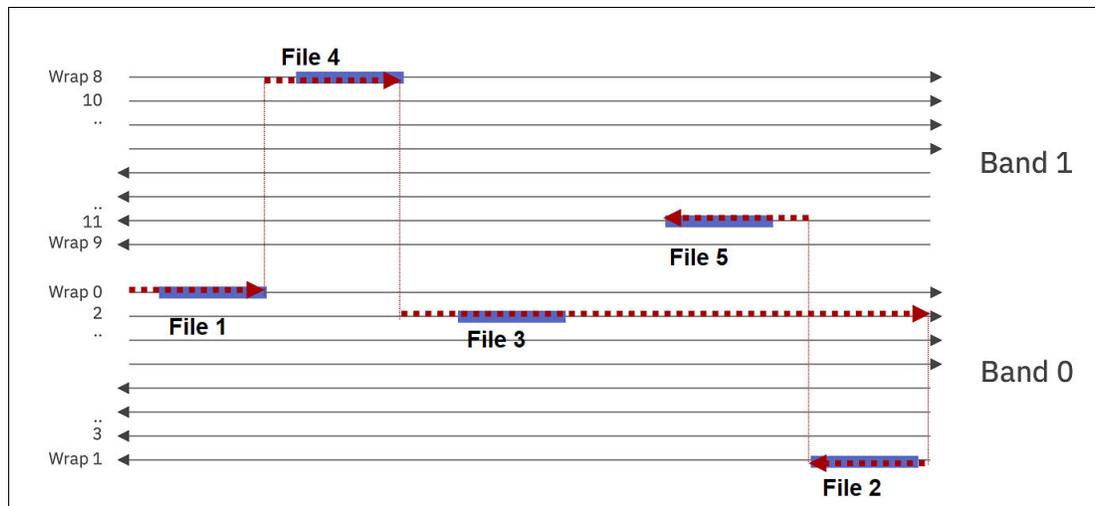


Figure 6-2 A conceptual diagram displaying the seek path of files based on RAO

Though RAO is a powerful feature, there are cases where it will not be effective. As addressed above, one example of such cases is when there are massive numbers of files to read from a cartridge. In such cases, a block sorted read of a cartridge is the more reasonable choice.

To address such cases, IBM Spectrum Archive EE provides an automatic transition of the two methods during recall requests.

This automatic transition includes the following features:

- ▶ File queuing on recall requests to same cartridges
When possible, multiple recall requests to a single cartridge will be joined into a single RAO request queue. In this case, the RAO requests will be joined up to a 2,700 file queue in prospect to reduce the total reading time.
- ▶ RAO usage decision on cases where large number of files are recalled
IBM Spectrum Archive EE will automatically switch to a block order recalling method when recall requests sum to 5,400 file mark for a single cartridge.²
- ▶ RAO usage decision on supported cartridges and drives
RAO function will not be used if there are no compatible drives and media available in the library. Refer to the [IBM Documentation - Supported tape drives and media](#) for compatible drives and media.

Note: For more information on RAO interface specification, refer to the [INCITS 503-202x: Information technology - SCSI Stream Commands - 5 \(SSC-5\) Overview](#) standard document.

Manual configuration of RAO

IBM Spectrum Archive EE provides a RAO global and local enabling/disabling switch for manually controlling the RAO feature.

Leaving the RAO usage on (or in “auto”) is generally recommended in most cases, though turning the RAO feature off may come to consideration if:

- ▶ Recalls are always made with large numbers of files (more than 5400 files) from a single cartridge, at a time.
- ▶ IBM Spectrum Archive EE is unable to support the RAO feature since there are no supported drives or media in the library.

RAO function can be globally enabled or disabled with the `recall_use_rao` option in the `eeadm cluster set` command. The default setting is “auto”. Example shows how to disable and enable the RAO feature. Globally configuring the RAO feature

```
[root@mikasa1 ~]# eeadm cluster set -a recall_use_rao -v disable
2021-10-11 16:29:58 GLESL802I: Updated attribute recall_use_rao.
```

```
[root@mikasa1 ~]# eeadm cluster set -a recall_use_rao -v auto
2021-10-11 16:30:58 GLESL802I: Updated attribute recall_use_rao.
```

Current settings can be found with the `cluster show` command (see Example 6-79).

Example 6-79 Showing current settings of the RAO feature

```
[root@mikasa1 ~]# eeadm cluster show
Attribute          Value
..
recall_use_rao    auto
```

² The 5,400 file mark is based on laboratory experimental data. In the experiment, we compared the effect of recalling N number of 20MB files with the RAO method versus the starting block order method. The total net recall time indicated that the starting block order method becomes more efficient just over $N=5,000$ files, and the RAO performance are significantly less effective by $N=10,000$ files.

The global RAO usage switch can be overridden locally in a recall command with the `--use-rao` option (see Example 6-80).

Example 6-80 Overriding global RAO switch with local recall command

```
[root@mikasa1 recallsample]# eeadm recall migreclist.txt --use-rao auto
..
```

6.15 Recalling files to their resident state

This section describes the `eeadm recall` command using the `--resident` option. However, this command should rarely be used. The `eeadm recall --resident` 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 `eeadm recall --resident` command will recall the files back initially if the files are migrated then mark the files resident and remove the link between disk and tape. If the files are already premigrated the recall operation will be skipped and will just mark the files resident. This option removes metadata on IBM Spectrum Scale, which is used for keeping the file/object state.

A typical usage of the `eeadm recall --resident` command is a user accidentally migrates/premigrates one or more files to the wrong tape or has forgotten to make a copy.

Example 6-81 shows the output of making a file resident.

Example 6-81 Making a migrated file resident again

```
[root@mikasa1 recallsample]# eeadm file state migrecfile2.txt
Name: /ibm/gpfs/recallsample/migrecfile2.txt
State: migrated
ID: 11151648183451819981-8790189057502350089-1392721861-159785-0
Replicas: 1
Tape 1: MB0241JE@Arnoldpool@lib1 (tape state=appendable)

[root@mikasa1 recallsample]# eeadm recall migreclist.txt --resident
2021-03-21 07:17:26 GLESL268I: 1 file name(s) have been provided to recall.
2021-03-21 07:17:26 GLESL700I: Task selective_recall was created successfully,
task ID is 1089.
2021-03-21 07:17:29 GLESL839I: All 1 file(s) has been successfully processed.
2021-03-21 07:17:29 GLESL845I: Succeeded: 1 resident, 0 already_resident

[root@mikasa1 recallsample]# eeadm file state migrecfile2.txt
Name: /ibm/gpfs/recallsample/migrecfile2.txt
State: resident
```

Note: migreclist.txt contains the list of files to be recalled. In this case, it contains only 'migrecfile2.txt'.

6.16 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. Note however that the 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 capacity 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 tasks 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 synchronization and update to the IBM Spectrum Archive EE tapes can be time consuming and normally only required before any **eadm tape export** command. Normally if files are deleted from the IBM Spectrum Scale file systems, merely the update for the amount of reclaimable space is sufficient information. In other words, as files are deleted, administrators only want to know how much space can be reclaimed. Thus starting with v1.3.0.0, the default for **eadm tape reconcile** is to only update internal metadata information in order to correctly reflect the **Reclaimable%** column of the **eadm tape list** command. In order to update the tape (like previous to v1.3.0.0), the **--commit-to-tape** option should be supplied as part of the **eadm tape reconcile** command.

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 **eadm tape reconcile** command fails and the following failure message appears:

GLESL098E: The same type of task or conflicting task was previously requested and it is running. Wait for completion of the task and try again.

2. After a reconciliation process is started, new migration tasks are prevented until the reconciliation process completes on the reconciling tapes. However, if any migration tasks are running, the reconciliation process does not begin until all migration tasks 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 the **Reclaimable%** column after that file is deleted from the IBM Spectrum Scale file system. Example 6-82 shows the initial file state of a single file on tape.

Example 6-82 Display the file state of a migrated file

```
[root@ueno ~]# eadm file state
/gpfs/gpfs0/cmt/md1/files/LTFS_EE_FILE_mQH4oCZXQZKSXc6cYZjAuXPnAFq56QBWJmY66PcuVnA
IW6BkrK_gzCwGJ.bin
Name:
/gpfs/gpfs0/cmt/md1/files/LTFS_EE_FILE_mQH4oCZXQZKSXc6cYZjAuXPnAFq56QBWJmY66PcuVnA
IW6BkrK_gzCwGJ.bin
State: premigrated
ID: 9226311045824247518-17452468188422870573-2026950408-6649759-0
Replicas: 3
Tape 1: P1B064JE@ueno@perf_lib (tape state=appendable)
Tape 2: P1B076JE@kanda@perf_lib (tape state=appendable)
Tape 3: P1B073JE@shimbashi@perf_lib (tape state=appendable)
```

The file is also present on the IBM Spectrum Scale file system, as shown in Example 6-83.

Example 6-83 List the file on the IBM Spectrum Scale file system

```
[root@ueno ~]# ls -hl
/gpfs/gpfs0/cmt/md1/files/LTFS_EE_FILE_mQH4oCZXQZKSXc6cYZjAuXPnAFq56QBWJmY66PcuVnA
IW6BkrK_gzCwGJ.bin
-rw----- 1 root root 1.8G Nov 12 15:29
/gpfs/gpfs0/cmt/md1/files/LTFS_EE_FILE_mQH4oCZXQZKSXc6cYZjAuXPnAFq56QBWJmY66PcuVnA
IW6BkrK_gzCwGJ.bin
```

At this point, this file can be removed with the `rm` command. After the file is removed, IBM Spectrum Archive EE has the reclaimable space information from `eeadm tape list` for these tapes (see Example 6-84).

Example 6-84 Tape list in IBM Spectrum Archive EE

```
[root@ueno ~]# eeadm tape list | egrep "Tape|P1B064JE|P1B076JE|P1B073JE"
```

Tape ID	Status	State	Usable(GiB)	Used(GiB)	Available(GiB)
P1B073JE	ok	appendable	18147	18137	9
0%	shimbashi	perf_lib homeslot	-		
P1B064JE	ok	appendable	18147	18134	12
0%	ueno	perf_lib drive	-		
P1B076JE	ok	appendable	18147	18140	6
0%	kanda	perf_lib homeslot	-		

If you perform a reconciliation of the tape now using the default settings, IBM Spectrum Archive EE only updates the amount of reclaimable space on the tape within the internal metadata structure, as shown in Example 6-85.

Example 6-85 Reconcile the tape

```
[root@ueno ~]# eeadm tape reconcile P1B064JE -p ueno
2018-12-07 12:31:21 GLESL700I: Task reconcile was created successfully, task id is 5585.
2018-12-07 12:31:21 GLESS016I: Reconciliation requested.
2018-12-07 12:31:22 GLESS050I: GPFS file systems involved: /gpfs/gpfs0 .
2018-12-07 12:31:22 GLESS210I: Valid tapes in the pool: MB0247JE P1B064JE P1B067JE P1B063JE P1B074JE .
2018-12-07 12:31:22 GLESS049I: Tapes to reconcile: P1B064JE .
2018-12-07 12:31:22 GLESS134I: Reserving tapes for reconciliation.
2018-12-07 12:31:22 GLESS135I: Reserved tapes: P1B064JE .
2018-12-07 12:31:22 GLESS054I: Creating GPFS snapshots:
2018-12-07 12:31:22 GLESS055I: Deleting the previous reconcile snapshot and creating a new one for /gpfs/gpfs0 ( gpfs0 ).
2018-12-07 12:31:25 GLESS056I: Searching GPFS snapshots:
2018-12-07 12:31:25 GLESS057I: Searching GPFS snapshot of /gpfs/gpfs0 ( gpfs0 ).
2018-12-07 12:31:44 GLESS060I: Processing the file lists:
2018-12-07 12:31:44 GLESS061I: Processing the file list for /gpfs/gpfs0 ( gpfs0 ).
2018-12-07 12:33:52 GLESS141I: Removing stale DMAPI attributes:
2018-12-07 12:33:52 GLESS142I: Removing stale DMAPI attributes for /gpfs/gpfs0 ( gpfs0 ).
2018-12-07 12:33:52 GLESS063I: Reconciling the tapes:
2018-12-07 12:33:52 GLESS248I: Reconcile tape P1B064JE.
2018-12-07 12:37:19 GLESS002I: Reconciling tape P1B064JE complete.
2018-12-07 12:37:19 GLESS249I: Releasing reservation of tape P1B064JE.
2018-12-07 12:37:19 GLESS058I: Removing GPFS snapshots:
2018-12-07 12:37:19 GLESS059I: Removing GPFS snapshot of /gpfs/gpfs0 ( gpfs0 ).
```

If you now list the tapes with `eeadm tape list`, you may see that the reclaimable space percentage be increased due to the deleted files (because it is in percentage and not in GBs, the size of deletion need to be large to see a percentage change).

6.17 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 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.

6.17.1 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 6.16, “Reconciliation” on page 197. If this is not performed, the reclamation might fail with a message that recommends to perform a reconcile.

- ▶ 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 task is stopped for the recall. After the recall is complete, the reclamation resumes automatically.

- ▶ Tape drives used for reclamation

To reclaim a tape cartridge, two tape drives with the “g” role attribute on the same node will be used. When multiple tapes cartridges are specified for reclamation, the reclaim function will try to use as many tape drives as it needs to process the multiple reclamation in parallel. There is an option to the command to limit the number of concurrent reclamation process thus limiting the number of tape drives that will be used at once.

Use the **eeadm tape reclaim** command to start reclamation of a specified tape cartridge pool or of certain tape cartridges within a specified tape cartridge pool. The **eeadm tape 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 6-86 shows the results of reclaiming a single tape cartridge 058AGWL5.

Example 6-86 Reclamation of a single tape cartridge

```
[root@mikasa1 ~]# eeadm tape reclaim JCB370JC -p pool1 -l lib2
2021-03-24 04:08:49 GLESL700I: Task reclaim was created successfully, task ID is
1149.
```

```

2021-03-24 04:08:49 GLESR216I: Multiple processes started in parallel. The maximum
number of processes is unlimited.
2021-03-24 04:08:49 GLESL084I: Start reclaiming 1 tapes in the following list of
tapes:
2021-03-24 04:08:49 GLESR212I: Source candidates: JCB370JC .
2021-03-24 04:08:49 GLESL677I: Files on tape JCB370JC will be copied to tape
JCA561JC.
2021-03-24 04:14:30 GLESL085I: Tape JCB370JC successfully reclaimed, it remains in
tape pool Arno1dpool12.
2021-03-24 04:14:30 GLESL080I: Reclamation complete. 1 tapes reclaimed, 0 tapes
unassigned from the tape pool.

```

At the end of the process, the tape cartridge is reformatted. The tape remains in the tape cartridge pool unless the “--unassign” option is specified. For more information, see , “The `eeadm <resource type> --help` command” on page 316.

6.18 Checking and repairing tapes

There are three tape states that can occur on a tape which helps inform the user to perform a check. Once the root cause of the status is determined use the `eeadm tape validate` command to perform a check on the tapes and update the state back to `appendable`. All three states will prevent further migration and recall tasks to those tapes effected, therefore it is important to resolve the issues found and clear the states. The following are the check tape states:

► **check_tape_library**

tapes fall into this state when they fail to get mounted to a drive, or is stuck in a drive and failed to get unmounted.

► **require_validate**

tapes fall into this state when the system detects some metadata mismatches on the tape.

► **check_key_server**

tapes fall into this state when there is a single tape that failed a write or read requiring an encryption key from an encryption key server.

Example 6-87 shows a tape in the `check_key_server` state and using the `eeadm tape validate` to restore its state back to `appendable` or `full` after fixing the encryption key server.

Example 6-87 Restoring a `check_key_server` tape back to `appendable`

```

[root@saitama2 ~]# eeadm tape list -l lib_mikasa
Tape ID   Status State           Usable(GiB) Used(GiB) Available(GiB)
Reclaimable% Pool Library      Location Task ID
MB0021JM error  check_key_server 0          0          0
0% pool2 lib_mikasa homeslot -

[root@saitama2 ~]# eeadm tape validate MB0021JM -p pool2 -l lib_mikasa
2019-01-10 11:27:24 GLESL700I: Task tape_validate was created successfully, task
id is 7125.
2019-01-10 11:28:41 GLESL388I: Tape MB0021JM is successfully validated.

[root@saitama2 ~]# eeadm tape list -l lib_mikasa

```

Tape ID	Status	State	Usable(GiB)	Used(GiB)	Available(GiB)
Reclaimable%	Pool	Library	Location	Task ID	
MB0021JM	ok	appendable	4536	4536	0
lib_mikasa	homeslot	-			0% pool2

Two other tape states exist that require the user to replace the tapes. These two states occur when a tape encounters a read or write failure, and can be remedied by using the **eadm tape replace** command. This command moves the data off the bad tape and onto an appendable tape within the same pool maintaining the migration order and finally removing the bad tape from the pool. The following are the two replacement tape states:

► **need_replace**

The IBM Spectrum Archive EE system detected one or more permanent read errors on this tape

► **require_replace**

The IBM Spectrum Archive EE system detected a permanent write error on the tape.

Example 6-88 shows the output of running the **eadm tape replace** command

Example 6-88 eadm tape replace

```
[root@mikasa1 prod]# eadm tape replace JCB206JC -p test -l lib_mikasa
2018-10-09 10:51:58 GLESL700I: Task tape_replace was queued successfully, task id
is 27382.
2018-10-09 10:51:58 GLESL755I: Kick reconcile before replace against 1 tapes.
2018-10-09 10:52:57 GLESS002I: Reconciling tape JCB206JC complete.
2018-10-09 10:52:58 GLESL756I: Reconcile before replace was finished.
2018-10-09 10:52:58 GLESL753I: Starting tape replace for JCB206JC.
2018-10-09 10:52:58 GLESL754I: Found a target tape for tape replace (MB0241JE).
2018-10-09 10:55:05 GLESL749I: Tape replace for JCB206JC is successfully done.
```

For more information about the various tape cartridge states, see 10.1.4, “Tape status codes” on page 320.

6.19 Importing and exporting

The import and export processes are the mechanisms for moving data on the LTFS written tape cartridges into or out of the IBM Spectrum Archive EE environment.

6.19.1 Importing

Import tape cartridges to your IBM Spectrum Archive EE system by running the **eadm tape import <list_of_tapes> -p <pool> [-l <library>] [OPTIONS]** command.

When you import a tape cartridge, the **eadm tape import** 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
3. Adds the file stubs in an import directory within the IBM Spectrum Scale file system

Example 6-89 shows the import of an LTFS tape cartridge that was created on a different LTFS system into a directory that is called FC0257L8 in the /ibm/gpfs file system.

Example 6-89 Import an LTFS tape cartridge

```
[root@ginza ~] eeadm tape import FC0257L8 -p pool1 -P /ibm/gpfs/  
2019-02-26 09:52:15 GLESL700I: Task import was created successfully, task id is  
2432  
2019-02-26 09:54:33 GLESL064I: Import of tape FC0257L8 complete.
```

Importing file paths

The default import file path for the **eeadm tape import** command is **/{GPFS file system}/IMPORT**. As shown in Example 6-90, 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 6-90 Import by using default parameters

```
[root@saitama2 ~]# eeadm tape import JCB610JC -p test3 -l lib_saitama  
2019-01-10 13:33:12 GLESL700I: Task import was created successfully, task id is  
7132.  
2019-01-10 13:34:51 GLESL064I: Import of tape JCB610JC complete.
```

```
[root@saitama2 ~]# ls -las /ibm/gpfs/IMPORT/  
JCB610JC/
```

Example 6-91 on page 203 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 6-91 Import by using the -P parameter

```
[root@saitama2 gpfs]# eeadm tape import JCB610JC -p test3 -l lib_saitama -P  
/ibm/gpfs/alternate  
2019-01-10 14:20:54 GLESL700I: Task import was created successfully, task id is  
7139.  
2019-01-10 14:22:30 GLESL064I: Import of tape JCB610JC complete.  
[root@saitama2 gpfs]# ls /ibm/gpfs/alternate/  
JCB610JC
```

With each of these parameters, you have the option of renaming imported files by using the **--rename** parameters. This option will only rename any imported files with conflicting names by appending the suffix “**_i**” where **i** is a number from 1 to **n**.

Importing offline tape cartridges

For more information about offline tape cartridges, see 6.19.2, “Exporting tape cartridges” on page 204. Offline tape cartridges can be reimported to the IBM Spectrum Scale namespace by running the **eeadm tape online** command.

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.

Example 6-92 shows an example of making an offline tape cartridge online.

Example 6-92 Online an offline tape cartridge

```
[root@saitama2 ~]# eeadm tape online JCB610JC -p test3 -l lib_saitama
2019-01-10 14:42:36 GLESL700I: Task import was created successfully, task id is
7143.
2019-01-10 14:44:11 GLESL064I: Import of tape JCB610JC complete.
```

6.19.2 Exporting tape cartridges

Export tape cartridges from your IBM Spectrum Archive EE system by running the **eeadm tape export** command. When you export a tape cartridge, the process 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 the tape cartridges are unchanged by the export, and are accessible by other LTFS systems.

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 6-93 shows the typical output from the **export** command.

Example 6-93 Export a tape cartridge

```
[root@mikasa1 recallsample]# eeadm tape export JCB370JC -p test3-1 lib2 --remove
2021-03-24 03:30:16 GLESL700I: Task export_remove was created successfully, task
ID is 1139.
2021-03-24 03:30:16 GLESS134I: Reserving tapes.
2021-03-24 03:30:16 GLESS269I: JCB370JC is mounted. Moving to homeslot.
2021-03-24 03:31:13 GLESS135I: Reserved tapes: JCB370JC .
2021-03-24 03:31:13 GLESL719I: Reconcile as a part of an export is starting.
2021-03-24 03:32:48 GLESS002I: Reconciling tape JCB370JC complete.
2021-03-24 03:32:49 GLESL632I: Reconcile as a part of an export finishes
2021-03-24 03:32:49 GLESL073I: Tape export for JCB370JC was requested.
2021-03-24 03:34:19 GLESM399I: Removing tape JCB370JC from pool test3 (Force).
2021-03-24 03:34:19 GLESL762I: Tape JCB370JC was forcefully unassigned from pool
test3 during an export operation. Files on the tape cannot be recalled. Run the
"eeadm tape import" command for the tape to recall the files again.
2021-03-24 03:34:19 GLESL074I: Export of tape JCB370JC complete.
2021-03-24 03:34:19 GLESL490I: The export command completed successfully for all
tapes.
```

Example 6-94 on page 205 shows how the normal exported tape is displayed as exported by running the **eeadm info tapes** command.

Example 6-94 Display status of normal export tape cartridge

```
# eeadm tape list
```

Tape ID	Status	State	Usable(GiB)	Used(GiB)	Available(GiB)	Reclaimable%	Pool	Library	Location	Task ID
FC0260L8	ok	appendable	10907	7	10899	0%	temp	liba	homeslot	-
UEF108M8	ok	appendable	5338	46	5292	0%	pool4	liba	homeslot	-
DV1993L7	ok	appendable	5338	37	5301	0%	pool4	liba	homeslot	-
FC0255L8	ok	exported	0	0	0	0%	-	liba	homeslot	-

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 9.5, “Software” on page 303 on how to clean up the remaining files on the IBM Spectrum Scale file system.

Regarding full replica support, Export/Import does not depend on the primary/redundant copy. When *all* copies are exported, the file is exported.

Table 6-1 lists 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 6-1 Export operations use case scenario of file with three tapes

Operation	File
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.

6.19.3 Offlining tape cartridges

Offline tape cartridges from your IBM Spectrum Archive EE system by running the **eeadm tape offline** command. This process marks all files from the tape cartridges or tape cartridge pool that are specified and marks them offline 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 making online the tape cartridge by using the **eeadm tape online** command.

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 6-95, tape JCB610JC contains redundant copies of five MPEG files that must be moved off-site.

Example 6-95 Offlining a tape cartridge

```
[root@saitama2 ~]# eeadm tape offline JCB610JC -p test3 -l lib_saitama -o "Moved to storage room B"
2019-01-10 14:39:55 GLESL700I: Task tape_offline was created successfully, task id is 7141.
2019-01-10 14:39:55 GLESL073I: Offline export of tape JCB610JC has been requested.
```

If you run the **eeadm tape list** command, you can see the offline status of the tape cartridge, as shown in Example 6-96.

Example 6-96 Display status of offline tape cartridges

```
[root@saitama2 ~]# eeadm tape list -l lib_saitama
Tape ID  Status  State      Usable(GiB)  Used(GiB)  Available(GiB)  Reclaimable%  Pool  Library  Location  Task ID
JCA561JC ok      offline    0            0           0              0%         pool2  lib_saitama  homeslot -
JCA224JC ok      appendable 6292         0           6292          0%         pool1  lib_saitama  homeslot -
JCC093JC ok      appendable 6292         496         5796          0%         pool1  lib_saitama  homeslot -
JCB745JC ok      append_fenced 6292         0           0              0%         pool2  lib_saitama  homeslot -
```

It is now possible to physically remove tape JCA561JC from the tape library so that it can be sent to the off-site storage location.

Regarding full replica support, Offlining/Onlining does not depend on the primary/redundant copy. When *all* copies are offlined, the file is offline.

Table 6-2 lists a use case example where a file was migrated to three physical tapes: TAPE1, TAPE2, TAPE3. The file behaves as shown by offline operations.

Table 6-2 Export operations use case scenario of file with three tapes

Operation	File
TAPE1 is offline.	File is available (can recall).
TAPE1/TAPE2 are offline.	File is available (can recall).
TAPE1/TAPE2/TAPE3 are offline.	File is offline (cannot recall).
TAPE1/TAPE2 are offline exported, and then TAPE3 is exported.	File is offline (cannot recall, and the file will be only recalled from TAPE1/TAPE2 once they are online).

Example 6-97 shows the file state when the tape host has been offlined.

Example 6-97 Offline file state

```
[root@saitama2 prod]# eadm file state *.bin
Name: /ibm/gpfs/prod/file1
State: migrated (check tape state)
ID: 11151648183451819981-3451383879228984073-1435527450-974349-0
Replicas: 1
Tape 1: JCB610JC@test3@lib_saitama (tape state=offline)
```

6.20 Drive Role settings for task assignment control

IBM Spectrum Archive EE allows users to configure their drives to allow or disallow specific tasks. Each of the attributes corresponds to the tape drive's capability to perform a specific type of task. Here are the attributes:

- ▶ Migration
- ▶ Recall
- ▶ Generic

Table 6-3 lists the available IBM Spectrum Archive EE drive attributes for the attached physical tape drives.

Table 6-3 IBM Spectrum Archive EE drive attributes

Attributes	Description
Migration	If the <code>Migration</code> attribute is set for a drive, that drive can process migration tasks. If not, IBM Spectrum Archive EE never runs migration tasks by using that drive. Save tasks 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 <code>Migration</code> .
Recall	If the <code>Recall</code> attribute is set for a drive, that drive can process recall tasks. If not, IBM Spectrum Archive EE never runs recall tasks 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 <code>Recall</code> .
Generic	If the <code>Generic</code> attribute is set for a drive, that drive can process generic tasks. If not, IBM Spectrum Archive EE never runs generic tasks by using that drive. IBM Spectrum Archive EE creates and runs miscellaneous generic tasks for administrative purposes, such as formatting tape, checking tape, reconciling tape, reclaiming a tape, and validating a tape. Some of those tasks 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 <code>Generic</code> . For reclaiming tape, at least two tape drives are required, so at least two drives need the <code>Generic</code> 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:

```
eadm drive assign <drive serial> -n <node_id> -r
```

The `-r` option requires a decimal numeric parameter. A logical **OR** applies to set the three attributes: **Migrate (4)**, **Recall (2)**, and **Generic (1)**. For example, a number of 6 for `-r` allows migration and recall task while copy and generic task are disallowed. All of the attributes are set by default.

To check the current active drive attributes, the `eadm drive list` command is useful. This command shows each tape drive's attributes, as shown in Example 6-98.

Example 6-98 Check current IBM Spectrum Archive EE drive attributes

```
[root@saitama2 prod]# eadm drive list -l lib_saitama
```

Drive S/N	State	Type	Role	Library	Node ID	Tape	Node Group	Task ID
0000078PG24E	mounted	TS1160	mrg	lib_saitama	6	JD0321JD	G0	-
0000078PG20E	not_mounted	TS1160	mrg	lib_saitama	2	-	G0	-
0000078D9DBA	not_mounted	TS1155	mrg	lib_saitama	2	-	G0	-
00000000A246	not_mounted	TS1155	mrg	lib_saitama	2	-	G0	-
0000078PG24A	not_mounted	TS1160	mrg	lib_saitama	6	-	G0	-
0000078D82F4	unassigned	-	---	lib_saitama	-	-	-	-

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.

The role of an assigned drive can be modified by using the `eadm drive set` command. For example, the following command changes the role to Migrate and Recall.

```
eadm drive set <drive serial> -a role -v mr
```

The configuration change takes effect after the on-going process has completed.

Drive attributes setting hint: 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.

6.21 Tape drive intermix support

This section describes the physical tape drive intermix support

This enhancement has these objectives:

- ▶ Use IBM LTO-9 tapes and drives in mixed configuration with older IBM LTO (LTO-8, M8, 7, 6, and 5) generations
- ▶ Use 3592 JC/JD/JE cartridges along with IBM TS1160, TS1155, TS1150, and TS1140 drives in mixed environments

Note: An intermix of LTO and TS11xx tape drive technology and media on a single library is not supported by IBM Spectrum Archive EE.

The following main use cases are expected to be used by this 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-9, 8, M8, 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. For more information, see this [IBM Documentation web page](#).

When 3592 JC, JD, and JE tapes are used in a tape library and IBM TS1160, TS1155, TS1150, and TS1140 drives are used correct tapes and drives are selected by IBM Spectrum Archive EE to read or write the required data.

With this 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. The TS1160 supports only recalls of JC media formatted using TS1140.

The **eeadm tape assign** command can be used for pool configuration when new physical tape cartridges are added to your IBM Spectrum Archive EE setup:

```
eeadm tape assign <list_of_tapes> -p <pool> [-l <library>] [OPTIONS]
```

WORM support for the IBM TS1160, 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. 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 TS1160, 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 [this website](#).

6.21.1 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.

6.21.2 Function overview for WORM tape support

The following features are present to support 3592 WORM tapes:

- ▶ 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 6-99 shows how to set the WORM attribute to an IBM Spectrum Archive EE pool by using the **eadm pool create** command.

Example 6-99 Set the WORM attribute to an IBM Spectrum Archive EE pool

```
[root@saitama2 prod]# eadm pool create 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 that is shown in Example 6-100.

Example 6-100 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
```

6.21.3 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.

► **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 V5.1.2 immutable and appendOnly limitations, see [IBM Documentation](#).

Example 6-101 shows the output that you receive while working, showing, and changing the IBM Spectrum Scale immutable or appendOnly file attributes.

Example 6-101 Set or change an IBM Spectrum Scale file immutable file attribute

```
[root@ltfsee_node0]# echo "Jan" > jan_jonas.out
[root@ltfsee_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:1tfssn2

[root@ltfsee_node0]# mmchattr -i yes jan_jonas.out

[root@ltfsee_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:HMACSHA512'
  type: wrapped FEK WrpPar 'AES:KWRAP' CmbPar 'XORHMACSHA512'
  KEY-a65222ad-e8bc-4fcc-89a3-4a7c54d4181c:1tfssn2

[root@ltfsee_node0]# echo "Jonas" >> jan_jonas.out
```

```
-bash: jan_jonas.out: Read-only file system
[root@l1tfsee_node0]#
```

These `immutable` or `appendOnly` file attributes can be changed at any time by the IBM Spectrum Scale administrator, so IBM Spectrum Scale is not able to provide an complete 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 TS1160, TS1155, TS1150, and TS1140 tape drives (3592 JV, JY, JZ).
- ▶ If IBM Spectrum Scale `immutable` attributes are changed to *yes* *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 10.1, “Command-line reference” on page 316.

6.22 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 `eadm file state` 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, how many replicas exist, and the health state of the tape.

Example 6-102 shows the typical output of the `eadm file state` command. Some files are on multiple tape cartridges, some are in a migrated state, and others are premigrated only.

Example 6-102 Files location

```
[root@saitama2 prod]# eadm file state *.bin
Name: /ibm/gpfs/prod/file1.bin
State: premigrated
ID: 11151648183451819981-3451383879228984073-1435527450-974349-0
```

```

Replicas: 2
Tape 1: JCB610JC@test3@lib_saitama (tape state=appendable)
Tape 2: JD0321JD@test4@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/file2.bin
State: migrated
ID: 11151648183451819981-3451383879228984073-2015134857-974348-0
Replicas: 2
Tape 1: JCB610JC@test3@lib_saitama (tape state=appendable)
Tape 2: JD0321JD@test4@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/file3.bin
State: migrated
ID: 11151648183451819981-3451383879228984073-599546382-974350-0
Replicas: 2
Tape 1: JCB610JC@test3@lib_saitama (tape state=appendable)
Tape 2: JD0321JD@test4@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/file4.bin
State: migrated
ID: 11151648183451819981-3451383879228984073-2104982795-3068894-0
Replicas: 1
Tape 1: JD0321JD@test4@lib_saitama (tape state=appendable)

```

For more information about supported characters for file names and directory path names, see [IBM Documentation](#).

6.23 Obtaining system resources, and tasks information

This section describes how to obtain resource inventory information and information about ongoing migration and recall tasks with IBM Spectrum Archive EE. You can use the **eeadm task list** command to obtain information about current tasks and the **eeadm task show** to see detailed information about the task. To view IBM Spectrum Archive EE system resources use any of the following commands:

- ▶ eeadm tape list
- ▶ eeadm drive list
- ▶ eeadm pool list
- ▶ eeadm node list
- ▶ eeadm nodegroup list

Example 6-103 shows the command that is used to display all IBM Spectrum Archive EE tape cartridge pools.

Example 6-103 Tape cartridge pools

```

[root@saitama2 prod]# eeadm pool list
Pool Name Usable(TB) Used(TB) Available(TB) Reclaimable% Tapes Type Library
Node Group
myWORM          0.0      0.0          0.0          0%          0  -
lib_saitama G0
pool1          64.1      3.9          60.2          0%          10 3592
lib_saitama G0

```

```
test2          6.1          0.0          6.1          0%          1 3592
lib_saitama  G0
```

Example 6-104 shows the serial numbers and status of the tape drives that are used by IBM Spectrum Archive EE.

Example 6-104 Drives

```
[root@saitama2 prod]# eeadm drive list
Drive S/N      State      Type      Role      Library      Node ID      Tape      Node Group      Task
ID
0000078PG24E  mounted    TS1160    mrg      lib_saitama  6           JD0321JD  G0              -
0000078PG20E  not_mounted TS1160    mrg      lib_saitama  2           -         G0              -
0000078D9DBA  not_mounted TS1155    mrg      lib_saitama  2           -         G0              -
```

To view all the IBM Spectrum Archive EE tape cartridges, run the command that is shown in Example 6-105.

Example 6-105 Tape cartridges

```
[root@saitama2 prod]# eeadm tape list
Tape ID      Status      State      Usable(GB)  Used(GB)  Available(GB)  Reclaimable%  Pool  Library      Location      Task ID
JCA561JC    ok          offline    0            0          0              0%        pool2  lib_saitama  homeslot     -
JCA224JC    ok          appendable 6292         0          6292           0%        pool1  lib_saitama  homeslot     -
JCC093JC    ok          appendable 6292         496       5796           0%        pool1  lib_saitama  homeslot     -
```

Regularly monitor the output of IBM Spectrum Archive EE tasks to ensure that tasks are progressing as expected by using the **eeadm task list** and **eeadm task show** command. Example 6-106 shows a list of active tasks.

Example 6-106 Active tasks

```
[root@saitama2 prod]# eeadm task list
TaskID      Type      Priority      Status      #DRV      CreatedTime(-0700)  StartedTime(-0700)
7168      selective_recall  H          running    0          2019-01-10_16:27:30  2019-01-10_16:27:30
7169      selective_recall  H          waiting    0          2019-01-10_16:27:30  2019-01-10_16:27:30
```

The **eeadm node list** command (see Example 6-107) provides a summary of the state of each IBM Spectrum Archive EE component node.

Example 6-107 eeadm node list

```
[root@saitama2 prod]# eeadm node list
Node ID      State      Node IP      Drives      Ctrl Node      Library      Node Group      Host Name
4           available  9.11.244.44  2           yes(active)    lib_saitama  G0              saitama2
2           available  9.11.244.43  3           yes            lib_saitama  G0              saitama1
3           available  9.11.244.42  1           yes(active)    lib_mikasa   G0              mikasa2
1           available  9.11.244.24  2           yes            lib_mikasa   G0              mikasa1
```

The **eeadm task show** command (see Example 6-108) provides a summary of the specified active task in IBM Spectrum Archive EE. The task id corresponds to the task id that is reported by the **eeadm task list** command.

Example 6-108 Detailed outlook of recall task

```
[root@saitama2 prod]# eeadm task show 7168 -v
```

```

=== Task Information ===
Task ID:          7168
Task Type:       selective_recall
Command Parameters: eeadm recall mig -l lib_saitama
Status:         running
Result:         -
Accepted Time:   Thu Jan 10 16:27:30 2019 (-0700)
Started Time:    Thu Jan 10 16:27:30 2019 (-0700)
Completed Time:  -
In-use Libraries: lib_saitama
In-use Node Groups: G0
In-use Pools:    test3
In-use Tape Drives: 0000078PG20E
In-use Tapes:    JCB610JC
Workload:       3 files, 5356750 bytes in total to recall in this task.
Progress:       1 completed (or failed) files / 3 total files.
Result Summary:  -
Messages:
  2019-01-10 16:27:30.421251 GLESM332W: File /ibm/gpfs/prod/LTFS_EE_FILE_2dEPRHh_M.bin is not migrated.

```

6.24 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 **eeadm cluster start** and **eeadm cluster 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 6-4 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 6-4 SNMP error message information

Description	Name	Severity	OID
GLESV100I: 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
GLESV102W: Part of IBM Spectrum Archive EE nodes failed to start.	ibmsaDegradeV102PartialStart	Warning	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

Description	Name	Severity	OID
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.	ibmsaInfoV301GPFSEOperational	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
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

Description	Name	Severity	OID
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
GLESM609W: Pool space is going to be small.	ibmsaWarnM609PoolLowSpace	Warning	1.3.6.1.4.1.2.6.246.1.2.31.0.609
GLESM613E: There is not enough space available on the tapes in a pool for migration.	ibmsaErrM613NoSpaceForMig	Error	1.3.6.1.4.1.2.6.246.1.2.31.0.613

6.25 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
```

## 6.25.1 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@l1tfs97 ~]# 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
```

## 6.25.2 Example of an SNMP trap

Example 6-109 shows the type of trap information that is received by the SNMP server.

*Example 6-109 SNMP trap example of an IBM Spectrum Archive EE node that has a low pool threshold*

---

```
2018-11-26 09:08:42 tora.tuc.stglabs.ibm.com [UDP: [9.11.244.63]:60811->[9.11.244.63]:162]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (147206568) 17 days, 0:54:25.68
SNMPv2-MIB::snmpTrapOID.0 = OID: IBMSA-MIB::ibmsaWarnM609PoolLowSpace
IBMSA-MIB::ibmsaMessageSeverity.0 = INTEGER: warning(40) IBMSA-MIB::ibmsaJob.0 = INTEGER:
other(7) IBMSA-MIB::ibmsaEventNode.0 = STRING: "tora.tuc.stglabs.ibm.com"
IBMSA-MIB::ibmsaMessageText.0 = STRING: "GLESM609W: Pool space is going to be small, library:
lib_tora, pool: je_pool1, available capacity: 23.6(TiB), threshold:
30(TiB)"
```

---

## 6.26 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.

## 6.26.1 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`.
- ▶ **mode: <string>**  
The current operational mode of the pool. Access to the member tapes is temporarily disabled when this field is set to “disabled”, “relocation\_source”, or “relocation\_destination”. Under normal operating conditions, the field is set to “normal”. If an internal error occurs, the field is set to an empty string.
- ▶ **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.
- ▶ **non\_appendable\_space: <number>**  
The total capacity, in bytes, of tapes in the pool that cannot be written to in the format specified for the pool, and that don't match the `media_restriction` value for the pool. The format and the `media_restriction` values are provided as attributes of the pool.
- ▶ **num\_of\_tapes: <number>**  
Number of tapes assigned to the pool.
- ▶ **format\_class: <string>**  
The format class of the pool.

- ▶ **media\_restriction: <string>**  
The **media\_restriction** setting is stored in the regular expression of the bar code value. The **media\_restriction** is used to define the type of cartridge that can be used for writing. The cartridge media type that is represented by the last two letters of the cartridge bar code is used in this field. The string can be either “**^{6}XX\$**”, “**^{8}\$**”, or “unknown”. The “**^{6}XX\$**” represents any 6 characters followed by type “XX “, where “XX “ is one of the following cartridge media types: L5, L6, L7, L8, M8, L9, JB, JC, JD, JE, JK, JL, JY, or JZ. The “**^{8}\$**” represents any 8 characters, and means that any cartridge media type is acceptable. A value of “unknown” means that there is an error condition.
- ▶ **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`, `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.
- ▶ **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.
- ▶ **library\_name: <string>**  
Library name to which the pool belongs.
- ▶ **library\_id: <string>**  
Library ID (serial number) to which the pool belongs.
- ▶ **node\_group: <string>**  
Node group name to which the pool belongs.

The following parameters are available to be passed in to filter specific pools:

- ▶ **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 6-110 shows how to request the pools resource through `curl` commands.

*Example 6-110 REST pool command*

```
[root@tora ~]# curl -X GET 'http://localhost:7100/ibmsa/v1/pools?pretty=true'
[
 {
 "active_space": 0,
 "capacity": 0,
 "device_type": "3592",
 "fill_policy": "Default",
 "format_class": "60F",
 "free_space": 0,
 "id": "f244d0eb-e70d-4a7f-9911-0e3e1bd12720",
 "library_id": "65a7cbb5-8005-4197-b2a5-31c0d6f6e1c0",
 "library_name": "lib_tora",
 "low_space_warning_enable": false,
 "low_space_warning_threshold": 0,
 "media_restriction": "^.{8}$",
 "mode": "normal",
 "mount_limit": 0,
 "name": "pool3",
 "no_space_warning_enable": false,
 "nodegroup_name": "G0",
 "non_appendable_space": 0,
 "num_of_tapes": 1,
 "owner": "System",
 "reclaimable%": 0,
 "reclaimable_space": 0,
 "used_space": 0,
 "worm": "no"
 },
 {
 "active_space": 0,
 "capacity": 0,
 "device_type": "",
 "fill_policy": "Default",
 "format_class": "E08",
 "free_space": 0,
 "id": "bab56eb2-783e-45a9-b86f-6eaf4e8d316",
 "library_id": "65a7cbb5-8005-4197-b2a5-31c0d6f6e1c0",
 "library_name": "lib_tora",
 "low_space_warning_enable": false,
 "low_space_warning_threshold": 2199023255552,
 "media_restriction": "^.{8}$",
 "mode": "normal",
 "mount_limit": 0,
 "name": "pool1",
 "no_space_warning_enable": false,
 "nodegroup_name": "G0",
 "non_appendable_space": 0,
 "num_of_tapes": 0,
 "owner": "System",
 "reclaimable%": 0,
 "reclaimable_space": 0,
 "used_space": 0,
 }
]
```

```
 "worm": "no"
 }
]
```

---

Example 6-111 shows how to call pools with specifying specific fields and sorting the output in descending order.

*Example 6-111 REST API pools endpoint*

---

```
[root@tora ~]# 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": 0,
 "device_type": "3592",
 "free_space": 0,
 "library_name": "lib_tora",
 "name": "pool3",
 "num_of_tapes": 1
 },
 {
 "capacity": 0,
 "device_type": "",
 "free_space": 0,
 "library_name": "lib_tora",
 "name": "pool1",
 "num_of_tapes": 0
 }
]
```

---

## 6.26.2 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.

► **state:** <string>

The string indicates the state of the tape. The string can be either “appendable”, “append\_fenced”, “offline”, “recall\_only”, “unassigned”, “exported”, “full”, “data\_full”, “check\_tape\_library”, “need\_replace”, “require\_replace”, “require\_validate”, “check\_key\_server”, “check\_hba”, “inaccessible”, “non\_supported”, “duplicated”, “missing”, “disconnected”, “unformatted”, “label\_mismatch”, “need\_unlock”, and “unexpected\_cond”. If the string is “unexpected\_cond”, it is probably an error occurred.

- ▶ **status: <string>**  
The string indicates the severity level of the tape's status. The string can be either "error", "degraded", "warning", "info", or "ok".
- ▶ **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.
- ▶ **capacity: <number>**  
Capacity of the tape, in bytes.  $capacity = used\_space + free\_space$ .
- ▶ **Appendable: <string>**  
A tape that can be written in the format that is specified by the pool attributes, and on the cartridge media type that is specified by the pool attributes, is appendable. The format and the cartridge media type are provided as attributes of the pool to which the tape belongs. The string can be either "yes", "no", or it can be empty. If the tape falls into a state such as "append\_fenced" or "inaccessible", the string becomes "no". If the string is empty, the tape is not assigned to any pool.
- ▶ **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.
- ▶ **drive\_id: <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 tape offline.
- ▶ **task\_id: <string>**  
The `task_id` of the task which using this tape.

- ▶ **location\_type: <string>**  
The location type of the tape. The string can be either “robot”, “homeslot”, “ieslot”, “drive”, or empty. If the string is empty, the tape is missing.
- ▶ **library\_id: <string>**  
Library ID (serial number) to which the pool is belongs.
- ▶ **library\_name: <string>**  
Library name that the pool is belongs to
- ▶ **pool\_id: <string>**  
Pool ID to which this tape is assigned. Empty string if the tape is not assigned to any pool.

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.

### 6.26.3 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.
- ▶ **serial: <string>**  
The serial number of the library.
- ▶ **scsi\_vendor\_id: <string>**  
The vendor ID of the library.

- ▶ **scsi\_firmware\_revision:** <string>  
The firmware revision of the library.
- ▶ **num\_of\_drives:** <number>  
The number of tape drives that are assigned to the logical library.
- ▶ **num\_of\_ieslots:** <number>  
The number of I/E slots that are assigned to the logical library.
- ▶ **num\_of\_slots:** <number>  
The number of storage slots that are assigned to the logical library.
- ▶ **host\_device\_name:** <string>  
The Linux device name of the library.
- ▶ **host\_scsi\_address:** <string>  
The current data path to the tape drive from the assigned node, in the decimal notation of host.bus.target.lun.
- ▶ **errors:** <array>  
An array of strings that represents errors. If no error is found, an empty array is returned.

The available filtering parameter for the `libraries` endpoint is the name of the library.

## 6.26.4 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) to which the node group belongs.
- ▶ **library\_name:** <string>  
The name of the library to which the node group belongs.

The available filtering parameters for **nodegroups** are `name` (nodegroup name), and `library_name`. These parameters filter out nodegroups that do not meet the values that were passed in.

## 6.26.5 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 value in the 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).

▶ **port:** <number>

The port number for LTFS.

▶ **state:** <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.

▶ **enabled:** <bool>

True, if the node is enabled

▶ **library\_id:** <string>

The library ID of the library to which the node is attached.

▶ **library\_name:** <string>

The name of the library to which the node is attached.

▶ **nodegroup\_id:** <string>

The ID of the nodegroup to which the node is belongs.

▶ **nodegroup\_name:** <string>

The name of the nodegroup to which the node is belongs.

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.

## 6.26.6 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.
- ▶ **state: <string>**  
The drive state. For more information, see Table 10-1, “Status and state codes for eadm drive list” on page 317.
- ▶ **status: <string>**  
The string indicates the severity level of the tape's 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.
- ▶ **address: <number>**  
The address of the drive within the library.
- ▶ **tape\_barcode: <string>**  
The barcode of a tape to which the drive is mounted. Empty if no tape is mounted on the drive.
- ▶ **task\_id: <string>**  
The task\_id of the task which is using this drive.
- ▶ **library\_id: <string>**  
The ID of the library to which the drive belongs.
- ▶ **library\_name: <string>**  
The name of the library to which the drive belongs.
- ▶ **nodegroup\_name: <string>**  
The name of a nodegroup to which the drive is assigned.
- ▶ **node\_id: <string>**  
The ID of the node to which the drive is assigned.
- ▶ **node\_hostname: <string>**  
The host name of the node to which the drive is assigned. This field can be empty if a drive is not assigned to any node.
- ▶ **scsi\_vendor\_id: <string>**  
The vendor ID of the drive.
- ▶ **scsi\_product\_id: <string>**  
The product ID of the drive.

- ▶ **scsi\_firmware\_revision:** <string>  
The firmware revision of the drive.
- ▶ **host\_device\_name:** <string>  
The linux device name of the drive.
- ▶ **host\_scsi\_address:** <string>  
The current data path to the tape drive from the assigned node, in the decimal notation of `host.bus.target.lun`.

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.

## 6.26.7 Task endpoint

The tasks endpoint returns information about all or specific active tasks. The following is an example of calling the tasks endpoint:

```
curl -X GET 'http://localhost:7100/ibmsa/v1/tasks/'
curl -X GET 'http://localhost:7100/ibmsa/v1/tasks/<id>'
```

The following is the response data when requesting the tasks endpoint:

- ▶ **id:** <string>  
The task ID. Because the task id is overwritten after the task id reaches the upper limit, the id is in the format of `<task_id>@<created_time>`.
- ▶ **task\_id:** <number>  
The same value as the corresponding value in the IBM Spectrum Archive task ID.
- ▶ **type:** <string>  
The task type which meets the eeadm commands. A value of “unknown” means that there is an error condition.
- ▶ **cm\_param:** <string>  
The eeadm command that user input. If the type is “transparent\_recall”, the field is empty.
- ▶ **result:** <string>  
The result for completed tasks. The string can be either “succeeded”, “failed”, “aborted”, “canceled”, “suspended” or it can be empty. If the string is empty, the task is not completed.
- ▶ **status:** <string>  
The status for the task. The string can be either “waiting”, “running”, “interrupted”, “suspending”, “canceling”, “completed” or it can be empty. If the string is empty, an internal error occurs.
- ▶ **inuse\_libs:** <string array>  
The libraries that the task is currently using. The value is a string array of `library.name`.
- ▶ **inuse\_node\_groups:** <string array>  
The node groups that the task is currently using. The value is a string array of `nodegroup.name`.
- ▶ **inuse\_drives:** <string array>  
The drives that the task is currently using. The value is a string array of `drive.id`.

- ▶ **inuse\_pools:** <string array>  
The pools that the task is currently using. The value is a string array of pool.name.
- ▶ **inuse\_tapes:** <string array>  
The tapes that the task is currently using. The value is a string array of tape.barcode.
- ▶ **node\_hostname:** <string>  
The host name of the node to which the drive is assigned. This field can be empty if a drive is not assigned to any node.

All timestamps response data return in the following UTC format:

<yy>-<MM>-<dd>T<HH>:<mm>:<ss>.<SSS>Z

- ▶ **created\_time:** <string>  
The time when the task was accepted by MMM.
- ▶ **started\_time:** <string>  
The time when the task was started.
- ▶ **completed\_time:** <string>  
The time when the task was determined to be completed by MMM.

The following are filtering parameters for the tasks endpoint:

- ▶ **task\_id:** <number>  
If you specify the parameter and the task id was overwritten, you can get the latest one.  
If you specify the parameter and one of start\_created\_time and end\_created\_time, start\_created\_time and end\_created\_time do not work.
- ▶ **start\_created\_time:** <string>  
Filter out the tasks which created\_time is earlier than this parameter.  
The format is <yyyy>-<MM>-<dd>T<HH>:<mm>
- ▶ **end\_created\_time:** <string>  
Filter out the tasks that created\_time is later than this parameter.  
The format is <yyyy>-<MM>-<dd>T<HH>:<mm>
- ▶ **type:** <string>  
Filter the list of tasks by the task type.
- ▶ **result:** <string>  
Filter the list of tasks by the task result.
- ▶ **status:** <string>  
Filter the list of tasks by the task status.

## 6.27 File system migration

When refreshing the disk storage and/or IBM Spectrum Scale nodes, the IBM Spectrum Scale file system may need to be migrated to a new file system. The file system migration can be within the IBM Spectrum Scale cluster or between the two different IBM Spectrum Scale clusters.

The migration process of a file system managed by IBM Spectrum Archive used to be time consuming because all the files on the source file system needed to be recalled from tapes and copied to the target file system. From version 1.3.2.0, it supports a migration procedure using SOBAR feature of IBM Spectrum Scale. By using SOBAR, the stub files on the source file system are re-created on the target file system as stub files. The file system migration can be completed without recalling/copying file data from tapes.

The procedure is referred to as System Migration using SOBAR feature and requires prior review by IBM. For more information, see this [IBM Documentation web page](#).





## 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 6, “Managing daily operations of IBM Spectrum Archive Enterprise Edition” on page 129, and Chapter 9, “Troubleshooting IBM Spectrum Archive Enterprise Edition” on page 293. However, it is important to list them here in the context of hints, tips, and preferred practices.

This chapter includes the following topics:

- ▶ 7.1, “Preventing migration of the .SPACEMAN and metadata directories” on page 235
- ▶ 7.2, “Maximizing migration performance with redundant copies” on page 235
- ▶ 7.3, “Changing the SSH daemon settings” on page 237
- ▶ 7.4, “Setting mmapplypolicy options for increased performance” on page 237
- ▶ 7.5, “Preferred inode size for IBM Spectrum Scale file systems” on page 239
- ▶ 7.6, “Determining the file states for all files within the GPFS file system” on page 239
- ▶ 7.7, “Memory considerations on the GPFS file system for increased performance” on page 242
- ▶ 7.8, “Increasing the default maximum number of inodes in IBM Spectrum Scale” on page 242
- ▶ 7.9, “Configuring IBM Spectrum Scale settings for performance improvement” on page 243
- ▶ 7.10, “Use cases for mmapplypolicy” on page 244
- ▶ 7.11, “Capturing a core file on Red Hat Enterprise Linux with the Automatic Bug Reporting Tool” on page 247
- ▶ 7.12, “Anti-virus considerations” on page 248
- ▶ 7.13, “Automatic email notification with rsyslog” on page 248

- ▶ 7.14, “Overlapping IBM Spectrum Scale policy rules” on page 248
- ▶ 7.15, “Storage pool assignment” on page 250
- ▶ 7.16, “Tape cartridge removal” on page 250
- ▶ 7.17, “Reusing LTFS formatted tape cartridges” on page 251
- ▶ 7.18, “Reusing non-LTFS tape cartridges” on page 253
- ▶ 7.19, “Moving tape cartridges between pools” on page 254
- ▶ 7.20, “Offline tape cartridges” on page 254
- ▶ 7.21, “Scheduling reconciliation and reclamation” on page 255
- ▶ 7.22, “License Expiration Handling” on page 255
- ▶ 7.23, “Disaster recovery” on page 256
- ▶ 7.24, “IBM Spectrum Archive EE problem determination” on page 261
- ▶ 7.25, “Collecting IBM Spectrum Archive EE logs for support” on page 262
- ▶ 7.26, “Backing up files within file systems that are managed by IBM Spectrum Archive EE” on page 264
- ▶ 7.27, “IBM TS4500 Automated Media Verification with IBM Spectrum Archive EE” on page 266
- ▶ 7.28, “How to disable commands on IBM Spectrum Archive EE” on page 271
- ▶ 7.29, “LTO 9 Media Optimization” on page 272

**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/ltfsee/bin) to the PATH variable of the operating system.

## 7.1 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 IBM General Parallel File System (GPFS) file system by excluding these directories by having an IBM Spectrum Scale policy rule in place. Example 7-1 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/.l1tfsee"`.

*Example 7-1 IBM Spectrum Scale sample directory exclude statement in the migration policy file*

---

```
define(
 user_exclude_list,
 (
 PATH_NAME LIKE '/ibm/glues/.l1tfsee/%'
 OR PATH_NAME LIKE '%/.SpaceMan/%'
 OR PATH_NAME LIKE '%/.snapshots/%'
)
)
```

---

For more information and detailed examples, see 6.11.2, “Threshold-based migration” on page 168 and 6.11.3, “Manual migration” on page 173.

## 7.2 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:

```
eeadm pool show <poolname> [-l <libraryname>] [OPTIONS]
```

To set the mount limit setting for a tape cartridge pool, run the following command:

```
eeadm pool set <poolname> [-l <libraryname>] -a <attribute> -v <value>
```

To set the mount limit attribute to 2, run the **eeadm pool show** and **eeadm pool set** commands, as shown in Example 7-2.

*Example 7-2 Set the mount limit attribute to 2*

---

```
[root@saitama1 ~]# eeadm pool show pool1 -l lib_saitama
Attribute Value
poolname pool1
poolid 813ee595-2191-4e32-ae0e-74714715bb43
devtype 3592
mediarestriction none
format E08 (0x55)
worm no (0)
nodegroup G0
fillpolicy Default
owner System
mountlimit 0
lowspacewarningenable yes
lowspacewarningthreshold 0
nospacewarningenable yes
mode normal

[root@saitama1 ~]# eeadm pool set pool1 -l lib_saitama -a mountlimit -v 2

[root@saitama1 ~]# eeadm pool show pool1 -l lib_saitama
Attribute Value
poolname pool1
poolid 813ee595-2191-4e32-ae0e-74714715bb43
devtype 3592
mediarestriction none
format E08 (0x55)
worm no (0)
nodegroup G0
fillpolicy Default
owner System
mountlimit 2
lowspacewarningenable yes
lowspacewarningthreshold 0
nospacewarningenable yes
mode normal
```

---

## 7.3 Changing the SSH daemon settings

The default values for `MaxSessions` and `MaxStartups` are too low and must increase to allow for successful operations with IBM Spectrum Archive EE. `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. More connections are dropped until authentication succeeds or the `LoginGraceTime` expires for a connection. The default is 10:30:100, which indicates:

- ▶ 10 (start): Threshold of unauthenticated connections. The daemons will start to drop connections from this point on
- ▶ 30 (rate): Percentage chance of dropping once the start value is reached (increases linearly beyond start value)
- ▶ 100 (full): Maximum number of connections. All connection attempts will be dropped beyond this point

To change `MaxSessions` to 60 and `MaxStartups` to 1024, complete the following steps:

1. Edit the `/etc/ssh/sshd_config` file to set the `MaxSessions` and `MaxStartups` values:

```
MaxSessions = 60
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
```

## 7.4 Setting `mmappolicy` options for increased performance

The default values of the `mmappolicy` 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`, `mmappolicy` 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 `mmappolicy` 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` option, specify only the node class defined for IBM Spectrum Archive EE nodes. This does not apply for cases where `-I defer` or `-I prepare` option is used.

► **-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), the global work directory is the `.mmSharedTmpDir` directory 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 `mmappolicy` 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.

## 7.5 Preferred inode size for IBM Spectrum Scale file systems

When you create the GPFS file systems, an option is available that is called `-i InodeSize` for the `mmcrfs` command. The option specifies the byte size of inodes. By default, the inode size is 4 KB and it consists of a fixed 128 byte header, plus data, such as disk addresses pointing to data, or indirect blocks, or extended attributes.

The supported inode sizes are 512, 1024, and 4096 bytes. Regardless of the file sizes, the preferred inode size is 4096 for all IBM Spectrum Scale file systems for IBM Spectrum Archive EE. This size should include user data file systems and the IBM Spectrum Archive EE metadata file system.

## 7.6 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 `eeadm file state` command. However, it is not practical to run this command for every file on the GPFS file system.

In Example 7-3, the file is in the migrated state and is only on the tape cartridge JD0321JD.

*Example 7-3 Example of the `eeadm info files` command*

---

```
[root@saitama1 prod]# eeadm file state LTFS_EE_FILE_2dEPRHhh_M.bin
Name: /ibm/gpfs/prod/LTFS_EE_FILE_2dEPRHhh_M.bin
State: premigrated
ID: 11151648183451819981-3451383879228984073-1435527450-974349-0
Replicas: 1
Tape 1: JD0321JD@test4@lib_saitama (tape state=appendable)
```

---

Thus, use list rules in an IBM Spectrum Scale policy instead. Example 7-4 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 7-4 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 7-5. If the policy is run daily, this script can be modified to keep several versions to be used for history purposes.

*Example 7-5 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 `mmapp1ypol1cy` 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 7-6.

*Example 7-6 Sample output of the /root/file\_states.txt file*

---

```
355150 165146835 0 dir -- /ibm/gpfs/prod
974348 2015134857 0 premigrated 1
JD0321JD@1d85a188-be4e-4ab6-a300-e5c99061cec4@ebc1b34a-1bd8-4c86-b4fb-bee7b60c24c7
-- /ibm/gpfs/prod/LTFS_EE_FILE_9_rzu.bin
974349 1435527450 0 premigrated 1
JD0321JD@1d85a188-be4e-4ab6-a300-e5c99061cec4@ebc1b34a-1bd8-4c86-b4fb-bee7b60c24c7
-- /ibm/gpfs/prod/LTFS_EE_FILE_2dEPRHhh_M.bin
974350 599546382 0 premigrated 1
JD0321JD@1d85a188-be4e-4ab6-a300-e5c99061cec4@ebc1b34a-1bd8-4c86-b4fb-bee7b60c24c7
-- /ibm/gpfs/prod/LTFS_EE_FILE_XH7Qwj5y9j2wqV4615rCxPMir039xL1t
```

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.

## 7.7 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.

## 7.8 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.

Consider the following points when managing inodes:

- ▶ 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.
- ▶ Excessively increasing the value for the maximum number of inodes 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]
```

## 7.9 Configuring IBM Spectrum Scale settings for performance improvement

The performance results in 3.7.2, “Planning for LTO-9 Media Initialization/Optimization” on page 63 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
- ▶ `maxFilesToCache` = 128k

For example, a 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.

Refer to IBM Spectrum Scale [Documentation](#) for more details on cache related parameters such as `maxFilesToCache` and `maxStatCache`.

## 7.10 Use cases for mmapplypolicy

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.

### 7.10.1 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 176 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 7-7 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 example, if you need to have files on three storage pools, modify the `OPTS` parameter to include a third `<pool>@<library>`.

*Example 7-7 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/eadm'
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
```

---

## 7.10.2 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 176.

Example 7-8 shows a simple premigration policy for files greater than 5 MB.

*Example 7-8 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/eedm'
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
```

---

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 commands are used for setting a fail-safe policy and calling `mmadcallback`:

- ▶ `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 7-9 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 7-9 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/eadm'
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
```

---

### 7.10.3 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 extra exclude directories wherever migrations are occurring. Example 7-10 uses the same migration policy that is shown in Example 7-7 on page 244 with the addition of extra exclude and check parameters.

*Example 7-10 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/eadm'
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
```

---

## 7.11 Capturing a core file on Red Hat Enterprise Linux with the Automatic Bug Reporting Tool

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
```

## 7.12 Anti-virus 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 Hierarchical Storage Management (HSM) work directories:
  - The library mount point (the /l t f s 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.

## 7.13 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 [this website](#).

Here is an example 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 7-11.

*Example 7-11 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
```

---

## 7.14 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 7-12, an attempt is made to migrate four files to tape cartridge pool pool2. Before the migration attempt, Tape ID JCB610JC is already in tape cartridge pool pool1, and Tape ID JD0321JD in pool2 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 **eeadm file state** command in Example 7-12.

*Example 7-12 Display the state of files by using the eeadm file state command*

---

```
[root@saitama1 prod]# eeadm file state *.bin
Name: /ibm/gpfs/prod/fileA.ppt
State: migrated
ID: 11151648183451819981-3451383879228984073-1435527450-974349-0
Replicas: 1
Tape 1: JCB610JC@pool1@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/fileB.ppt
State: migrated
ID: 11151648183451819981-3451383879228984073-2015134857-974348-0
Replicas: 1
Tape 1: JCB610JC@pool1@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/fileC.ppt
State: migrated
ID: 11151648183451819981-3451383879228984073-599546382-974350-0
Replicas: 1
Tape 1: JD0321JD@pool2@lib_saitama (tape state=appendable)

Name: /ibm/gpfs/prod/fileD.ppt
State: premigrated
ID: 11151648183451819981-3451383879228984073-2104982795-3068894-0
Replicas: 1
Tape 1: JD0321JD@pool2@lib_saitama (tape state=appendable)
```

---

The mig scan list file that is used in this example contains these entries, as shown in Example 7-13.

*Example 7-13 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 7-14.

*Example 7-14 Migration of files by running the eeadm migration command*

---

```
[root@saitama1 prod]# eeadm migrate mig -p pool2@lib_saitama
2021-12-21 01:19:12 GLESL700I: Task migrate was created successfully, task ID is 1074.
2021-12-21 01:19:13 GLESM896I: Starting the stage 1 of 3 for migration task 1074 (qualifying the state of migration candidate files).
```

```
2021-12-21 01:19:13 GLESM897I: Starting the stage 2 of 3 for migration task 1074
(copying the files to 1 pools).
2021-12-21 01:19:13 GLESM898I: Starting the stage 3 of 3 for migration task 1074
(changing the state of files on disk).
2021-12-21 01:19:13 GLESL840E: Failed to process the requested 4 file(s), with 2
succeeding and 2 failing.
2021-12-21 01:19:13 GLESL841I: Succeeded: 1 migrated, 1 already_migrated.
2021-12-21 01:19:13 GLESL843E: Failed: 0 duplicate, 2 wrong_pool, 0 not_found, 0
too_small, 0 too_early, 0 other_failure
```

---

The files on Tape ID JCB610JC (fileA.ppt and fileB.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 JD0321JD, the attempt to migrate fileC.ppt file produces a already\_migrated or a migration result fail for some cases with multiple library environments because the file is already migrated. Only the attempt to migrate the pre-migrated fileD.ppt file succeeds. Therefore, one operation succeeded and three other operations result with wrong\_pool or already\_migrated.

## 7.15 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.

## 7.16 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 6.8.2, “Moving tape cartridges” on page 156, and “The eeadm <resource type> --help command” on page 316.

## 7.16.1 Reclaiming tape cartridges before you remove or export them

To avoid failed recall operations, it is recommended that the tape cartridges are reclaimed before removing or exporting a cartridge.

When a cartridge that can be appended is planned for removal, use one of the following methods to perform reclaim before the removal:

- ▶ Run the **eeadm tape reclaim** command before you remove it from the LTFS file system (by running the **eeadm tape unassign** command)
- ▶ Export it from the LTFS library (by running the **eeadm tape export** command. This command will internally run the reclaim during its task.)

If tape cartridges are in a `need_replace` or a `require_replace` state, use the **eeadm tape replace** command. This command will internally run the **reclaim** command during its procedure.

The **eeadm tape unassign --safe-remove** can be used for cases where the `replace` command fails. Note that the **--safe-remove** option will recall all the active files on the tape back to an IBM Spectrum Scale file system that has adequate free space, and those files need to be manually migrated again to a good tape.

## 7.16.2 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.

## 7.17 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, the data must be moved off the cartridge or deleted from the file system then the cartridges must be reformatted before they are added to an IBM Spectrum Archive EE tape cartridge pool. This task can be done by running the **eeadm tape reclaim** or **eeadm tape unassign -E** commands. Note that the tapes removed with the **-E** option will need a **-f** option when re-assigning with the **eeadm tape assign** command.

## 7.17.1 Reformatting LTFS tape cartridges through eeadm 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 7-15.

*Example 7-15 Display content of a used LTFS tape cartridge (non-IBM Spectrum Archive EE)*

---

```
[root@l1tfs97 ~]# ls -la /l1tfs/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 7-16. This example indicates that this cartridge was previously used as an IBM Spectrum Archive EE tape cartridge.

*Example 7-16 Display content of a used LTFS tape cartridge (IBM Spectrum Archive EE)*

---

```
[root@l1tfs97 l1tfs]# ls -lsa /l1tfs/JD0321JD
total 0
0 drwxrwxrwx 4 root root 0 Jan 9 14:33 .
0 drwxrwxrwx 7 root root 0 Dec 31 1969 ..
0 drwxrwxrwx 3 root root 0 Jan 9 14:33 ibm
0 drwxrwxrwx 2 root root 0 Jan 10 16:01 .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 `eeadm tape list` command in the status unassign, as shown in Example 7-17.

*Example 7-17 Output of the eeadm tape list command*

---

```
[root@saitama2 l1tfs]# eeadm tape list -l lib_saitama
Tape ID Status State Usable(GiB) Used(GiB) Available(GiB) Reclaimable% Pool Library Location Task ID
JCA561JC ok offline 0 0 0 0% pool2 lib_saitama homeslot -
JCA224JC ok appendable 6292 0 6292 0% pool1 lib_saitama homeslot -
JCC093JC ok appendable 6292 496 5796 0% pool1 lib_saitama homeslot -
JCB141JC ok unassigned 0 0 0 0% - lib_saitama homeslot -
```

---

## Reformatting and reusing an LTFS SDE/LE tape cartridge

In this case, you run the **eeadm tape assign** command to add this tape cartridge to an IBM Spectrum Archive EE tape cartridge pool and format it at the same time, as shown in the following example and in Example 7-18:

```
eeadm tape assign <list_of_tapes> -p <pool> [OPTIONS]
```

If the format fails then there is data most likely written to the tape and a force format is required by appending the **-f** option to the **eeadm tape assign** command.

*Example 7-18 Reformat a used LTFS SDE/LE tape cartridge*

---

```
[root@saitama2 ~]# eeadm tape assign JCB141JC -p pool1 -l lib_saitama -f
2019-01-15 08:35:21 GLESL700I: Task tape_assign was created successfully, task id
is 7201.
2019-01-15 08:38:09 GLESL087I: Tape JCB141JC successfully formatted.
2019-01-15 08:38:09 GLESL360I: Assigned tape JCB141JC to pool pool1 successfully.
```

---

## Reformatting and reusing an IBM Spectrum Archive EE tape cartridge

If you want to reuse an IBM Spectrum Archive EE tape cartridge, the user can either reclaim the tape so the data is preserved onto another cartridge within the same pool. If the data on the cartridge is no longer needed then the user must delete all files on disk that have been premigrated/migrated to the cartridge and run the **eeadm tape unassign** command with the **-E** option. After the tape has been removed from the pool using the **eeadm tape unassign -E** command add the tape to the new pool using the **eeadm tape assign** command with the **-f** option to force a format.

## 7.18 Reusing non-LTFS tape cartridges

For your IBM Spectrum Archive EE setup, in some scenarios, you might want to reuse tape cartridges that were used before as non-LTFS formatted media in another server setup behind your tape library (such as backup tape cartridges from 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 6.8.1, “Adding tape cartridges” on page 154.

## 7.19 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 6.8.2, “Moving tape cartridges” on page 156.

### 7.19.1 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. This is not internally allowed in IBM Spectrum Archive EE.

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 6.9, “Tape storage pool management” on page 162.

### 7.19.2 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 7.19.1, “Avoiding changing assignments for tape cartridges that contain files” on page 254.

For more information, see 6.9, “Tape storage pool management” on page 162 and 6.17, “Reclamation” on page 200.

## 7.20 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.

### 7.20.1 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.

### 7.20.2 Solving problems

For more information about solving problems that are caused by trying to import a tape cartridge in offline state that was modified while it was outside the library, see “Importing offline tape cartridges” on page 203.

## 7.21 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 6.15, “Recalling files to their resident state” on page 196 and 6.17, “Reclamation” on page 200.

## 7.22 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`

Once in this state, some commands will return errors with messages indicating licence expiration.

- ▶ 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`
- ▶ `save`

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. Other task that inherit such behaviors will also fail, since these type of commands that include writing to tape will be invalid.

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.

## 7.23 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.

### 7.23.1 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 listed in Table 7-1.

*Table 7-1 Summary of disaster recovery tiers (SHARE)*

| <b>Tier</b> | <b>Description</b>                                     |
|-------------|--------------------------------------------------------|
| 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 (off-site 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 off-site storage facility, and determines some recovery requirements. As shown in Figure 7-1 on page 257, backups are taken that are stored at an off-site 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.

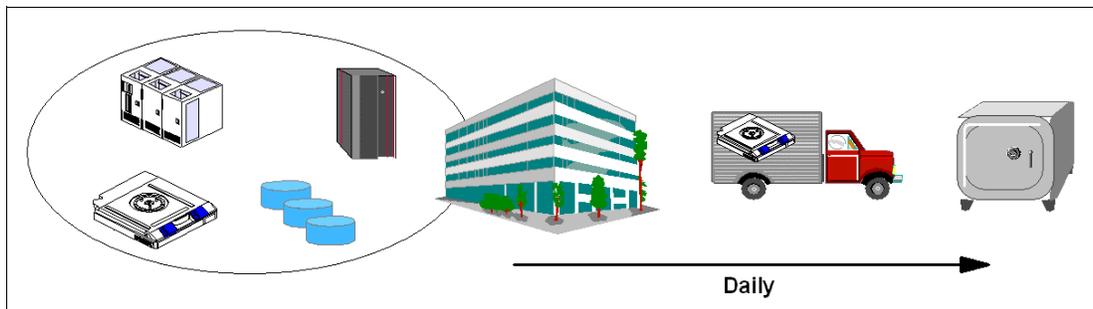


Figure 7-1 Tier 1 - offsite vaulting (PTAM)

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.

## 7.23.2 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 6.11.4, “Replicas and redundant copies” on page 177). 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 copy 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 copy 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 copy 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 copy 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 copy 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 7.23.3, “IBM Spectrum Archive EE tier 1 DR procedure” on page 259 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 7-19 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 176.

*Example 7-19 Sample of a migration policy*

---

```
define(user_exclude_list,(PATH_NAME LIKE '/ibm/gpfs/.lzfsee/%' 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/lzfsee/bin/eadm'
OPTS '-p primary@lib_lzfseevm,copy@lib_lzfseevm,DR@lib_lzfseevm'
SIZE(20971520)

RULE 'LTFSEE_FILES_RULE' MIGRATE FROM POOL 'system' TO POOL 'LTFSEE_FILES' WHERE
(
is_premigrated
AND NOT user_exclude_list
)
```

---

After you create redundant copies 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 6.19.2, “Exporting tape cartridges” on page 204.

**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 **eadm tape move** command with the option **-L ieslot**. For more information, see 6.8, “Tape library management” on page 154 and 10.1, “Command-line reference” on page 316.

### 7.23.3 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 exported tapes that was exported from IBM Spectrum Archive EE must be stored in an offline vault. 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:
  - # eadm node list
  - # eadm tape list
  - # eadm pool list
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:

```
eeadm library rescan
```

6. Move the inserted tapes to homeslot with the following command:

```
eeadm tape move <tape id> -L homeslot
```

7. Import the DR tape cartridges into the IBM Spectrum Archive EE environment by running the **eeadm tape import** command. The **eeadm tape import** 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 10, “Reference” on page 315.

When you rebuild from one or more tape cartridges, the **eeadm tape import** 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.

## 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) that are located in an exported tape can be restored from tapes by running the **eeadm tape import** command<sup>1</sup>.

The **eeadm tape import** command reinstantiates the stub files in IBM Spectrum Scale for migrated files. The state of those files changes to the migrated state. Also, the **eeadm tape import** 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 **eeadm tape import** 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.

### *Recovery procedure by using the eeadm tape import command*

Here is a typical user scenario for recovering migrated files and saved file system objects from tape by running the **eeadm tape import** 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 **eeadm tape import** command:

```
eeadm tape import LTFS01L6 LTFS02L6 LTFS03L6 -p PrimPool -P
/gpfs/ltfsee/rebuild
```

/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.

---

<sup>1</sup> Note that only exported tapes are valid for this command.

### ***Import processing for unexported tapes that are not reconciled***

The **eeadm tape import** 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 **eeadm tape import** 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 **eeadm tape import** command restores files or file system objects that were removed from IBM Spectrum Scale.

### ***Rebuild and restore considerations***

While the **eeadm tape import** 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.

For more information about the **eeadm tape import** command, see “Importing” on page 202.

## **7.24 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 7-20, 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 7-20 Check the ltfsee.log file***

---

```
[root@mikasa1 gpfs]# eeadm migrate mig -p myfirstpool@lib_saitama
2019-01-21 08:37:54 GLESL700I: Task migrate was created successfully, task id is 7217.
2019-01-21 08:37:55 GLESM896I: Starting the stage 1 of 3 for migration task 7217
(qualifying the state of migration candidate files).
2019-01-21 08:37:55 GLESM897I: Starting the stage 2 of 3 for migration task 7217 (copying
the files to 1 pools).
2019-01-21 08:38:36 GLESM898I: Starting the stage 3 of 3 for migration task 7217 (changing
the state of files on disk).
2019-01-21 08:38:36 GLESL159E: Not all migration has been successful.
2019-01-21 08:38:36 GLESL038I: Migration result: 1succeeded, 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 gpfs]#
[root@ltfs97 gpfs]# vi /var/log/ltfsee.log
2019-01-21T08:37:55.399423-07:00 saitama2 mmm[22236]: GLESM148E(00704): File
/ibm/gpfs/document20.txt is already migrated and will be skipped.
2019-01-21T08:38:36.694378-07:00 saitama2 mmm[22236]: GLESL159E(00142): Not all migration
has been successful.
```

---

In Example 7-20, 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:

```
2019-01-21T08:37:55.399423-07:00 saitama2 mmm[22236]: GLESM148E(00704): File
/ibm/gpfs/document20.txt is already migrated and will be skipped.
```

For more information about problem determination, see Chapter 9, “Troubleshooting IBM Spectrum Archive Enterprise Edition” on page 293.

### 7.24.1 Rsyslog log suppression by rate-limiting

IBM Spectrum Archive uses rsyslogd and journald of the Red Hat Enterprise Linux system for logging. By default, rate-limiting of the log messages is enabled. However, this causes problems because all the logs are needed for problem analysis.

It is highly recommended to disable the rate-limiting so that no logs are suppressed, as follows:

1. Open `/etc/systemd/journald.conf` and add the following lines:

```
RateLimitInterval=0
RateLimitBurst=0
```

2. Open `/etc/rsyslog.conf`, and in the Global Directives section, add the following lines:

```
$imjournalRateLimitInterval 0
$imjournalRateLimitBurst 0
```

3. Restart the services:

```
systemctl restart systemd-journald
systemctl restart rsyslog.service
```

## 7.25 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:

```
1tfsee_log_collection
```

Example 7-21 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 7-21 The ltfsee\_log\_collection command*

---

```
[root@kyoto ~]# ltfsee_log_collection
```

```
IBM Spectrum Archive Enterprise Edition - log collection program
```

```
This program collects the following information from your IBM Spectrum Scale (GPFS) cluster.
```

- (1) Log files that are generated by IBM Spectrum Scale (GPFS), and IBM Spectrum Archive EE
- (2) Configuration information that is configured to use IBM Spectrum Scale (GPFS) and IBM Spectrum Archive EE.
- (3) System information including OS distribution and kernel information, hardware information (CPU and memory) and process information (list of running processes).
- (4) Task information files under the following subdirectory <GPFS mount point>/ltfsee/statesave

```
If you agree to collect all the information, input 'y'.
```

```
If you agree to collect only (1) and (2), input 'p' (partial).
```

```
If you agree to collect only (4) task information files, input 't'.
```

```
If you don't agree to collect any information, input 'n'.
```

```
The following files are collected only if they were modified within the last 90 days.
```

- /var/log/messages\*
- /var/log/ltfsee.log\*
- /var/log/ltfsee\_trc.log\*
- /var/log/ltfsee\_mon.log\*
- /var/log/ltfs.log\*
- /var/log/ltfsee\_install.log
- /var/log/ltfsee\_stat\_driveperf.log\*
- /var/log/httpd/error\_log\*
- /var/log/httpd/rest\_log\*
- /var/log/ltfsee\_rest/rest\_app.log\*
- /var/log/logstash/\*

```
You can collect all of the above files, including files modified within the last 90 days, with an argument of 'all'.
```

```
 #./ltfsee_log_collection all
```

```
If you want to collect the above files that were modified within the last 30 days.
```

```
 #./ltfsee_log_collection 30
```

```
The collected data will be zipped in the ltfsee_log_files_<date>_<time>.tar.gz file. You can check the contents of the file before submitting it to IBM.
```

```
Input > y
```

```
Creating a temporary directory '/root/ltfsee_log_files'...
```

```
The collection of local log files is in progress.
```

```
Removing collected files...
```

```
Information has been collected and archived into the following file.
```

```
ltfsee_log_files_20190121_084618.tar.gz
```

---

## 7.26 Backing up files within file systems that are managed by IBM Spectrum Archive EE

The IBM Spectrum Protect Backup/Archive client and the IBM Spectrum Protect 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 files within the GPFS or IBM Spectrum Scale file systems. The **mmbackup** command can be used to back up some or all of the files of a GPFS or IBM Spectrum Scale file system to IBM Spectrum Protect servers using the IBM Spectrum Protect Backup-Archive client. After files have been backed up, you can restore them using the interfaces provided by IBM Spectrum Protect.

The **mmbackup** command utilizes all the scalable, parallel processing capabilities of the **mmapplypolicy** command to scan the file system, evaluate the metadata of all the objects in the file system, and determine which files need to be sent to backup in IBM Spectrum Protect, as well which deleted files should be expired from IBM Spectrum Protect. Both backup and expiration take place when running **mmbackup** in the incremental backup mode.

The **mmbackup** command can inter-operate with regular IBM Spectrum Protect commands for backup and expire operations. However if after using **mmbackup**, any IBM Spectrum Protect incremental or selective backup or expire commands are used, **mmbackup** needs to be informed of these activities. Use either the **-q** option or the **--rebuild** option in the next **mmbackup** command invocation to enable **mmbackup** to rebuild its shadow databases.

These databases shadow the inventory of objects in IBM Spectrum Protect so that only new changes will be backed up in the next incremental **mmbackup**. Failing to do so will needlessly back up some files again. The shadow database can also become out of date if **mmbackup** fails due to certain IBM Spectrum Protect server problems that prevent **mmbackup** from properly updating its shadow database after a backup. In these cases it is also required to issue the next **mmbackup** command with either the **-q** option or the **--rebuild** options.

The **mmbackup** command provides the following benefits:

- ▶ A full backup of all files in the specified scope.
- ▶ An incremental backup of only those files that have changed or been deleted since the last backup. Files that have changed since the last backup are updated and files that have been deleted since the last backup are expired from the IBM Spectrum Protect server.
- ▶ Utilization of a fast scan technology for improved performance.
- ▶ The ability to perform the backup operation on a number of nodes in parallel.
- ▶ Multiple tuning parameters to allow more control over each backup.
- ▶ The ability to backup the read/write version of the file system or specific global snapshots.
- ▶ Storage of the files in the backup server under their GPFS root directory path independent of whether backing up from a global snapshot or the live file system.
- ▶ Handling of unlinked filesets to avoid inadvertent expiration of files.

For more information, see [IBM Documentation](#).

## 7.26.1 Considerations

Consider the following points when you are using the IBM Spectrum Protect Backup/Archive client in the IBM Spectrum Archive EE environment:

- ▶ IBM Spectrum Protect requirements for backup
- ▶ Update the `dsm.sys` and `dsm.opt` files to support both IBM Spectrum Protect and IBM Spectrum Archive EE operations:

**Note:** In `dsm.sys` file, when multiple server stanza are defined for IBM Spectrum Archive and IBM Spectrum Protect, the following lines need to be placed at the beginning of the file. Otherwise, IBM Spectrum Archive will not be able to migrate or recall files.

```
HSMBACKENDMODE TSMFREE
ERRORLOGNAME /opt/tivoli/tsm/client/hsm/bin/dsmerror.log
MAXRECALLDAEMONS 64
MINRECALLDAEMONS 64
HSMmigZEROBLOCKFILES YES
errorlogretention 180 S
```

- ▶ Ensure the files are backed up first with IBM Spectrum Protect followed by archiving the files with IBM Spectrum Archive EE to avoid recall storms.

## 7.26.2 Backing up a GPFS or IBM Spectrum Scale environment

The best practice is to always backup the files using IBM Spectrum Protect first and then archiving the files using IBM Spectrum Archive EE. The primary reason is that attempting to back up the stub of a file that was migrated to IBM Spectrum Archive EE 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.

The `mmbackup` command is used to back up the files of a GPFS or IBM Spectrum Scale file system to IBM Spectrum Protect servers by using the IBM Spectrum Protect Backup/Archive Client of the IBM Spectrum Protect family. In addition, the `mmbackup` command can operate with regular IBM Spectrum Protect backup commands for backup. After a file system is backed up, you can restore files by using the interfaces that are provided by the IBM Spectrum Protect family.

Starting with IBM Spectrum Archive EE v1.3.0.0 and IBM Spectrum Scale v5.0.2.2, a new option has been added to the `eeadm migrate` command called `--mmbackup`. With the `--mmbackup` option is supplied, IBM Spectrum Archive EE will first verify that current backup versions of the files exist within IBM Spectrum Protect before it will archive them to IBM Spectrum Archive EE. If those files are not backed up, those files will be filtered out and thus not archived to tape. This ensures that there will be no recall storm due to the backup of files.

The `--mmbackup` option of the `eeadm migrate` command takes 1 option, which is the location of the `mmbackup` shadow database. This location is normally the same device or directory option of the `mmbackup` command.

## 7.27 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 more 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 7-2 shows an example of such a setup.

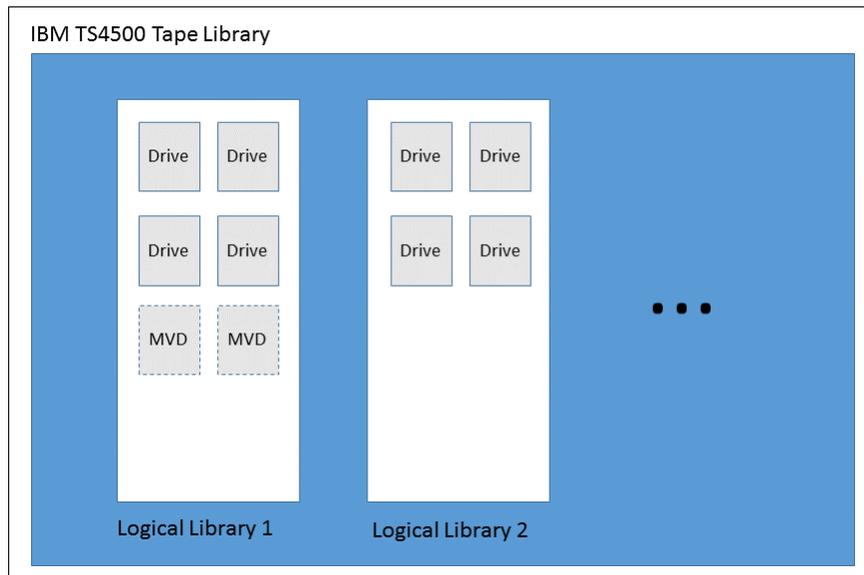


Figure 7-2 TS4500 with one logical library showing two MVDs configured

**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 7-3 - Figure 7-6 on page 269 are examples 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 to the logical library (if it was not assigned before). Complete the following steps:

1. 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 7-3.

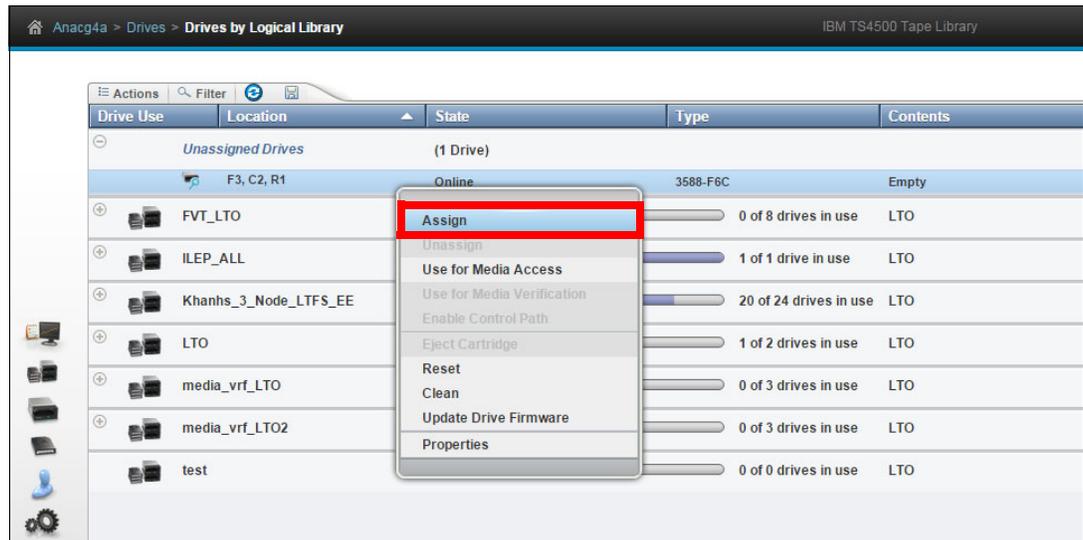


Figure 7-3 Assign a tape drive to a logical library through the TS4500 web interface (step 1)

2. 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 7-4.

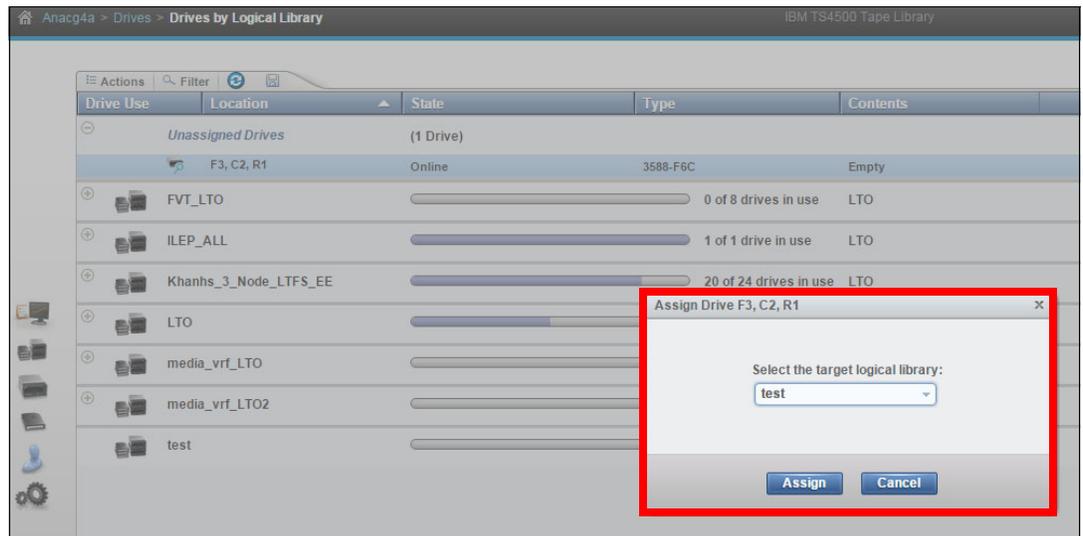


Figure 7-4 Assign a tape drive to a logical library through the TS4500 web interface (step 2)

- If the drive to be used as an MVD is configured within the logical library, change its role, as shown in Figure 7-5 and Figure 7-6 on page 269.

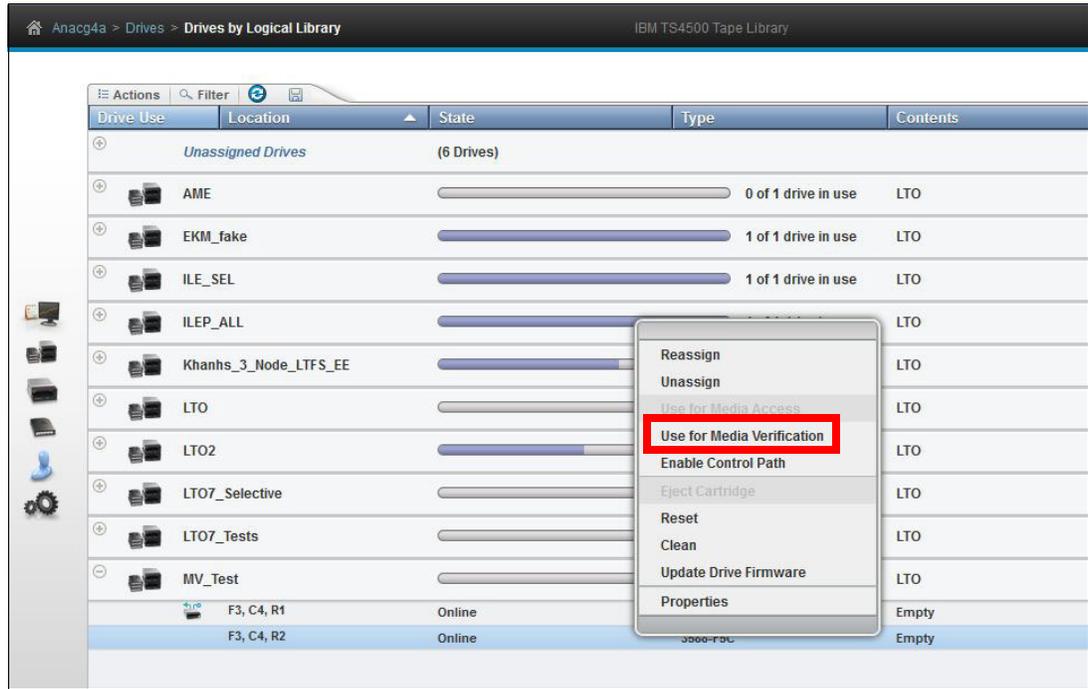


Figure 7-5 Reserve a tape drive as the media verification drive through the TS4500 web interface

- 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 7-6).

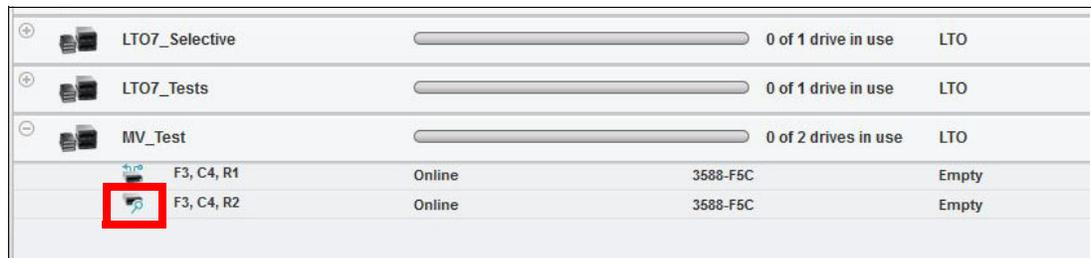


Figure 7-6 Display a tape drive as the media verification drive through the TS4500 web interface

**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 7-7 shows you the TS4500 web interface dialog box for enabling automatic media verification on a 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.

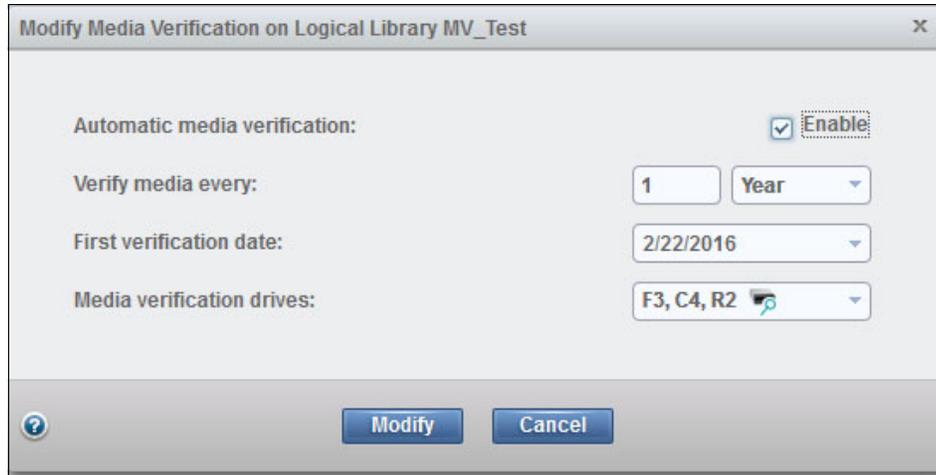


Figure 7-7 Modify Media Verification dialog box to set up a schedule

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 7-8.

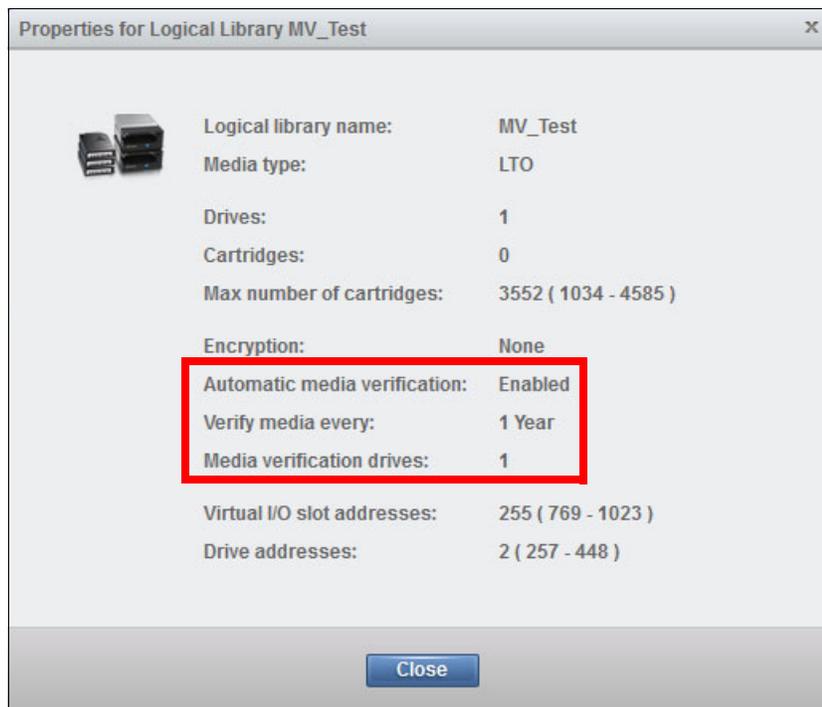


Figure 7-8 TS4500 properties for a logical library

For more information and the usage of the TS4500 media verification functions, see IBM TS4500 R8 Tape Library Guide, SG24-8235 and IBM TS4500 documentation at [IBM Documentation](#).

## 7.28 How to disable commands on IBM Spectrum Archive EE

Some IBM Spectrum Archive EE commands can be disabled or enabled, such as the transparent recall command.

The control is done by disabling or enabling corresponding task types by using the **eeadm cluster set** command. The commands, corresponding task types, and corresponding attribute names that are used by the **eeadm cluster set** command are listed in Table 7-2.

Table 7-2 *eeadm commands, task types, and attribute names*

| File access or eeadm command            | Task type          | Attribute name           |
|-----------------------------------------|--------------------|--------------------------|
| File access that would trigger a recall | transparent_recall | allow_transparent_recall |
| eeadm migrate                           | migrate            | allow_migrate            |
| eeadm premigrate                        | premigrate         | allow_premigrate         |
| eeadm recall                            | selective_recall   | allow_selective_recall   |
| eeadm save                              | save               | allow_save               |
| eeadm tape assign                       | tape_assign        | allow_tape_assign        |
| eeadm tape datamigrate                  | tape_datamigrate   | allow_tape_datamigrate   |
| eeadm tape export                       | tape_export        | allow_tape_export        |
| eeadm tape import                       | tape_import        | allow_tape_import        |
| eeadm tape offline                      | tape_offline       | allow_tape_offline       |
| eeadm tape online                       | tape_online        | allow_tape_online        |
| eeadm tape reconcile                    | tape_reconcile     | allow_tape_reconcile     |
| eeadm tape reclaim                      | tape_reclaim       | allow_tape_reclaim       |
| eeadm tape replace                      | tape_replace       | allow_tape_replace       |
| eeadm tape unassign                     | tape_unassign      | allow_tape_unassign      |

The attributes can be set to “yes” or “no” by using the **eeadm cluster set** command. The current setting can be verified by using the **eeadm cluster show** command.

When a task type is disabled, the next command immediately fails. The failed task can be verified by using the **eeadm task list -c** command and the **eeadm task show** command.

Example 7-22 shows the results of the `eeadm cluster show` and `eeadm task show` commands when the `allow_transparent_recall` option is set to `no`.

*Example 7-22 Disabling transparent recall*

---

```
[root@server dir1]# eeadm cluster set -a allow_transparent_recall -v no
2019-11-26 19:52:39 GLESL802I: Updated attribute allow_transparent_recall.
```

```
[root@server dir1]# eeadm cluster show | grep -E "Attribute|transparent"
Attribute Value
allow_transparent_recall no
```

```
[root@server dir1]# cat file2
cat: file2: Permission denied
```

```
[root@server dir1]# eeadm task show 1622
=== Task Information ===
Task ID: 1622
Task Type: transparent_recall
Command Parameters: dsmrecalld
Status: completed
Result: failed
Accepted Time: Thu Nov 14 21:40:34 2019 (+0900)
Started Time: Thu Nov 14 21:40:34 2019 (+0900)
Completed Time: Thu Nov 14 21:40:34 2019 (+0900)
Workload: 7 bytes of file (name: /ibm/gpfs0/archive/dir1/file2, inode:
215804)
Progress: -
Result Summary: Disabled
```

---

## 7.29 LTO 9 Media Optimization

Each new LTO 9 cartridge requires a one-time initialization called media optimization prior to commencing read/write operations. The media optimization has been implemented in LTO 9 to optimize data placement to each LTO 9 cartridge characteristics. **The optimization averages between 35 to 52 minutes per first load of a cartridge to a tape drive.** Most initializations will complete within 60 minutes, but the whole process may take up to 2 hours. IBM recommends that all media optimizations be performed in the destination ecosystem with a drive or drives in the acclimated environment.

For more information about LTO 9 media optimization, see [this web page](#).

When using LTO 9 drives/cartridges in IBM Spectrum Archive managed systems, the media optimization is performed as part of `eeadm tape assign` command. When deploying LTO-9 cartridges, the extra time the command may take due to the media optimization needs to be accounted for.

One method for deploying multiple LTO 9 tapes into a large tape library managed by IBM Spectrum Archive is as follows:

1. Use all of the available LTO 9 drives to perform the `eeadm tape assign` command to prepare enough LTO 9 tapes for several days.
2. Use IBM Spectrum Archive drive role feature to dedicate one or a few drives to keep performing the `eeadm tape assign` command while using other drives for migration and recall. The drive role can be set by the `eeadm drive set` command.
  - Set "g" to the drives that perform tape assign
  - Set "mr" to the other drives for migration/recall

**Note:** Consider the following points:

- ▶ Even if a pre-optimized LTO 9 cartridge (re-using or initialized somewhere else) is deployed, the `eeadm tape assign` command will re-format and re-optimize the tape.
- ▶ The `eeadm tape assign` command first determines if it is OK to format the specified tape by checking if the tape already has LTFS files stored. The check is a safe-guard to avoid unexpectedly formatting a tape that has valid LTFS data. However, when deploying a brand new LTO 9 tape, the check results in performing the media optimization twice. Therefore, when assigning a brand new LTO 9 tape, it is recommended to specify a `--force-format` or `-f` option to the command to bypass the initial check.

Also, the `eeadm tape reclaim` command formats the source tape at the last phase of the command. If the source tape is LTO 9, the format will trigger the media optimization.





# IBM Spectrum Archive Enterprise Edition use cases

This chapter describes various use case examples for IBM Spectrum Archive Enterprise Edition (IBM Spectrum Archive EE).

This chapter includes the following topics:

- ▶ 8.1, “Overview of use cases” on page 276
- ▶ 8.2, “Media and Entertainment” on page 279
- ▶ 8.3, “Media and Entertainment” on page 280
- ▶ 8.4, “High-Performance Computing” on page 281
- ▶ 8.5, “Healthcare” on page 283
- ▶ 8.6, “Genomics” on page 284
- ▶ 8.7, “Archive of research and scientific data for extended periods” on page 285
- ▶ 8.8, “University Scientific Data Archive” on page 286
- ▶ 8.9, “Oil and gas” on page 287
- ▶ 8.10, “S3 Object Interface” on page 288
- ▶ 8.11, “AFM use cases” on page 290

## 8.1 Overview of use cases

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 8-1.

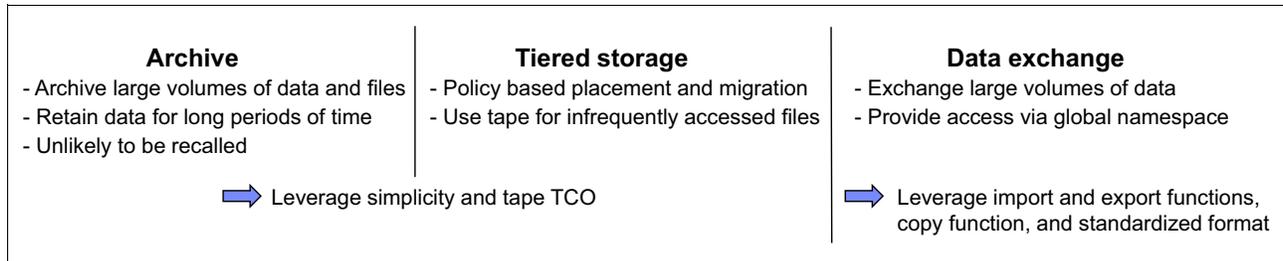


Figure 8-1 Typical use cases for IBM Spectrum Archive EE

For more information about each use case, see Figure 8-2, Figure 8-3 on page 277, and Figure 8-4 on page 278.

### 8.1.1 Use case for archive

Figure 8-2 summarizes the requirements, solution, and benefits of an IBM Spectrum Archive EE use case for archiving data.

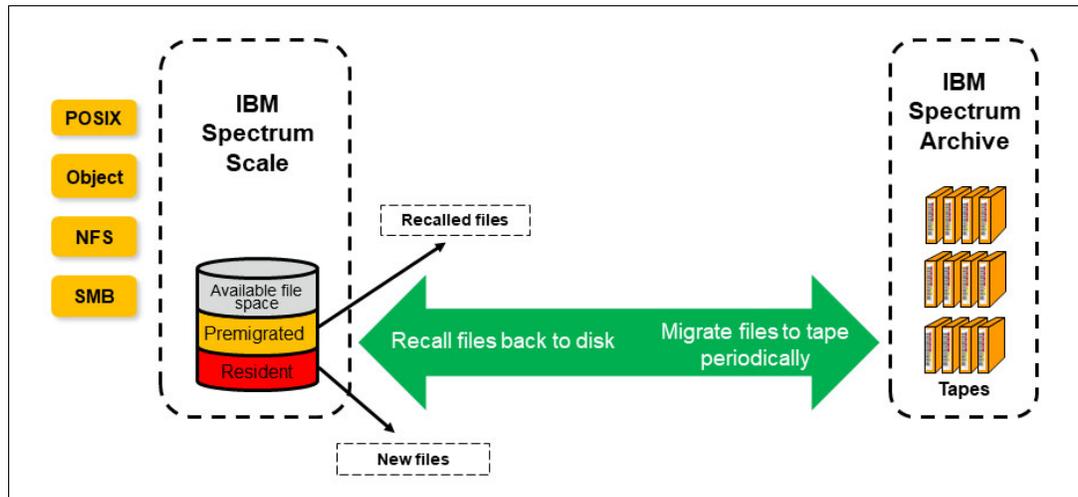


Figure 8-2 Use case for archive

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 archive storage that is 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

## 8.1.2 Use case for tiered and scalable storage

Figure 8-3 summarizes the requirements, solution, and benefits of an IBM Spectrum Archive EE use case for tiered and scalable storage.

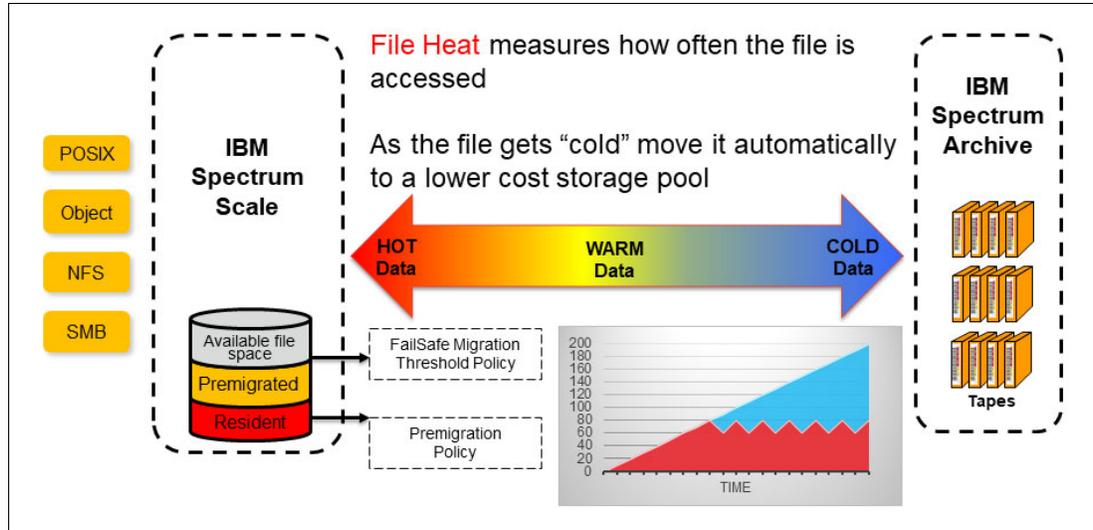


Figure 8-3 Use case for tiered and scalable storage

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 archive storage that is 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

### 8.1.3 Use case data exchange

Figure 8-4 summarizes the requirements, solution, and benefits of an IBM Spectrum Archive EE use case for data exchange.

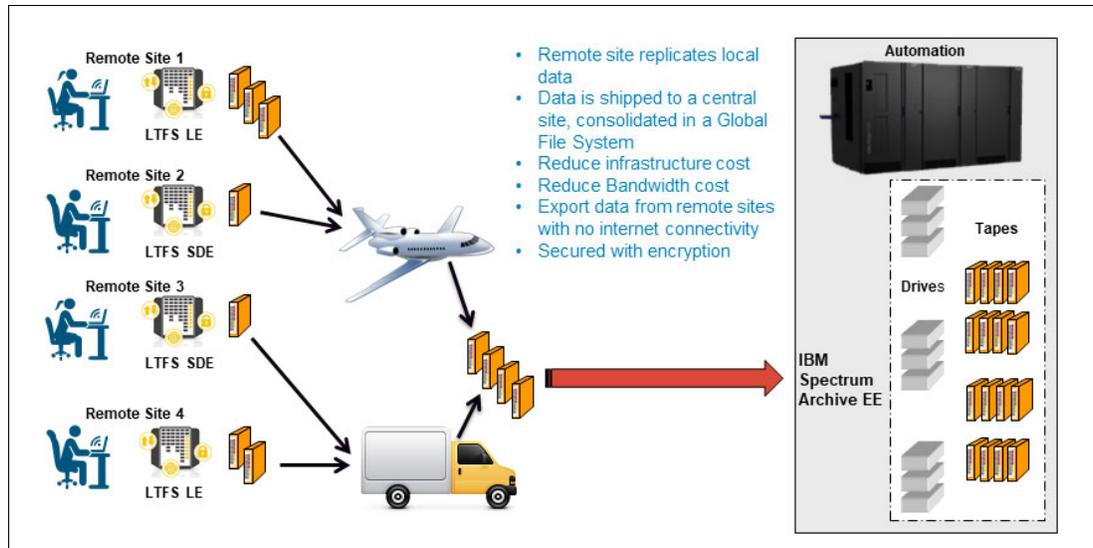


Figure 8-4 Use case for data exchange

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 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

## 8.2 Media and Entertainment

As visual effects get more sophisticated, the computational demands of media and entertainment rose sharply. In response, Pixit Media became the United Kingdom's leading provider of post-production software-defined solutions that are based on IBM Spectrum Storage technology, which provides reliable performance and cost-efficiency to support blockbuster growth.

Their business challenge is to keep audiences spellbound with their output. Media and Entertainment companies must share, store, and access huge files that traditional infrastructure solutions are failing to deliver. To meet this transformation, Pixit Media put consistent performance and chart-topping scalability in the limelight through software-defined solutions that are based on IBM Spectrum Storage technology that helps customers create hit after hit.

Figure 8-5 shows how a studio's workflow and applications can talk to IBM Spectrum Scale, which is backed by industry-standard servers and disk arrays that serve as the global single namespace file system.

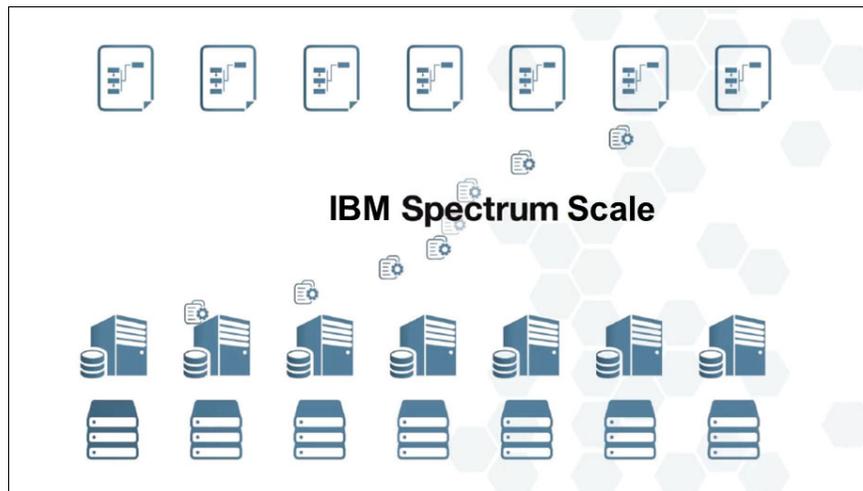


Figure 8-5 Pixit Media workflow using IBM Spectrum Scale

Pixit Media also began offering solutions that are based on IBM Spectrum Archive EE. They commented that:

“Most of our clients have complex requirements when it comes to archiving a project. They want to move it to a reliable media on a project-by-project basis but without having to manually manage the library. IBM Spectrum Archive EE can offer these customers a worry-free, centralized approach to managing this process, which can be set up in just two or three days and scale out extremely quickly.”<sup>1</sup>

<sup>1</sup> <https://www.ibm.com/case-studies/pixit-media-systems-spectrum>

## 8.3 Media and Entertainment

The customer is the nation's premier regional sports network providing a unique and special experiences for local sports fans, sponsors, teams, and media partners across 4 different regions of the United States. All production and post production videos are stored on high-speed storage. However many of these post production videos will never be accessed again so there is no need to occupy space on the high-speed storage. The customer will migrate these post production videos to tape using IBM Spectrum Archive EE.

If the post production videos are ever needed again, they can transparently be recalled back to the high-speed storage. The user data files can be viewed as a normal file system to the Media Asset Management (MAM) system while providing a seamless integration into environments. The mission is to preserve these assets and provide rapid access. Figure 8-6 shows an IBM Spectrum Archive use case for media and entertainment.

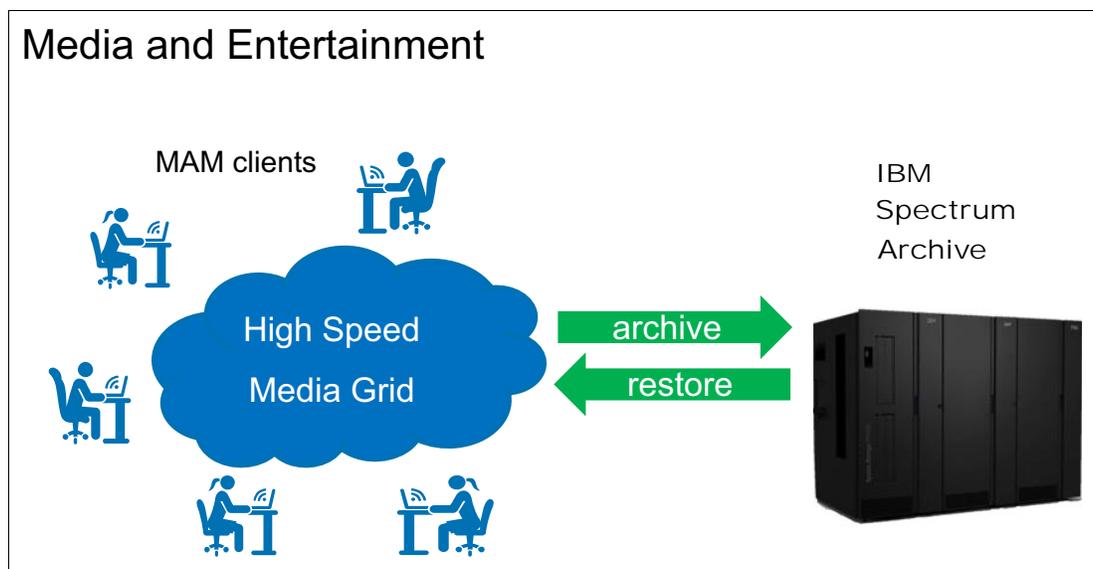


Figure 8-6 IBM Spectrum Archive use case for media and entertainment

## 8.4 High-Performance Computing

Institutions and universities have long used High-Performance Computing (HPC) resources to gather, store, and process their deep wealth of data which is often kept indefinitely. One key step within this process is the ability to easily transfer, share, and discuss the data within their own research teams and others. But this is often a major hurdle for the researchers.

Data transfers, with Globus, overcome this hurdle by providing a secure and unified interface to their research data. Globus handles the complexities behind the details of large-scale data transfers (such as tuning performance parameters), maintains security, monitors progress, and validates correctness. Globus includes a “fire and forget” model where users submit the data transfer task and are notified by Globus when the task is complete. Therefore, researchers can concentrate only on performing their research.

Every data transfer task has a source endpoint and a destination endpoint. The endpoints are described as the different locations where data can be moved to or from using the Globus transfer, sync, and sharing service. For the IBM Spectrum Scale file system and IBM Spectrum Archive EE tape usage, the following terminology is used:

- ▶ Data is *archived* when you use the IBM Spectrum Scale file system as a destination endpoint and it is cached until it is migrated to tape.
- ▶ Data is *restored* or *recalled* when you use the IBM Spectrum Scale file system as a source endpoint and might require a bulk recall of data from tape before the actual data transfers can occur.

When recalls from tape through IBM Spectrum Archive EE are required, the ability to optimize the bulk recalls is important because it can require a significant amount of time to complete. Without optimizations, the tape recalls look randomized because there is no queue to be able to group tape recalls which are on the same tape or within the same area of tape. Therefore, a lot of time will be consumed by locating or rewinding times on the tape, and on the unmounting or mounting of the target tape.

Therefore, starting with Globus Connect Server (GCS) version 5.4, Globus introduced a feature called “posix staging” to allow the files to be prestaged to a disk cache before performing the data transfer. With IBM Spectrum Scale and IBM Spectrum Archive EE, this feature allows the optimization of bulk recalls from tape, prestaging them to the IBM Spectrum Scale file system before being accessed by Globus. This process is done through Globus by calling a specific “staging app” to generate the bulk recalls inside IBM Spectrum Archive EE. A bulk recall can contain up to 64 staging requests per data transfer task.

Figure 8-7 shows Globus integration with IBM Spectrum Archive for archiving and recalling, including prestaging capabilities. Globus will likely interleave staging and transfer operations during the processing of the transfer task to improve performance.

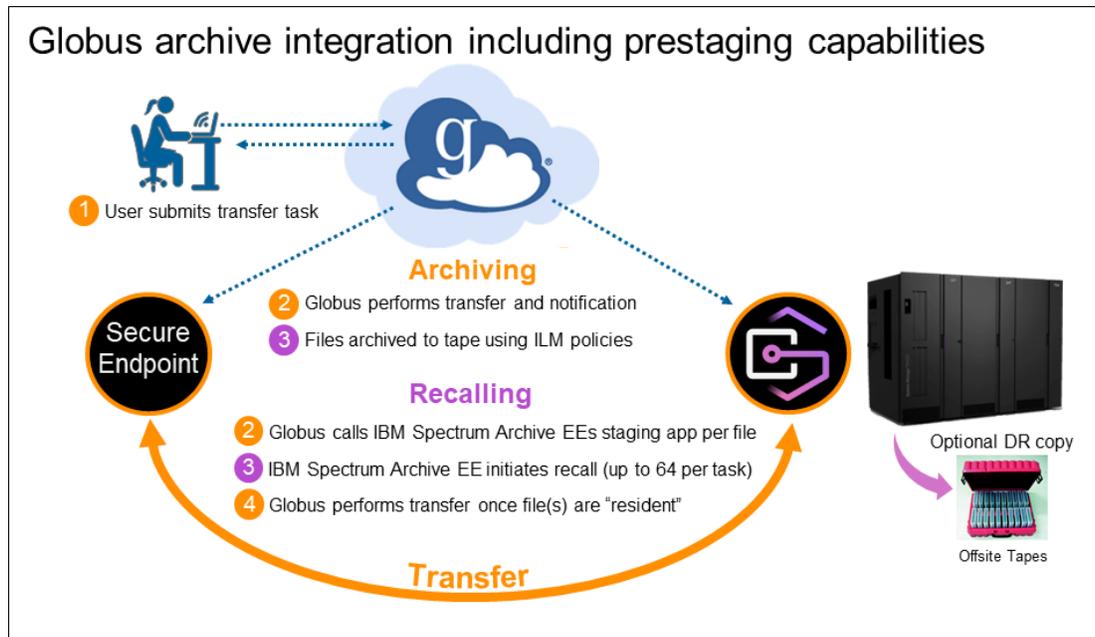


Figure 8-7 Globus archive integration with IBM Spectrum Archive EE

For more information about the Globus posix staging feature, see [this web page](#).

## 8.5 Healthcare

Amsterdam University Medical Center (UMC), location VUmc is enabling ground breaking research with scalable, cost-effective storage for big data. With its storage requirements skyrocketing that must be kept for decades, they needed to reevaluate its infrastructure to continue conducting cutting-edge research in a secure and cost-effective way. Working with a partner, the medical center helped researchers migrate from NAS drives to a centralized storage platform that is based on IBM Spectrum Storage solutions.

With IBM Spectrum Scale and IBM Spectrum Archive EE solutions at the heart of its centralized storage environment, Amsterdam UMC, location, VUmc supports its clinicians, researchers, and administrators with the resources they need to work effectively. The promise is to deliver a solution with which users can find the exact data they are looking for easily and quickly, when and where they need it, and without disruptions.

When users or researchers have a new storage request, they are presented with a menu that contains three storage options: gold, silver, and bronze. The gold option is on disk, while the silver and bronze options are on tape (see Figure 8-8). Even when the data is on tape, it is accessible from the online tape archive that provides the lower cost per TB and makes the entire solution environmentally friendly and green.

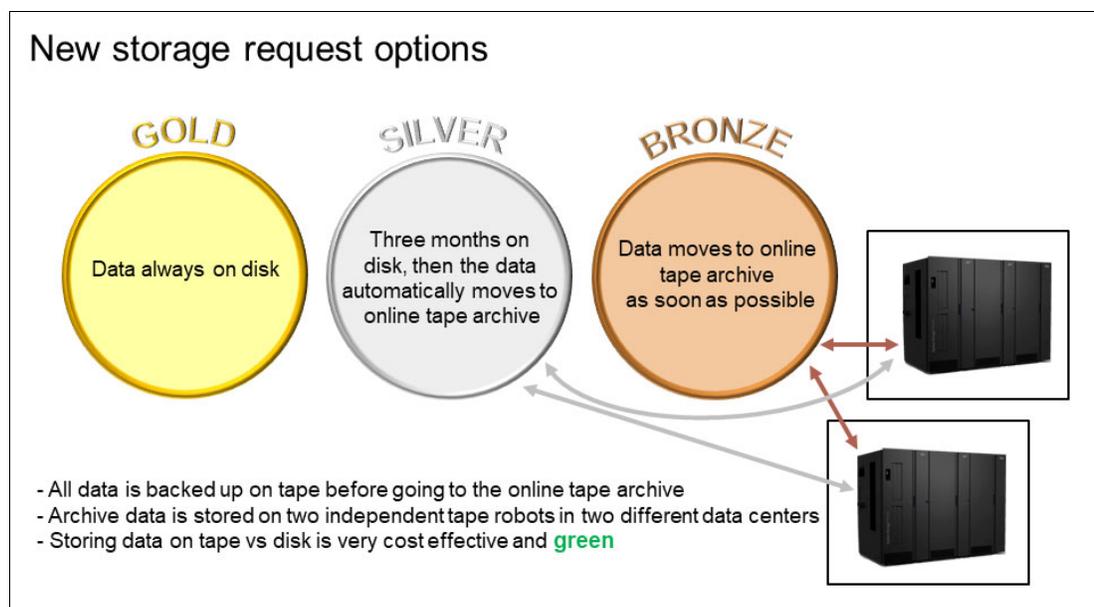


Figure 8-8 Storage request options based on policies

Now the solution helps to achieve 99% faster data migrations, which enables IT to focus on value-added development. The centralized data architecture ensures VUmc can fully support its clinicians, researchers, and administrators with the resources they need to accelerate discoveries and conduct innovative research.

For more information, see [this IBM Support web page](#).

## 8.6 Genomics

The customer is one of the largest genomics research facilities in North America, integrating sequencing, bioinformatics, data management, and genomics research. Their mission is to deliver analysis to support personalized treatment to individual cancer patients where the Standard of Care has failed. Hundreds of PBs of data has been moved from old filers to the IBM Spectrum Archive EE system, at a rate of 1.2 PB per month to tape. 2 sets of tape storage pools are used to be able to exchange/share the data with a remote education institute.

The software engineers have also optimized the recall of massive genomics data for their researchers, allowing for quick access to TBs of their migrated genomics data. Figure 8-9 shows a high-level archive for a genomics data archive.

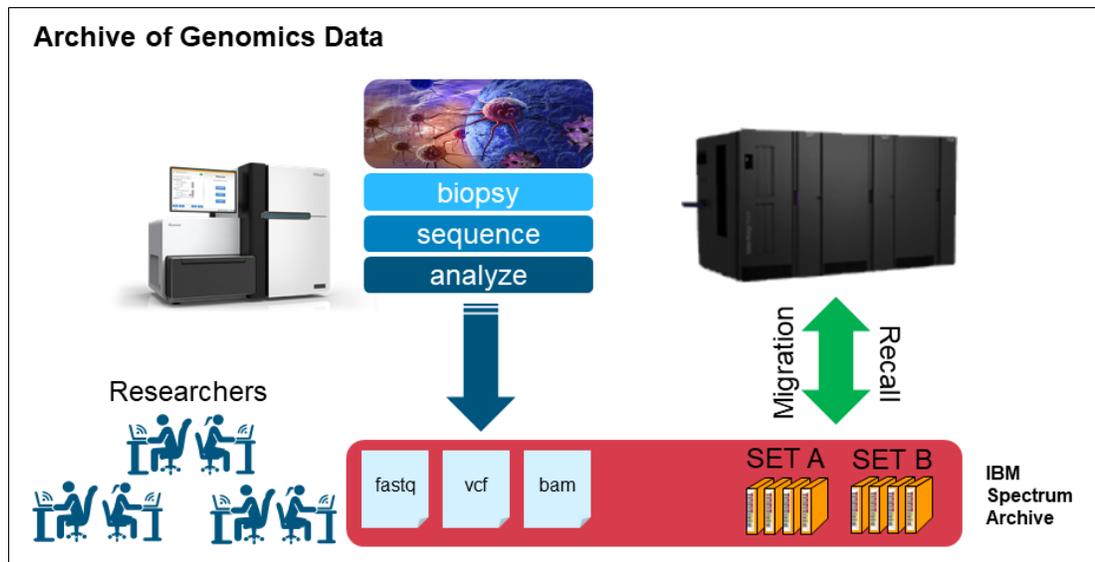


Figure 8-9 IBM Spectrum Archive use case for genomics data archive

## 8.7 Archive of research and scientific data for extended periods

This research institute routinely manages large volumes of data generated internally and collected by other institutes. To support its ongoing projects, the institute must archive and store the data for many years. However, because most of the data is infrequently accessed, the research institute was looking for a cost-efficient archiving solution that would allow transparent user access.

In addition, the research institute needed a solution that would facilitate the fast import and export of large data volumes. Figure 8-10 shows the high-level architecture to archive research and scientific data for long periods.

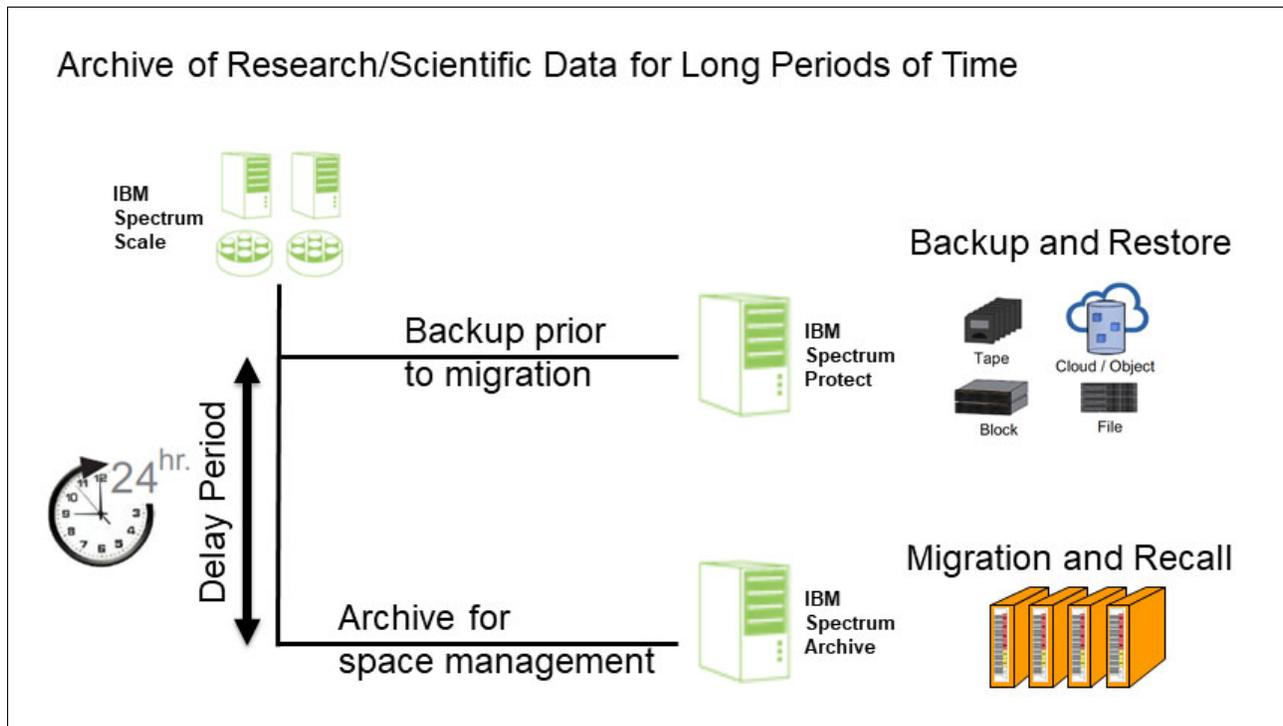


Figure 8-10 IBM Spectrum Archive use case for archiving research/scientific data for long periods of time

Figure 8-10 also shows the redundancy of the archive data from the backup solution. Both archive and backup solutions are storing data on lower cost tape storage. The copies are in two independent systems offering more options for stricter data security requirements.

## 8.8 University Scientific Data Archive

This university specializes in transportation research. This solution was designed to meet the long-term storage needs of the scientific research community, the university refers to it as the *Scientific Data Archive*. Scientific research frequently gathers data that need to be available for subsequent dissemination, follow-on research studies, compliance or review of provenance, and other purposes, sometimes with commitment to maintain these data sets for decades.

The proposed solution will be a storage system residing in two locations, providing network access to multiple organizational units within the university, each with their own respective permission models. The primary objective of the Scientific Data Archive is to provide cost-effective, resilient, long-term storage for research data and supporting research computing infrastructure. The archive will be a storage facility operated the university's Storage Management Team, in cooperation with other units on campus. This facility will deliver service to research organizations on campus. Figure 8-11 show the architecture for the university's archive.

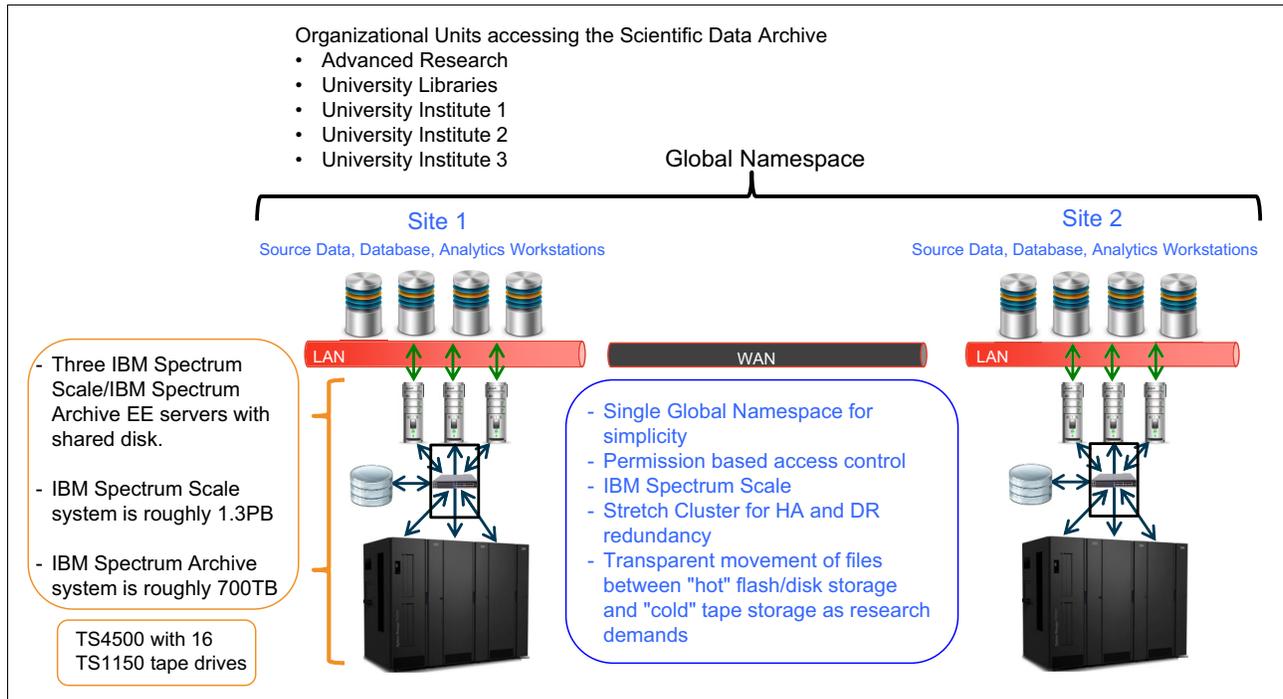


Figure 8-11 IBM Spectrum Archive use case for university Scientific Data Archive

## 8.9 Oil and gas

An oil and gas managed service provider collects offshore seismic data for analysis. The solution uses the leading technology based on IBM Power Linux, IBM Spectrum Scale, and IBM Spectrum Archive EE as a seismic data repository. The seismic data is collected from vessels. The technology for both acquisition and processing of seismic data has evolved dramatically over time. Today, a modern seismic vessel typically generates 4-5 TBs of new raw data per day, which once processed will generate 10 to 100 times more data in different formats.

For data that needs to be online but not accessed very frequently, tape is by far more attractive than spinning disk. Hybrid storage solutions with automated and policy-driven movement of data between different storage tiers including tape is required for such large data repositories.

Figure 8-12 shows an IBM Spectrum Archive Oil and Gas archive use case.

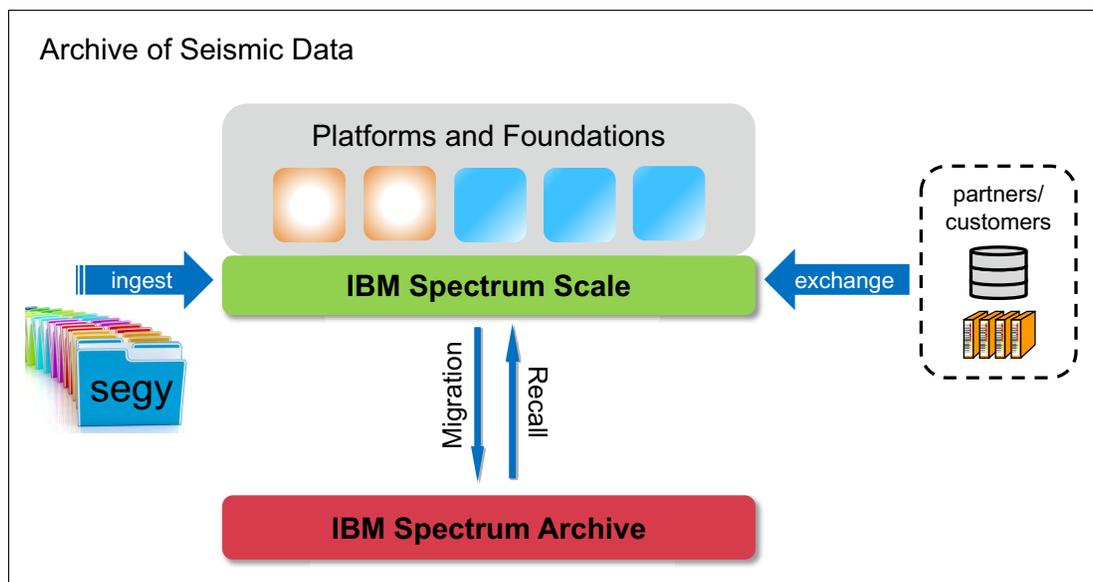


Figure 8-12 IBM Spectrum Archive use case for oil and gas

A video of every leading Managed Service Provider in the Nordics with thousands of servers using tape as an integral part of their seismic data management is available at [this website](#).

## 8.10 S3 Object Interface

For many traditional file-based workloads, tape is a popular option for its cold storage economics. But more recently, the idea of a “S3 on Tape” concept grew in popularity, where a hot or warm tier of object storage is used with a cold tier backed by tape. Going one step further, some customers even want to have a “unified data sharing” capability where the same data, file, or object is accessible across multiple protocols (such as NFS, SMB, S3) regardless of whether the file is on disk or on tape.

Therefore, as more and more companies leverage the advantage of object storage and applications, the storage of the data becomes an important role. Data must be easily accessed (regardless of protocol) and should not be tied to a specific tier. Access should be in parallel, support tape throughput, and always on access.

IBM Spectrum Scale object storage combines the benefits of IBM Spectrum Scale with OpenStack Swift, which is the most widely used open source object store. This object storage system uses a distributed architecture with no central point of control, providing greater scalability, redundancy, and the ability to access the objects via Swift API or the S3 API. But IBM Spectrum Scale object storage requires the deployment of “Protocol Nodes” within the IBM Spectrum Scale cluster.

For more information, see this [IBM Documentation web page](#).

Other than the IBM Spectrum Scale object storage, there is the option of using MinIO, which provides the “S3 on Tape” concept. MinIO is a high-performance, open source, S3 compatible, enterprise hardened object storage. In NAS Gateway mode, MinIO provides the S3-compatible environment that serves as the object storage endpoint. MinIO leverages the IBM Spectrum Scale file system including multiple instances of MinIO NAS Gateways as a distributed object storage.

The MinIO NAS Gateway is a simple translator for all S3 API calls, and writes the files on the IBM Spectrum Scale file system as normal files. This MinIO NAS Gateway provides an important feature called global 1-to-1 data sharing, which means that every object is a single file on IBM Spectrum Scale (one object to one file). Any S3 object can be seen as a file through IBM Spectrum Scale, including SMB/NFS, and any file created through IBM Spectrum Scale, including SMB/NFS, can be seen as an object using the S3 API. You can also create a separate NAS Gateway directory for only S3 data.

When you use MinIO with IBM Spectrum Scale and IBM Spectrum Archive EE, MinIO handles all the S3 API calls while IBM Spectrum Scale provides the lifecycle management of the objects to either remain on disk for hot or warm tier or on tape for the cold tier through the powerful policy engine (see Figure 8-13).

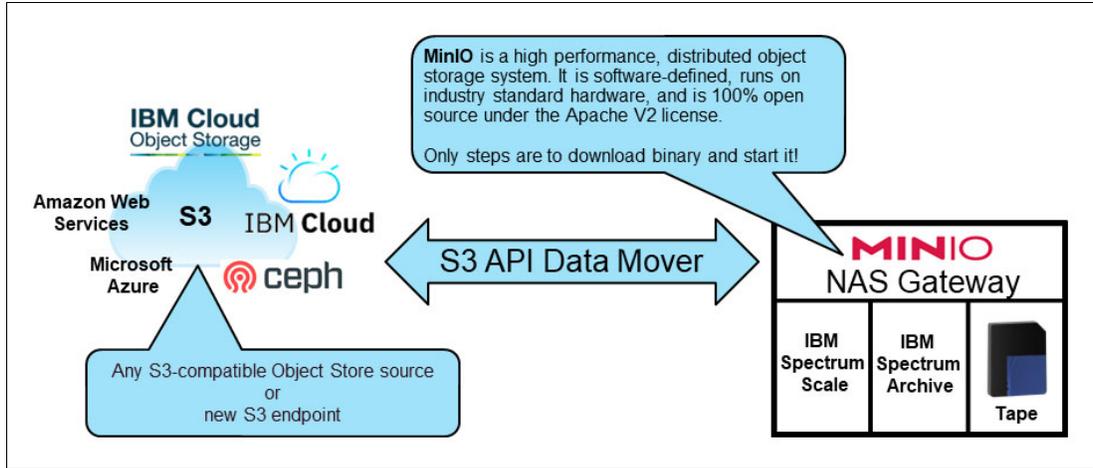


Figure 8-13 High level MinIO with IBM Spectrum Archive architecture

To the S3 object user, all objects appear as they are in the object storage, available on disk, and accessible using standard S3 GET or HEAD methods. Thus, it is simple to use it as any standard S3 object storage, but it offers the advantage of tape economics for older objects which have not been accessed for some time.

## 8.11 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).

### 8.11.1 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 create a centralized archive repository. This 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 8-14 shows a configuration of a single home cluster with multiple cache clusters to form a centralized archive repository.

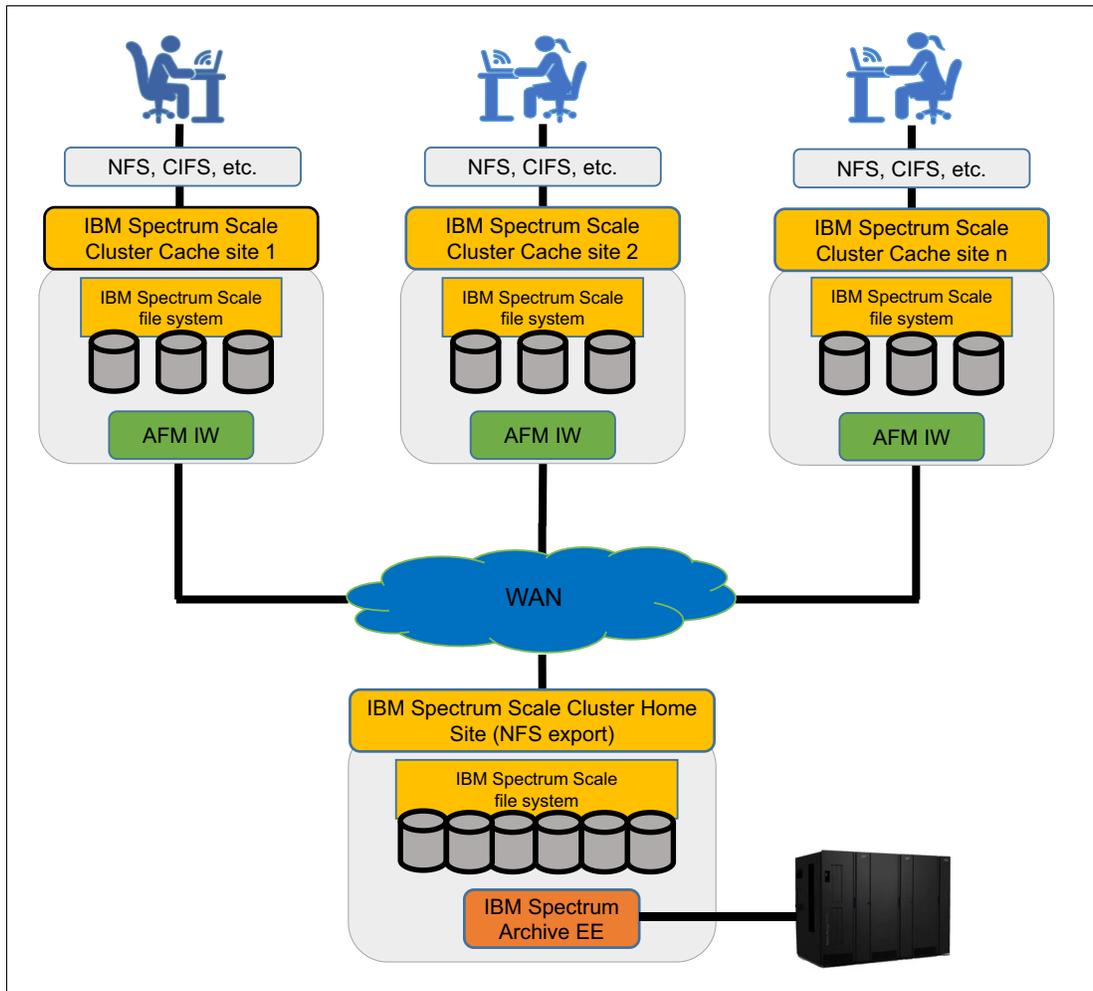


Figure 8-14 Centralized archive repository

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.

### 8.11.2 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 8-15 shows a configuration of an asynchronous archive replication solution between a home and cache site.

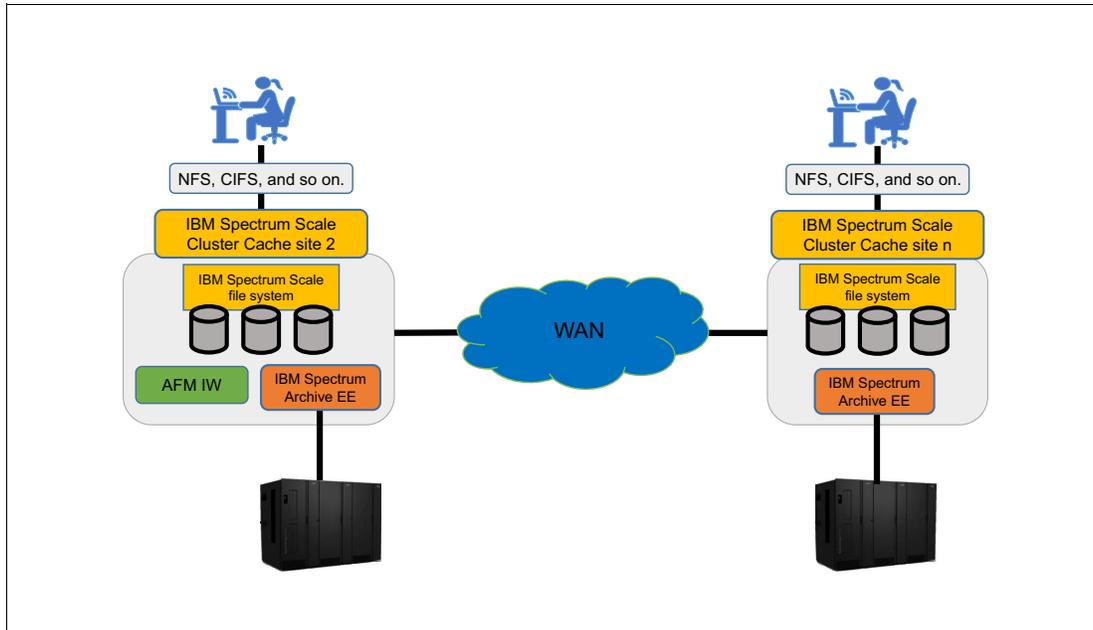


Figure 8-15 Asynchronous archive replication



# Troubleshooting IBM Spectrum Archive Enterprise Edition

This chapter describes the process that you can use to troubleshoot issues with IBM Spectrum Archive Enterprise Edition (IBM Spectrum Archive EE).

This chapter includes the following topics:

- ▶ 9.1, “Overview” on page 294
- ▶ 9.2, “Hardware” on page 297
- ▶ 9.3, “Recovering data from a write failure tape” on page 301
- ▶ 9.4, “Recovering data from a read failure tape” on page 302
- ▶ 9.5, “Software” on page 303
- ▶ 9.6, “Recovering from system failures” on page 312

## 9.1 Overview

This section provides a simple health check procedure for IBM Spectrum Archive EE.

### 9.1.1 Quick health check

If you are having issues with an IBM Spectrum Archive EE environment, Figure 9-1 shows a simple flowchart that you can follow as the first step to troubleshooting problems with the IBM Spectrum Archive EE components.

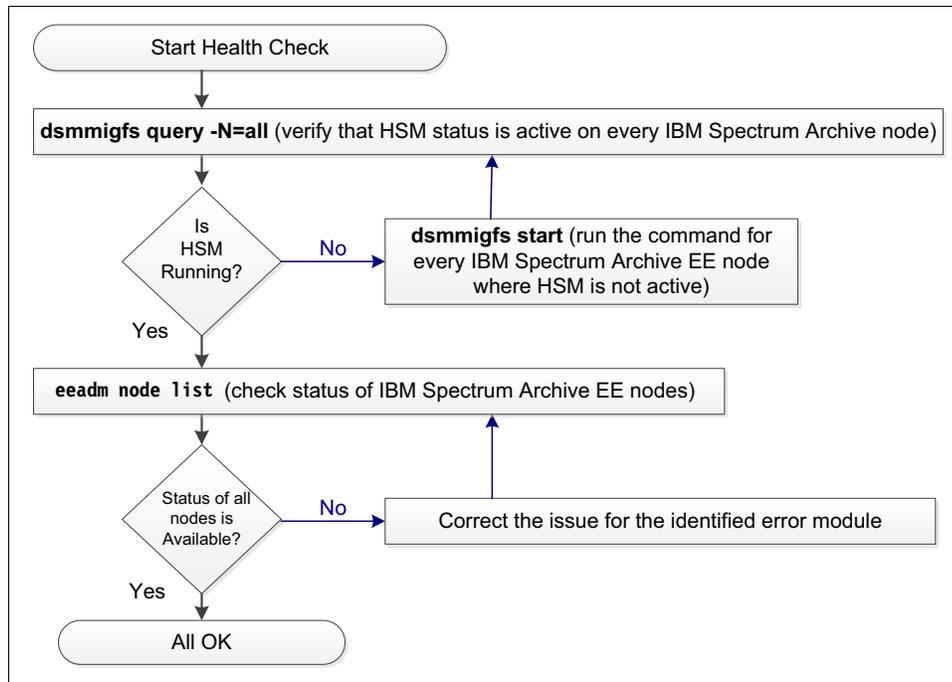


Figure 9-1 Quick health check procedure

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.

### 9.1.2 Common startup errors

IBM Spectrum Archive EE has multiple components it manages and requires to have a successful start before the system is ready for use. In addition for its own components multiple external components must be running and configured properly for IBM Spectrum Archive EE to have a proper startup. This section will walk through some of these components and the common startup errors that effect IBM Spectrum Archive EE from starting up correctly.

After running `eeadm cluster start` and the returned status code is error, IBM Spectrum Archive EE failed to start correctly and user actions are required to remedy the situation. To view the type of error that has occurred during startup and which nodes are effected run the `eeadm node list` command

## Failed startup caused by rpcbind

The `rpcbind` utility is required for MMM to function correctly, if `rpcbind` is not running IBM Spectrum Archive EE will start up with errors and will be unusable until the issue has been resolved. In most cases this is caused when `rpcbind` is not started up. If the server had been recently powered down for maintenance and was started back up `rpcbind` might not start up automatically.

Example 9-1 shows the output of a failed IBM Spectrum Archive EE startup due to `rpcbind` not running.

### *Example 9-1 rpcbind caused startup failure*

---

```
[root@tora ~]# eeadm cluster start
Library name: lib_tora, library serial: 0000013FA002040C, control node (ltfsee_md)
IP address: 9.11.244.63.
Starting - sending a startup request to lib_tora.
Starting - waiting for startup completion : lib_tora.
Starting - opening a communication channel : lib_tora.
.
Starting - waiting for getting ready to operate : lib_tora.
.....
2019-01-21 09:04:17 GLESL657E: Fail to start the IBM Spectrum Archive EE service
(MMM) for library lib_tora.
 Use the "eeadm node list" command to see the error modules.
 The monitor daemon will start the recovery sequence.
```

```
[root@tora ~]# eeadm node list
```

```
Spectrum Archive EE service (MMM) for library lib_tora fails to start or is not
running on tora.tuc.stglabs.ibm.com Node ID:1
```

```
Problem Detected:
Node ID Error Modules
 1 MMM; rpcbind;
```

---

To remedy this issue use the `systemctl start rpcbind` command to start the process up, and either wait for the IBM Spectrum Archive EE Monitor daemon to start up MMM or issue a `eeadm cluster stop` and `eeadm cluster start` to get the MMM started. After `rpcbind` is started, it can be verified by running the `systemctl status rpcbind` to see if it is running.

Example 9-2 shows how to startup `rpcbind` and waiting for the Monitor daemon to restart MMM.

### *Example 9-2 Starting rpcbind to remedy MMM startup failure*

---

```
[root@tora ~]# systemctl start rpcbind

[root@tora ~]# systemctl status rpcbind
? rpcbind.service - RPC bind service
 Loaded: loaded (/usr/lib/systemd/system/rpcbind.service; disabled; vendor
 preset: enabled)
 Active: active (running) since Mon 2019-01-21 09:08:37 MST; 1min 12s ago
 Process: 23628 ExecStart=/sbin/rpcbind -w $RPCBIND_ARGS (code=exited,
 status=0/SUCCESS)
 Main PID: 23629 (rpcbind)
 Tasks: 1
```

```
Memory: 572.0K
CGroup: /system.slice/rpcbind.service
+-23629 /sbin/rpcbind -w
```

```
Jan 21 09:08:37 tora.tuc.stglabs.ibm.com systemd[1]: Starting RPC bind service...
Jan 21 09:08:37 tora.tuc.stglabs.ibm.com systemd[1]: Started RPC bind service.
[root@tora ~]# eeadm node list
Node ID State Node IP Drives Ctrl Node Library Node Group Host
Name
1 available 9.11.244.63 0 yes(active) lib_tora G0
tora.tuc.stglabs.ibm.com
```

---

## Failed startup caused by LE

LE is another crucial component to IBM Spectrum Archive EE. If it is not started up correctly, MMM does not start. There are two common startup problems that can be remedied quickly. The first problem is caused by no drive visibility by the Spectrum Archive EE node. This problem can be fixed by connecting Fibre Channel cables to the node and verified by running the `ltfs -o device_list` command.

Example 9-3 shows the output of `ltfs -o device_list` with no drives connected.

### Example 9-3 No drives connected to server

---

```
[root@tora ~]# ltfs -o device_list
6b6c LTFS14000I LTFS starting, LTFS version 2.4.1.0 (10219), log level 2.
6b6c LTFS14058I LTFS Format Specification version 2.4.0.
6b6c LTFS14104I Launched by "/opt/IBM/ltfs/bin/ltfs -o device_list".
6b6c LTFS14105I This binary is built for Linux (x86_64).
6b6c LTFS14106I GCC version is 4.8.3 20140911 (Red Hat 4.8.3-9).
6b6c LTFS17087I Kernel version: Linux version 3.10.0-862.14.4.el7.x86_64
(mockbuild@x86-040.build.eng.bos.redhat.com) (gcc version 4.8.5 20150623 (Red Hat
4.8.5-28) (GCC)) #1 SMP Fri Sep 21 09:07:21 UTC 2018 i386.
6b6c LTFS17089I Distribution: NAME="Red Hat Enterprise Linux Server".
6b6c LTFS17089I Distribution: Red Hat Enterprise Linux Server release 7.5 (Maipo).
6b6c LTFS17089I Distribution: Red Hat Enterprise Linux Server release 7.5 (Maipo).
6b6c LTFS17085I Plugin: Loading "sg" changer backend.
6b6c LTFS17085I Plugin: Loading "sg" tape backend.
Changer Device list:.
Tape Device list:.
```

---

The second most common LE error is caused after connecting the drives to the IBM Spectrum Archive EE node and forgetting to set at least one of the drives a control path drive from the library GUI. IBM Spectrum Archive EE requires at least one control path drive so it can communicate with the library.

Example 9-4 shows the output of a failed IBM Spectrum Archive EE start up that is caused by LE.

### Example 9-4 LE failed MMM startup

---

```
[root@tora ~]# eeadm cluster start
Library name: lib_tora, library serial: 0000013FA002040C, control node (ltfsee_md)
IP address: 9.11.244.63.
Starting - sending a startup request to lib_tora.
Starting - waiting for startup completion : lib_tora.
Starting - opening a communication channel : lib_tora.
```

```

.
Starting - waiting for getting ready to operate : lib_tora.
...
2019-01-21 09:24:33 GLESL657E: Fail to start the IBM Spectrum Archive EE service
(MMM) for library lib_tora.
 Use the "eadm node list" command to see the error modules.
 The monitor daemon will start the recovery sequence.
[root@tora ~]# eadm node list

Spectrum Archive EE service (MMM) for library lib_tora fails to start or is not
running on tora.tuc.stglabs.ibm.com Node ID:1

Problem Detected:
Node ID Error Modules
 1 LE; MMM;

```

---

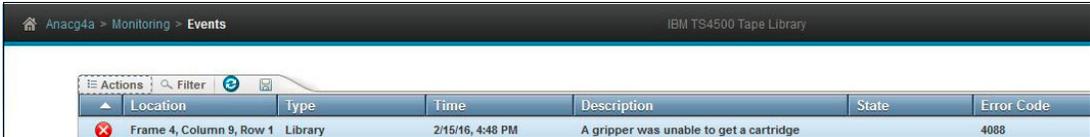
To remedy this failure, ensure that the node can see its drives and at least one drive is a control path drive.

## 9.2 Hardware

This section provides information that can help you to identify and resolve problems with the hardware that is used by IBM Spectrum Archive EE.

### 9.2.1 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 9-2 shows an example of a library error.



| 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 9-2 Tape library error log

For more information about how to solve tape library errors, see the *IBM TS4500 R8 Tape Library Guide*, [SG24-8235](#).

### 9.2.2 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 9-3 shows an example from the web interface of a tape drive that has an error and is no longer responding.

| 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 9-3 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, or call IBM H/W service. For more information about analyzing the operating system error logs, see 9.5.1, “Linux” on page 303.

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 299.

## 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 9-5 shows the format of the dump files.

### Example 9-5 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/ltfsle/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 10.3, “System calls and IBM tools” on page 327.

**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 **eeadm drive list** command. Example 9-6 shows the tape drives in use by IBM Spectrum Archive EE.

*Example 9-6 Identify the tape drive to remove*

---

```
[root@tora ~]# eeadm drive list
```

| Drive S/N    | State       | Type   | Role | Library  | Node ID | Tape | Node Group | Task ID |
|--------------|-------------|--------|------|----------|---------|------|------------|---------|
| 000000014A00 | not_mounted | TS1160 | mrg  | lib_tora | 1       | -    | G0         | -       |
| 0000078PG20C | not_mounted | TS1160 | mrg  | lib_tora | 1       | -    | G0         | -       |

---

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 **eadm drive unassign <drive serial number>** command. Example 9-7 shows the removal of a single tape drive from IBM Spectrum Archive EE.

*Example 9-7 Remove the tape drive*

---

```
[root@tora ~]# eadm drive unassign 000000014A00
2019-01-22 13:26:45 GLESL700I: Task drive_unassign was created successfully,
task id is 6432.
2019-01-22 13:26:51 GLESL121I: Drive serial 000000014A00 is removed from the
tape drive list.
2019-01-22 13:36:49 GLESL121I: Drive serial 000000014A00 is removed from the
tape drive list.
```

---

3. Check the success of the removal. Run the **eadm drive list** command and verify that the output shows that the MMM attribute for the drive is in the stock state. Example 9-8 shows the status of the drives after it is removed from IBM Spectrum Archive EE.

*Example 9-8 Check the tape drive status*

---

```
[root@tora ~]# eadm drive list
```

| Drive S/N    | State       | Type   | Role | Library  | Node ID | Tape | Node Group | Task ID |
|--------------|-------------|--------|------|----------|---------|------|------------|---------|
| 0000078PG20C | not_mounted | TS1160 | mrg  | lib_tora | 1       | -    | G0         | -       |
| 000000014A00 | unassigned  | -      | ---  | lib_tora | -       | -    | -          | -       |

---

4. Identify the primary device number of the drive for subsequent operations by running the **/opt/ibm/ltfsle/bin/ltfs -o device\_list** command. The command outputs a list of available drives. Example 9-9 shows the output of this command.

*Example 9-9 ltfs -o device\_list command output*

---

```
[root@tora ~]# /opt/ibm/ltfsle/bin/ltfs -o device_list
77d1 LTFS14000I LTFS starting, LTFS version 2.4.1.1 (10226), log level 2.
77d1 LTFS14058I LTFS Format Specification version 2.4.0.
77d1 LTFS14104I Launched by "/opt/IBM/ltfs/bin/ltfs -o device_list".
77d1 LTFS14105I This binary is built for Linux (x86_64).
77d1 LTFS14106I GCC version is 4.8.3 20140911 (Red Hat 4.8.3-9).
77d1 LTFS17087I Kernel version: Linux version 3.10.0-862.14.4.el7.x86_64
(mockbuil@x86-040.build.eng.bos.redhat.com) (gcc version 4.8.5 20150623 (Red
Hat 4.8.5-28) (GCC)) #1 SMP Fri Sep 21 09:07:21 UTC 2018 i386.
77d1 LTFS17089I Distribution: NAME="Red Hat Enterprise Linux Server".
77d1 LTFS17089I Distribution: Red Hat Enterprise Linux Server release 7.5
(Maipo).
77d1 LTFS17089I Distribution: Red Hat Enterprise Linux Server release 7.5
(Maipo).
77d1 LTFS17085I Plugin: Loading "sg" changer backend.
77d1 LTFS17085I Plugin: Loading "sg" tape backend.
Changer Device list:.
Device Name = /dev/sg11, Vender ID = IBM , Product ID = 03584L22 ,
Serial Number = 0000013FA002040C, Product Name = TS3500/TS4500.
Tape Device list:.
Device Name = /dev/sg1, Vender ID = IBM , Product ID = 0359260F ,
Serial Number = 0000078PG20C, Product Name = [0359260F].
```

Device Name = /dev/sg0, Vender ID = IBM , Product ID = 0359260F ,  
Serial Number = 000000014A00, Product Name =[0359260F].

---

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/sgnumber, 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/sgnumber, 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 **eeadm drive assign drive\_serial -n node\_id** command, where **node\_id** is the same node that the drive was assigned to originally in step 1 on page 299.

Example 9-10 shows the tape drive that is readded to IBM Spectrum Archive EE.

*Example 9-10 Add again the tape drive*

---

```
[root@tora ~]# eeadm drive assign 000000014A00 -n 1
2019-01-23 08:37:24 GLESL119I: Drive 000000014A00 assigned successfully.
```

---

Running the **eeadm drive list** command again shows that the tape drive is no longer in a “unassigned” state. Example 9-11 shows the output of this command.

*Example 9-11 Check the tape drive status*

---

```
[root@tora ~]# eeadm drive list
Drive S/N State Type Role Library Node ID Tape Node Group Task ID
000000014A00 not_mounted TS1160 mrg lib_tora 1 - G0 -
0000078PG20C not_mounted TS1160 mrg lib_tora 1 - G0 -
```

---

## 9.3 Recovering data from a write failure tape

A tape that suffers from a write failure goes into a **require\_replace** state. Complete the following steps to recover data from a **require\_replace** tape:

1. Verify that the tape is in the **require\_replace** state by running **eeadm tape list**.
2. Verify there is a tape cartridge in the appendable state within the same cartridge pool which has enough available space to hold the data from the **require\_replace** tape.
3. Run the **eeadm tape replace** command on the **require\_replace** tape to start the data transfer process onto an appendable tape within the same pool and unassigning it from the tape cartridge pool at the end.
4. Run the **eeadm tape move** command on the **require\_replace** tape to move the I/O station in order to dispose of the bad tape.

Example 9-12 shows the commands and output to replace a **require\_replace** tape.

*Example 9-12 Replacing data from a require\_replace tape*

---

```
[root@ginza prod]# eeadm tape list | grep pool1
FC0254L8 degraded require_replace 10907 0 0
0% pool1 liba homeslot -
FC0257L8 ok appendable 10907 0 10907
0% pool1 liba homeslot -

[root@ginza prod]# eeadm tape replace FC0254L8 -p pool1
2019-02-05 14:49:03 GLESL700I: Task tape_replace was created successfully, task id
is 2319.
2019-02-05 14:49:03 GLESL755I: Start a reconcile before starting a replace against
1 tapes.
2019-02-05 14:51:46 GLESS002I: Reconciling tape FC0254L8 complete.
2019-02-05 14:51:49 GLESL756I: Reconcile before replace finished.
2019-02-05 14:51:49 GLESL753I: Starting tape replace for FC0254L8.
2019-02-05 14:51:49 GLESL754I: Found a target tape for tape replace (FC0257L8).
2019-02-05 14:55:03 GLESL749I: The tape replace operation for FC0254L8 is
successful.

[root@ginza prod]# eeadm tape list | grep pool1
FC0257L8 ok appendable 10907 0 10906
0% pool1 liba drive -
```

---

In the rare case where the error on tape is permanent (several **eeadm tape replace** commands have failed), it is suggested to try the **eeadm tape unassigned --safe-remove** command instead. The **--safe remove** option recalls all the active files on the tape back to an IBM Spectrum Scale file system that has adequate free space. The files then need to be manually migrated again to a good tape again. Specify only one tape for use with the **--safe-remove** option.

## 9.4 Recovering data from a read failure tape

A tape that suffers from a read failure goes into a **need\_replace** state.

Complete the following steps to copy migrated jobs from a **need\_replace** tape to a valid tape within the same pool:

1. Identify a tape with a read failure by running **eeadm tape list** to locate the **need\_replace** tape.
2. Verify there is an appendable tape within the same pool which has enough available space to hold the data from the **need\_replace** tape cartridge.
3. Run the **eeadm tape replace** command to start the data transfer process onto an appendable tape within the same pool and unassigning it from the tape cartridge pool at the end.
4. Run **eeadm tape move** command on the **need\_replace** tape to move the I/O station in order to dispose of the bad tape.

Example 9-13 shows system output of the steps to recover data from a read failure tape.

*Example 9-13 Recovering data from a read failure*

---

```
[root@hakone prod3]# eeadm tape list | grep test1
IM1229L6 ok appendable 2242 0 2242
0% test1 liba homeslot -
IM1195L6 info need_replace 2242 0 0
0% test1 liba drive -

[root@hakone prod3]# eeadm tape replace IM1195L6 -p test1
2019-02-26 14:29:27 GLESL700I: Task tape_replace was created successfully, task id
is 1297.
2019-02-26 14:29:27 GLESL755I: Start a reconcile before starting a replace against
1 tapes.
2019-02-26 14:30:05 GLESS002I: Reconciling tape IM1195L6 complete.
2019-02-26 14:30:06 GLESL756I: Reconcile before replace finished.
2019-02-26 14:30:06 GLESL753I: Starting tape replace for IM1195L6.
2019-02-26 14:30:06 GLESL754I: Found a target tape for tape replace (IM1229L6).
2019-02-26 14:31:23 GLESL749I: The tape replace operation for IM1195L6 is
successful.
```

---

In the seldom case where the error on tape is permanent such that several **eeadm tape replace** commands have failed, it is suggested to try the **eeadm tape unassign --safe-remove** command instead. The **--safe remove** option recalls all the active files on the tape back to an IBM Spectrum Scale file system that has adequate free space, and those files need to be manually migrated again to a good tape again. Specify only one tape for the use with **--safe-remove** option.

## 9.5 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.

### 9.5.1 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 9-14 shows the changes to the `rsyslog.conf` file and the location of the two log files.

*Example 9-14 The `rsyslog.conf` file*

---

```
[root@l1fssn1 ~]# cat /etc/rsyslog.conf | grep ltfs
:msg, startswith, "GLES," /var/log/l1fsee_trc.log;gles_trc_template
:msg, startswith, "GLES" /var/log/l1fsee.log;RSYSLOG_FileFormat
:msg, regex, "LTFS[ID0-9][0-9]*[EWID]" /var/log/l1tfs.log;RSYSLOG_FileFormat
```

---

By default, after the `l1tfs.log`, `l1fsee.log`, and `l1fsee_trc.log` files reach the threshold size, they are rotated and four copies are kept. Example 9-15 shows the log file rotation settings. These settings can be adjusted as needed within the `/etc/logrotate.d/ibmsa-logrotate` control file.

*Example 9-15 Syslog rotation*

---

```
[root@l1fssn1 ~]# cat /etc/logrotate.d/ibmsa-logrotate
/var/log/l1fsee.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/l1fsee_trc.log {
 size 10M
 rotate 9
 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/l1fsee_mon.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 9-16.

*Example 9-16 Simple migration with informational messages*

---

```
[root@ginza prod]# eeadm migrate mig -p pool2
2019-02-07 18:15:07 GLESL700I: Task migrate was created successfully, task id is
2326.
2019-02-07 18:15:08 GLESM896I: Starting the stage 1 of 3 for migration task 2326
(qualifying the state of migration candidate files).
2019-02-07 18:15:08 GLESM897I: Starting the stage 2 of 3 for migration task 2326
(copying the files to 1 pools).
2019-02-07 18:15:08 GLESM898I: Starting the stage 3 of 3 for migration task 2326
(changing the state of files on disk).
2019-02-07 18:15:08 GLESL159E: Not all migration has been successful.
2019-02-07 18:15:08 GLESL038I: Migration result: 0 succeeded, 100 failed, 0
duplicate, 0 duplicate wrong pool, 0 not found, 0 too small to qualify for
migration, 0 too early for migration.
```

---

From the GLESL159E 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 9-17 shows the end of the `ltfsee.log` file immediately after the **failed migrate** command is run.

*Example 9-17 The ltfsee.log file*

---

```
tail /var/log/ltfsee.log
2019-02-07T18:20:03.301978-07:00 ginza mmm[14807]: GLESM600E(00412): Failed to
migrate/premigrate file /ibm/gpfs/prod/FILE1. The specified pool name does not
match the existing replica copy.
2019-02-07T18:20:03.493810-07:00 ginza mmm[14807]: GLESL159E(00144): 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 pool.

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 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.

## 9.5.2 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 9-18 shows normal operational messages that appear in the IBM Spectrum Scale log file.

*Example 9-18 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 9-19 shows the output of the error when trying to mount GPFS.

*Example 9-19 Superblock error from mounting GPFS*

```

[root@l1fsm11 ~]# 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 9-20 shows output of starting HSM, checking the status, and mounting GPFS.

*Example 9-20 Starting HSM and mounting GPFS*

```

[root@l1fsm11 ~]# systemctl start hsm
[root@l1fsm11 ~]# 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 l1fsm11.tuc.stglabs.ibm.com systemd[1]: Started HSM Service.
May 24 13:04:59 l1fsm11.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 dsmrecalld
root 9773 9748 0 12:57 ? 00:00:00 dsmrecalld
root 9774 9748 0 12:57 ? 00:00:00 dsmrecalld
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 dsmrecalld daemons are running by issuing **ps -afe | grep dsm**. If no dsmrecalld daemons are running, start them by issuing **dsmmigfs start**. After they have been started, GPFS can be successfully mounted.

### 9.5.3 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 9.5.1, “Linux” on page 303.

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 Documentation](#).

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 9-21 shows a problem analysis procedure for LTFS.

*Example 9-21 LTFS messages*

---

```
cat /var/log/ltfs.log
 2019-02-07T18:33:04.663564-07:00 ginza ltfs[12251]: 478d LTFS14787I Formatting
cartridge FC0252L8.
2019-02-07T18:33:42.724406-07:00 ginza ltfs[12251]: 478d LTFS14837I Formatting
cartridge FC0252L8 (0x5e, 0x00).
2019-02-07T18:33:42.729350-07:00 ginza ltfs[12251]: 478d LTFS14789E Failed to
format cartridge FC0252L8 (-1079).
2019-02-07T18:33:42.729543-07:00 ginza ltfs[12251]: 478d LTFSI1079E The operation
is not allowed.
```

---

The set of 10 characters represents the message ID, and the text that follows describes the operational state of LTFS. The fourth message ID (LTFSI1079E) 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 LTFSI1079E, is shown in the following example: Example 9-22 on page 309.

*Example 9-22 Example of message*

---

```
LTFS14789E Failed to format cartridge FC0252L8 (-1079).
LTFSI1079E The operation is not allowed.
The previous operation did not run due to tape and or drive issues.
In this case the drive and tape compatibility was incorrect.
```

---

Based on the description that is provided here, the tape cartridge in the library failed to format. Upon further investigation the tape cartridge and drive is incompatible. The required user action to solve the problem is to attach a compatible drive to the library and Spectrum Archive EE node and rerun the operation.

## 9.5.4 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 9-23 shows an example of this file.

*Example 9-23 The dsmerror.log file*

---

```
[root@ltfs97 /]# 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
```

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.

## 9.5.5 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 missing 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)
d. Task information files under the following subdirectory <GPFS mount
 point>/ltfsee/statesave.
```

If you want to collect all the information, enter `y`.

If you want to collect only a and b, enter `p` (partial).

If you agree to collect only (4) task information files, input `t`.

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 [IBM Documentation](#).

## 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.

## File status

Table 9-1 lists the seven possible status codes for files in IBM Spectrum Archive EE. They can be viewed for individual files by running the **eadm file state** command.

Table 9-1 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 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.                                                                                                                                                                                              |
| missing     | The missing status indicates that a file has the migrated or premigrated 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 missing because of tape corruption or if the tape cartridge was removed from the system without exporting. |

Files with the missing status are caused because of tape corruption or if the tape cartridge was removed from the system without exporting. If the cause is a corrupted index run the **eadm tape validate** command, otherwise if the tape is missing from the tape library bring the tape back into the library and run the **eadm library rescan** command.

## 9.6 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.

### 9.6.1 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 6.2.2, “IBM Spectrum Archive Library Edition component” on page 137.
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 `eeadm tape list` 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 6.18, “Checking and repairing tapes” on page 201.
4. Recover the inconsistent tape cartridge by running the `eeadm tape validate` command. For more information, see 6.18, “Checking and repairing tapes” on page 201.

### 9.6.2 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. Identify the issue on the tape library.
2. Manually repair the gating part.
3. Run the `eeadm tape validate` command on each effected tape.
4. Follow the procedure that is described in 9.6.1, “Power failure” on page 312.

**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 “Unassigned state” on page 322.

### 9.6.3 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 **eeadm tape list** command. For more information, see 6.23, “Obtaining system resources, and tasks information” on page 214.

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 **eeadm tape list** command.
5. Add the tape cartridge to the LTFS inventory by running the **eeadm tape assign** command.

#### 9.6.4 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:

```
ltfs -o release_device -o chanager_devname=[device_name]
```





# Reference

This chapter describes the commands that are used to operate IBM Spectrum Archive Enterprise Edition (EE), data, and metadata formats for IBM Spectrum Scale to IBM Spectrum Archive EE migrations, system calls, tools, and limitations.

This chapter includes the following topics:

- ▶ 10.1, “Command-line reference” on page 316
- ▶ 10.2, “Formats for IBM Spectrum Scale to IBM Spectrum Archive EE migration” on page 325
- ▶ 10.3, “System calls and IBM tools” on page 327
- ▶ 10.4, “IBM Spectrum Archive EE interoperability with IBM Spectrum Archive products” on page 330

## 10.1 Command-line reference

This section describes the IBM Spectrum Archive EE commands, IBM Spectrum Scale commands, and IBM Spectrum Protect space management commands.

### 10.1.1 IBM Spectrum Archive EE help guide for commands

This section describes the syntax, parameters, and function of the IBM Spectrum Archive EE commands by way of the help guide. By following the help guide, any IBM Spectrum Archive EE command can be run by reviewing the description and syntax of the command and using the provided examples.

#### The **eeadm** command

All IBM Spectrum Archive EE commands start with the **eeadm** command and all commands feature the following syntax:

```
eeadm <resource type> <action> [OPTIONS]
eeadm <subcommand> [OPTIONS]
```

The following resource types are available:

- ▶ **cluster**: Manages the IBM Spectrum Archive EE cluster
- ▶ **drive**: Manages the tape drives
- ▶ **file**: Manages the files on IBM Spectrum Scale
- ▶ **library**: Manages the tape libraries
- ▶ **node**: Manages the IBM Spectrum Archive EE servers
- ▶ **nodegroup**: Manages the node groups
- ▶ **pool**: Manages the tape storage pools
- ▶ **tape**: Manages the tape cartridges
- ▶ **task**: Manages the internal tasks

The following subcommands are available:

- ▶ **migrate**: Migrate files to the tapes and reclaim the disk space allocated to the file.
- ▶ **pre migrate**: Migrate files to the tapes but do not reclaim the disk space.
- ▶ **recall**: Recall files from the tapes to the disk storage.
- ▶ **save**: Saves the name of empty files, empty directories, and symbolic links to the tapes.

The **eeadm** command is initially used to get the high-level view of all available resource types or subcommands that are available. Then, following the remaining help guide to find the desired command to run.

#### The **eeadm <resource type> --help** command

The **eeadm <resource type> --help** command displays the available list of actions for each resource type.

For example, the **eeadm cluster** command manages the IBM Spectrum Archive EE cluster.

To display the available actions for the **eeadm cluster** command, enter the following command:

```
eeadm cluster --help
```

In this case, the following actions for **eeadm cluster** are available:

- ▶ eeadm cluster failover
- ▶ eeadm cluster set
- ▶ eeadm cluster show
- ▶ eeadm cluster start
- ▶ eeadm cluster stop
- ▶ eeadm cluster restore

Therefore, the **eeadm <resource type> --help** command should be the next help guide procedure that is used to find the available actions of each resource type.

### The eeadm <resource type> <action> --help command

After the wanted action for the resource type is found, enter the **eeadm <resource type> <action> --help** command to display the overall syntax, description, required parameters, and optional parameters for the command. The output also includes command examples that can be used to run the wanted action.

Therefore, the **eeadm <resource type> <action> --help** command should be your final help guide procedures used to run the wanted commands.

## 10.1.2 Drive status and state codes

The **eeadm drive list** command lists all of the configured drives. Drives can be listed with one of the states listed in Table 10-1 or the drive state can be displayed as empty.

Table 10-1 Status and state codes for eeadm drive list

| Drive status | Drive state | Description and next action                                                                                                                                                                                     |
|--------------|-------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ok           | mounted     | IBM Spectrum Archive EE can access this tape by using the file system interface, but no task is assigned to this drive, and a tape is in this drive.<br><br>Next Action: None                                   |
| ok           | standby     | IBM Spectrum Archive EE cannot access a tape by using the file system interface, but the tape is in this drive. No task is assigned to this drive.<br><br>Next Action: None                                     |
| ok           | mounting    | IBM Spectrum Archive EE is opening a file system interface to access a tape on this drive because of an assigned task.<br><br>Next Action: None                                                                 |
| ok           | in_use      | IBM Spectrum Archive EE is using this drive for a task.<br><br>Next Action: None                                                                                                                                |
| ok           | not_mounted | This drive does not include a tape.<br><br>Next Action: None                                                                                                                                                    |
| ok           | unmounting  | IBM Spectrum Archive EE is closing a file system interface access to the tape in this drive because of opening a file system access to another tape, or moving a tape to the homeslot.<br><br>Next Action: None |

| Drive status | Drive state      | Description and next action                                                                                                                                                                                                                         |
|--------------|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| info         | locked           | IBM Spectrum Archive EE is processing a tape in this drive that has a need_unlock status. The need_unlock state is resolved automatically by the IBM Spectrum Archive EE system.<br><br>Next Action: Wait until the state changes to another state. |
| info         | check_node       | A node that manages a drive is down.<br><br>Next Action: Check the nodes in IBM Spectrum Archive EE by using the <b>eeadm node list</b> command, and make the nodes available by using the <b>eeadm node up</b> command.                            |
| info         | unassigned       | This drive is not a resource of IBM Spectrum Archive EE.<br><br>Next Action: Use the <b>eeadm drive assign</b> command to include this drive as a resource of IBM Spectrum Archive EE.                                                              |
| info         | disabled         | The drive is assigned to a node. But the drive is not open and IBM Spectrum Archive EE cannot use this drive.<br><br>Next Action: Use the <b>eeadm drive up</b> command to enable this drive. The drive can be used if the command is successful.   |
| error        | check_tape_drive | The tape drive reports a hardware error.<br><br>Next Action: Check the drive error messages on the library operator panel or web GUI.                                                                                                               |
| error        | not_installed    | The library is configured, but the drive does not exist as part of the library.<br><br>Next Action: Check the drive error messages on the library operator panel or web GUI.                                                                        |
| error        | missing          | The drive is defined as a resource of IBM Spectrum Archive EE. But it is not found in the library.<br><br>Next Action: There is no degradation. Use the <b>eeadm drive unassign</b> command to clean up this entry.                                 |
| error        | check_drive_path | IBM Spectrum Archive EE cannot find any alternative path at path fail over.<br><br>Next Action: Check drive connectivity on the assigned node, and restore connectivity by using the <b>eeadm drive up</b> command.                                 |

The **eeadm file state** command displays the data placement of files. Each file is in a unique state, as listed in Table 10-2.

Table 10-2 Status codes for the *eeadm file state* 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 content of the file has not been truncated.                                                                                                                                                                             |
| 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.                                                                                                                                                                                              |
| missing     | The missing status indicates that a file has the migrated or premigrated 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 missing because of tape corruption or if the tape cartridge was removed from the system without exporting. |

### 10.1.3 Node status codes

The **eeadm node list** command lists the configuration and status of all the configured nodes, as listed in Table 10-3.

Table 10-3 Status codes for the *eeadm node list* 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 <b>ltfsee_config -m ADD_NODE</b> 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.                                                                                                                                                |

## 10.1.4 Tape status codes

The `eeadm tape list` command lists the configuration and status of all of the tapes, as listed in Table 10-4.

Table 10-4 *Tape status along with a description*

| New Status         | Old Status                        | Description                                                                                                                                                                                                               |
|--------------------|-----------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| appendable         | Valid                             | Tape is available and data can be appended.                                                                                                                                                                               |
| append_fenced      | Valid                             | Pool setting caused this tape to be fenced from appending data.                                                                                                                                                           |
| check_hba          | N/A                               | Temporary flag because of HBA-related error on mount, unmount, load, or unload of the tape to or from the tape drive.                                                                                                     |
| check_key_server   | N/A                               | Temporary flag because of encryption-related error on mount, unmount, load, or unload of the tape to or from the tape drive.                                                                                              |
| check_tape_library | Unusable                          | Temporary flag because of failed mount, unmount, load, or unload of the tape to or from the tape drive.                                                                                                                   |
| data_full          | Valid                             | The tape is almost full (only metadata changes are accepted).                                                                                                                                                             |
| disconnected       | Disconnect                        | The Linear Tape File System (LTFS) LE instance that has this tape is not running.                                                                                                                                         |
| duplicated         | Duplicate                         | Two or more tapes have the same barcode label in the logical library.                                                                                                                                                     |
| exported           | Exported                          | The tape was exported by a normal export.                                                                                                                                                                                 |
| full               | Valid                             | The tape is full and no other changes or updates can be made.                                                                                                                                                             |
| inaccessible       | Inaccessible                      | The tape library reported that the tape is inaccessible.                                                                                                                                                                  |
| label_mismatch     | N/A                               | The barcode and label data on tape are mismatched.                                                                                                                                                                        |
| missing            | Missing                           | The tape is missing from the logical library.                                                                                                                                                                             |
| need_replace       | Warning                           | An read error occurred on the tape.                                                                                                                                                                                       |
| need_unlock        | Critical                          | This is an intermittent state while the IBM Spectrum Archive Enterprise Edition system is processing a permanent write error. This state is automatically resolved by the IBM Spectrum Archive Enterprise Edition system. |
| non_supported      | Not supported                     | The tape is an unsupported media cartridge and should be removed from the logical library.                                                                                                                                |
| offline            | Offline                           | The tape was exported offline.                                                                                                                                                                                            |
| recall_only        | Write Protected                   | The tape is physically write-protected or write-protected by an advisory lock.                                                                                                                                            |
| require_replace    | Critical or Error or Write Fenced | The tape is in a read-only state because of a write error.                                                                                                                                                                |
| require_validate   | Invalid                           | The tape is invalid from an LTFS LE perspective.                                                                                                                                                                          |
| require_validate   | Unknown                           | The tape status is unknown and a tape validation is needed.                                                                                                                                                               |
| unassigned         | Unavailable                       | The tape is not assigned to any pool.                                                                                                                                                                                     |
| unformatted        | Unformatted                       | The tape is in a pool, but is not LTFS-formatted.                                                                                                                                                                         |

Table 10-5 lists the next available commands for each tape status.

Table 10-5 Next available actions for each tape status

| Status                        | migrate | pre migrate | recall | save | tape assign | tape unassign | tape export | tape import | tape offline | tape online | tape datamigrate | tape move | tape reclaim | tape reconcile | tape replace | tape validate |
|-------------------------------|---------|-------------|--------|------|-------------|---------------|-------------|-------------|--------------|-------------|------------------|-----------|--------------|----------------|--------------|---------------|
| appendable                    | x       | x           | x      | x    |             | x             | x           |             | x            |             | x                | x         | x            | x              | x            | x             |
| append_fenced                 |         |             | x      |      |             | x             | x           |             | x            |             | x                | x         | x            | x              |              | x             |
| check_hba                     |         |             |        |      |             |               |             |             |              |             |                  | x         |              |                |              | x             |
| check_key_server              |         |             |        |      |             |               |             |             |              |             |                  | x         |              |                |              | x             |
| check_tape_library            |         |             |        |      |             |               |             |             |              |             |                  | x         |              |                |              | x             |
| data_full                     |         |             | x      |      |             | x             | x           |             | x            |             | x                | x         | x            | x              |              | x             |
| disconnected                  |         |             |        |      |             |               |             |             |              |             |                  |           |              |                |              |               |
| duplicated                    |         |             |        |      |             |               |             |             |              |             |                  |           |              |                |              |               |
| exported                      |         |             |        |      |             |               |             | x           |              |             |                  | x         |              |                |              |               |
| full                          |         |             | x      |      |             | x             |             |             |              |             | x                | x         | x            | x              |              | x             |
| inaccessible                  |         |             |        |      |             |               |             |             |              |             |                  |           |              |                |              |               |
| label_mismatch                |         |             |        |      |             |               |             |             |              |             |                  | x         |              |                |              |               |
| missing                       |         |             |        |      |             |               |             |             |              |             |                  |           |              |                |              |               |
| need_replace                  |         |             | x      |      |             | x             |             |             |              |             |                  | x         |              |                | x            | x             |
| need_unlock (N/A next action) |         |             |        |      |             |               |             |             |              |             |                  |           |              |                |              |               |
| non_supported                 |         |             |        |      |             |               |             |             |              |             |                  | x         |              |                |              |               |
| offline                       |         |             |        |      |             |               |             |             | x            |             |                  | x         |              |                |              |               |
| recall_only                   |         |             | x      |      |             | x             |             |             |              |             | x                | x         | x            | x              |              | x             |
| require_replace               |         |             | x      |      |             | x             |             |             |              |             |                  | x         |              |                | x            | x             |
| require_validate              |         |             | x      |      |             |               |             |             |              |             |                  | x         |              |                |              | x             |
| unassigned                    |         |             |        |      | x           |               |             | x           |              |             |                  | x         |              |                |              |               |
| unformatted                   |         |             |        |      |             |               |             |             |              |             |                  | x         |              |                |              |               |

### Require\_validate state

A tape cartridge in the state is a temporary condition that can be caused when the tape within the pool becomes unknown to the system, or the metadata on the gpfs file system becomes inconsistent with the data cartridge.

### **Recall\_only state**

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, and if the write protection switch is on, 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 **eeadm tape validate** command to update the status of the tape to `appendable`.

### **Require\_replace state**

This status is caused by physical errors on writing data into the tape. The tape becomes a read-only tape and must be replaced as soon as possible. For more information about recovery procedures, see 9.3, “Recovering data from a write failure tape” on page 301 and 9.4, “Recovering data from a read failure tape” on page 302.

### **Need\_replace state**

This status is caused when a tape drive encounters an error trying to read data off the tape. This tape also must be replaced as soon as possible to mitigate future failures. For more information about recovery procedures, see “Recovering data from a read failure tape” on page 302.

### **Unassigned state**

This status is caused by a tape cartridge being removed from tape pool. The process of adding it to LTFS (see 6.8.1, “Adding tape cartridges” on page 154) changes the status back to `appendable`. Therefore, this message requires no other corrective action.

### **Unformatted status**

This status usually is observed when a scratch tape is added to LTFS without formatting it. It can be fixed by removing by **eeadm tape unassign** command and formatting it with the **eeadm tape assign** command, as described in 6.8.3, “Formatting tape cartridges” on page 158.

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 10.3.2, “Using the IBM LTFS Format Verifier” on page 328.

### **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 described in 6.8.2, “Moving tape cartridges” on page 156) should correct the `Inaccessible` status.

### **Non-supported status**

Only LTO-9, 8, M8, 7, 6, 5, and 3592 JB, JC, JD, JE, JK, JL, JM, JV, JY, and JZ tape cartridges are supported by IBM Spectrum Archive EE. This “tape state” or “status” indicates that the tape cartridge is not one of these types and must be removed from the tape library.

## 10.1.5 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 the following IBM Documentation web pages:

- ▶ [IBM Spectrum Scale](#)
- ▶ [GPFS](#)

For more information about the `mmapplypolicy` command and 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 323.

The GPFS `mmgetstate -a` command displays the nodes where GPFS is active.

## 10.1.6 IBM Spectrum Protect for Space Management commands

You can use a subset of IBM Spectrum Protect for Space Management commands 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 at this [IBM Documentation web page](#).

**Important:** The IBM Spectrum Archive EE license does not entitle customers to use any IBM Spectrum Protect components or products other than IBM Spectrum Protect for Space Management from the IBM Spectrum Protect family to migrate data to LTFS.

## Compatibility of IBM Spectrum Protect for Space Management commands with IBM Spectrum Archive EE

Only a subset of IBM Spectrum Protect for Space Management commands is compatible with the IBM Spectrum Archive EE environment. Use only compatible IBM Spectrum Protect 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 IBM Spectrum Protect for Space Management commands are compatible with the IBM Spectrum Archive EE environment:

- ▶ `dsmmigfs`
- ▶ `dsm\s`
- ▶ `dsmrecall`

### 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 HSM daemons are started with the same environment as the `dsmwatchd` watch daemon. Therefore, the `dsm.opt` and `dsm.sys` options files in the IBM Spectrum Protect 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 6.2.1, “IBM Spectrum Scale” on page 136.

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`

## 10.2 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 10-1 shows an example of original GPFS data.

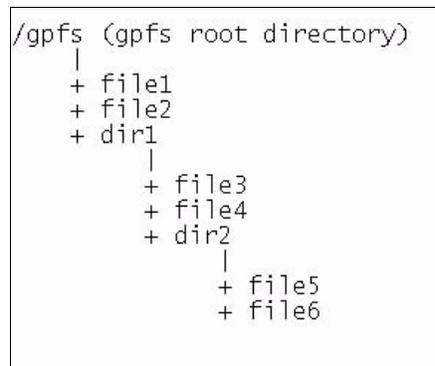


Figure 10-1 Example GPFS data directory structure

► Migrated data and symbolic links in LTFS

Figure 10-2 shows the LTFS data format of tapeA and tapeB after the following migrations of the GPFS files from Figure 10-1:

- file1, file2, and file3 to tapeA.
- file1 (which is a replica), file4, file5, and file6 to tapeB.

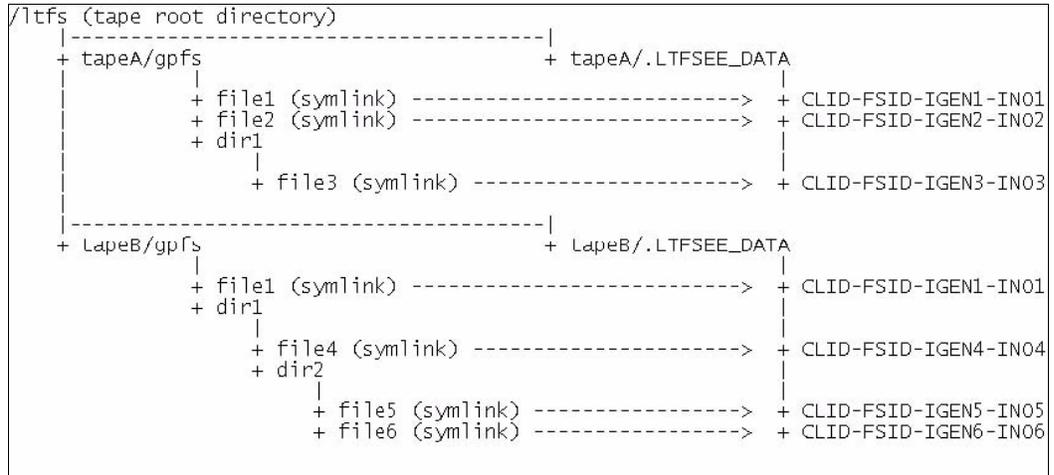


Figure 10-2 Example layout of migrated data and symbolic links on LTFS (on tapes)

► 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 10-2, 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 10-2, 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.

## 10.3 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).

### 10.3.1 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 **Installed version**.
7. Select your operating system from the **Platform** menu.
8. Click **Continue**.
9. (Optional) Narrow the search of available downloads according to your criteria.
10. Click **Continue** to view the list of available downloads.
11. Select the version that you want to download.
12. To download the new version, follow the instructions on the Fix Central download page.

## 10.3.2 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 **Installed version**.
7. Select your operating system from the **Platform** menu.
8. Click **Continue**.
9. Narrow the search of available downloads according to your criteria (this step can be skipped).
10. Click **Continue** to view the list of available downloads.
11. Select the version that you want to download.
12. 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 [IBM Fix Central website](#).

### 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.

## 10.4 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 at this [IBM Documentation web page](#).

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 Documentation:

```
rpm -e ltfs-license-2.1.0-[revision]
```

# 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 Tape Library Guide for Open Systems*, SG24-5946
- ▶ *IBM TS4500 R8 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

*IBM Tape Device Drivers Installation and User's Guide*, GC27-2130, also is relevant as further information source.

## Online resources

The following websites are also relevant as further information sources:

**Tip:** To learn more about IBM Spectrum Archive EE and try it in a virtual tape library environment, see [IBM Spectrum Archive Enterprise Edition Fundamentals & Lab Access](#) (log in required).

- ▶ IBM Product Support: TS3310:  
<https://www.ibm.com/support/pages/node/656001>
- ▶ IBM Documentation: TS4300:  
<https://www.ibm.com/docs/en/ts4300-tape-library>
- ▶ IBM Documentation: TS4500:  
<https://www.ibm.com/docs/en/ts4500-tape-library>
- ▶ IBM Spectrum Archive Enterprise Edition at IBM Documentation:  
<https://www.ibm.com/docs/en/spectrum-archive-ee>

- ▶ IBM Spectrum Archive Library Edition at IBM Documentation:  
<https://www.ibm.com/docs/en/spectrum-archive-1e>
- ▶ IBM Spectrum Archive Enterprise Edition Support:  
[https://www.ibm.com/support/home/product/5449353/IBM\\_Spectrum\\_Archive\\_Enterprise\\_Edition\\_\(EE\)](https://www.ibm.com/support/home/product/5449353/IBM_Spectrum_Archive_Enterprise_Edition_(EE))
- ▶ IBM Spectrum Scale at IBM Documentation:  
<https://www.ibm.com/docs/en/spectrum-scale>
- ▶ IBM Systems Storage interactive product guide:  
<https://www.ibm.com/downloads/cas/AXYWNN5Z>
- ▶ Linear Tape-Open Program:  
<https://www.lto.org/lto-generation-compatibility/>
- ▶ 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 V1.3.2.2: Installation and Configuration Guide**

SG24-8333-09

ISBN 0738460427



(0.5" spine)

0.475" x 0.873"

250 <-> 459 pages







SG24-8333-09

ISBN 0738460427

Printed in U.S.A.

Get connected

